



Reference Manual Volume 3: Procedures

**Adaptive Server Enterprise
12.5**

DOCUMENT ID: 36273-01-1250-01

LAST REVISED: June 2001

Copyright © 1989-2001 by Sybase, Inc. All rights reserved.

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

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

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase, the Sybase logo, ADA Workbench, Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Anywhere, Adaptive Server Enterprise, Adaptive Server Enterprise Monitor, Adaptive Server Enterprise Replication, Adaptive Server Everywhere, Adaptive Server IQ, Adaptive Warehouse, AnswerBase, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-FORMS, APT-Translator, APT-Library, Backup Server, ClearConnect, Client-Library, Client Services, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DB-Library, dbQueue, Developers Workbench, Direct Connect Anywhere, DirectConnect, Distribution Director, E-Anywhere, E-Whatever, Embedded SQL, EMS, Enterprise Application Studio, Enterprise Client/Server, Enterprise Connect, Enterprise Data Studio, Enterprise Manager, Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, EWA, Financial Fusion, Financial Fusion Server, Gateway Manager, ImpactNow, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, KnowledgeBase, MainframeConnect, Maintenance Express, MAP, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, MySupport, Net-Gateway, 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.stop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, PowerJ, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Report Workbench, Report-Execute, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Resource Manager, RW-DisplayLib, RW-Library, S-Designor, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILS, smart.partners, smart.parts, smart.script, SQL Advantage, SQL Anywhere, SQL Anywhere Studio, SQL Code Checker, SQL Debug, SQL Edit, SQL Edit/TPU, SQL Everywhere, SQL Modeler, SQL Remote, SQL Server, SQL Server Manager, SQL SMART, SQL Toolset, SQL Server/CFT, SQL Server/DBM, SQL Server SNMP SubAgent, SQL Station, SQLJ, STEP, SupportNow, Sybase Central, Sybase Client/Server Interfaces, Sybase Financial Server, Sybase Gateways, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, Transact-SQL, Translation Toolkit, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, Viewer, Visual Components, VisualSpeller, VisualWriter, VQL, WarehouseArchitect, Warehouse Control Center, Warehouse Studio, Warehouse WORKS, Watcom, Watcom SQL, Watcom SQL Server, Web Deployment Kit, Web.PB, Web.SQL, WebSights, WebViewer, WorkGroup SQL Server, XA-Library, XA-Server and XP Server are trademarks of Sybase, Inc. 3/01

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

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

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

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

Contents

About This Book	xi
CHAPTER 1	System Procedures 1
	Introduction to system procedures 1
	Permissions on system procedures 1
	Executing system procedures 2
	Entering parameter values 3
	Messages 4
	System procedure tables 4
	List of system procedures 4
CHAPTER 2	Procedures: sp_activeroles 15
	sp_activeroles 15
CHAPTER 3	Procedures: sp_addalias – sp_add_resource_limit 17
	sp_addalias 17
	sp_addauditrecord 18
	sp_addaudittable 20
	sp_addengine 21
	sp_addexeclasse 22
	sp_addextendedproc 23
	sp_addexternlogin 24
	sp_addgroup 26
	sp_addlanguage 27
	sp_addlogin 30
	sp_addmessage 32
	sp_addobjectdef 34
	sp_add_qpgroup 37
	sp_addremotelogin 38
	sp_add_resource_limit 40
CHAPTER 4	Procedures: sp_addsegment – sp_adduser 47

	sp_addsegment.....	47
	sp_addserver	48
	sp_addthreshold.....	51
	sp_add_time_range	55
	sp_addtype.....	58
	sp_addumpdevice	62
	sp_adduser	64
CHAPTER 5	Procedures: sp_altermessage – sp_autoconnect.....	67
	sp_altermessage.....	67
	sp_audit.....	68
	sp_autoconnect.....	72
CHAPTER 6	Procedures: sp_bindcache – sp_bindrule	75
	sp_bindcache	75
	sp_bindefault.....	78
	sp_bindexclass.....	80
	sp_bindmsg.....	83
	sp_bindrule.....	84
CHAPTER 7	Procedures: sp_cacheconfig – sp_cursorinfo.....	87
	sp_cacheconfig	87
	sp_cachestrategy	95
	sp_changedbowner.....	98
	sp_changegroup	99
	sp_checknames	100
	sp_checkreswords	101
	sp_checksourc.....	114
	sp_chgattribute.....	116
	sp_clearpsex.....	120
	sp_clearstats	121
	sp_cmp_all_qplans	122
	sp_cmp_qplans	124
	sp_commonkey	125
	sp_companion.....	127
	sp_configure.....	130
	sp_copy_all_qplans.....	135
	sp_copy_qplan	136
	sp_countmetadata.....	136
	sp_cursorinfo.....	138
CHAPTER 8	Procedures: sp_dboption – sp_displayroles	143

	sp_dboption.....	143
	sp_dbrecovery_order.....	150
	sp_dbremap.....	152
	sp_defaultloc.....	153
	sp_depends.....	156
	sp_deviceattr.....	162
	sp_diskdefault.....	163
	sp_displayaudit.....	164
	sp_displaylevel.....	168
	sp_displaylogin.....	169
	sp_displayroles.....	171
CHAPTER 9	Procedures: sp_dropalias – sp_dropmessage.....	175
	sp_dropalias.....	175
	sp_drop_all_qplans.....	175
	sp_dropdevice.....	176
	sp_dropengine.....	177
	sp_dropexeclass.....	177
	sp_dropextendedproc.....	178
	sp_dropexternlogin.....	179
	sp_dropglockpromote.....	180
	sp_dropgroup.....	180
	sp_dropkey.....	181
	sp_droplanguage.....	182
	sp_droplogin.....	183
	sp_dropmessage.....	184
CHAPTER 10	Procedures: sp_dropobjectdef – sp_drop_qplan.....	187
	sp_dropobjectdef.....	187
	sp_drop_qpgroup.....	188
	sp_drop_qplan.....	189
CHAPTER 11	Procedures: sp_dropremotelogin – sp_dropuser.....	191
	sp_dropremotelogin.....	191
	sp_drop_resource_limit.....	192
	sp_droprowlockpromote.....	195
	sp_dropsegment.....	196
	sp_dropserver.....	197
	sp_dropthreshold.....	198
	sp_drop_time_range.....	199
	sp_droptype.....	200
	sp_dropuser.....	200

CHAPTER 12	Procedures: sp_dumpoptimize	203
	sp_dumpoptimize	203
CHAPTER 13	Procedures: sp_estspace – sp_grantlogin	209
	sp_estspace	209
	sp_export_qpgroup	212
	sp_extendsegment	213
	sp_extengine	214
	sp_familylock	215
	sp_find_qplan	217
	sp_flushstats	219
	sp_forceonline_db	219
	sp_forceonline_object	221
	sp_forceonline_page	222
	sp_foreignkey	224
	sp_freedll	225
	sp_getmessage	226
	sp_grantlogin	227
CHAPTER 14	Procedures: sp_ha_admin	229
	sp_ha_admin	229
CHAPTER 15	Procedures: sp_help – sp_helpindex	231
	sp_help	231
	sp_helppartition	237
	sp_helpcache	239
	sp_helpconfig	240
	sp_helpconstraint	244
	sp_helppdb	248
	sp_helpdevice	251
	sp_helpextendedproc	252
	sp_helpexternlogin	253
	sp_helpgroup	254
	sp_helpindex	255
CHAPTER 16	Procedures: sp_helpjava – sp_helprotect	259
	sp_helpjava	259
	sp_helpjoins	261
	sp_helpkey	262
	sp_helplanguage	264
	sp_helplog	265
	sp_helpobjectdef	265

	sp_help_qpgroup	266
	sp_help_qplan	268
	sp_helpremotelogin	270
	sp_help_resource_limit	270
	sp_helprotect.....	273
CHAPTER 17	Procedures: sp_helpsegment – sp_helpuser.....	279
	sp_helpsegment.....	279
	sp_helpserver.....	281
	sp_helpsort.....	282
	sp_helptext.....	283
	sp_helpthreshold.....	285
	sp_helpuser.....	285
CHAPTER 18	Procedures: sp_hidetext	287
	sp_hidetext.....	287
CHAPTER 19	Procedures: sp_import_qpgroup – sp_logiosize.....	289
	sp_import_qpgroup	289
	sp_indsuspect	290
	sp_listsuspect_db.....	291
	sp_listsuspect_object.....	291
	sp_listsuspect_page.....	292
	sp_lock	293
	sp_locklogin	297
	sp_logdevice	298
	sp_loginconfig	300
	sp_logininfo	301
	sp_logiosize	302
CHAPTER 20	Procedures: sp_modifylogin – sp_object_stats	307
	sp_modifylogin	307
	sp_modify_resource_limit	310
	sp_modify_time_range.....	312
	sp_modifythreshold	314
	sp_monitor	318
	sp_monitorconfig.....	320
	sp_object_stats	324
CHAPTER 21	Procedures: sp_passthru – sp_procxmode	329
	sp_passthru.....	329

sp_password	331
sp_placeobject	332
sp_plan_dbccdb	333
sp_poolconfig	336
sp_primarykey	341
sp_processmail	342
sp_procqmode	344
sp_procxmode	346

CHAPTER 22	Procedures: sp_recompile – sp_role	349
	sp_recompile	349
	sp_remap	350
	sp_remoteoption	351
	sp_remotesql	353
	sp_rename	355
	sp_renamedb	356
	sp_rename_qpgroup	359
	sp_reportstats	359
	sp_revokelgin	361
	sp_role	361

CHAPTER 23	Procedures: sp_sendmsg – sp_transactions	363
	sp_sendmsg	363
	sp_serveroption	365
	sp_setlangalias	368
	sp_setpglockpromote	369
	sp_setpsexex	371
	sp_set_qplan	372
	sp_setrowlockpromote	373
	sp_setsuspect_granularity	376
	sp_setsuspect_threshold	378
	sp_showcontrolinfo	379
	sp_showexeclax	381
	sp_showplan	382
	sp_showpsexex	383
	sp_spaceused	385
	sp_ssladmin	387
	sp_syntax	389
	sp_sysmon	390
	sp_thresholdaction	393
	sp_transactions	395

CHAPTER 24	Procedures: sp_unbindcache – sp_who	403
	sp_unbindcache	403
	sp_unbindcache_all	405
	sp_unbindefault	405
	sp_unbindexeclass	407
	sp_unbindmsg	408
	sp_unbindrule	409
	sp_volchanged	410
	sp_who	413
CHAPTER 25	Catalog Stored Procedures	417
	Overview	417
	Specifying optional parameters	418
	Pattern matching	419
	System procedure tables	419
	ODBC datatypes	420
	sp_column_privileges	420
	sp_columns	422
	sp_databases	424
	sp_datatype_info	425
	sp_fkeys	426
	sp_pkeys	428
	sp_server_info	429
	sp_special_columns	432
	sp_sproc_columns	433
	sp_statistics	435
	sp_stored_procedures	436
	sp_table_privileges	437
	sp_tables	439
CHAPTER 26	System Extended Stored Procedures	441
	Permissions on system ESPs	441
	DLLs associated with system ESPs	442
	Using system ESPs	442
	xp_cmdshell	442
	xp_deletemail	444
	xp_enumgroups	444
	xp_findnextmsg	445
	xp_logevent	446
	xp_readmail	447
	xp_sendmail	450
	xp_startmail	453
	xp_stopmail	454

CHAPTER 27	dbcc Stored Procedures	457
	Specifying the Object Name and Date	458
	Specifying the Object Name	458
	Specifying the Date	458
	sp_dbcc_alterws	459
	sp_dbcc_configreport	460
	sp_dbcc_createws	461
	sp_dbcc_deletedb	463
	sp_dbcc_deletehistory	464
	sp_dbcc_differentialreport	465
	sp_dbcc_evaluatedb	466
	sp_dbcc_faultreport	467
	sp_dbcc_fullreport	469
	sp_dbcc_runcheck	470
	sp_dbcc_statisticsreport	471
	sp_dbcc_summaryreport	474
	sp_dbcc_updateconfig	478
Index		481

About This Book

The Adaptive Server Reference Manual is a four-volume guide to Sybase® Adaptive Server™ Enterprise and the Transact-SQL® language.

Volume 1, “*Building Blocks*,” describes the “parts” of Transact-SQL: datatypes, built-in functions, expressions and identifiers, reserved words, and SQLSTATE errors. Before you can use Transact-SQL successfully, you need to understand what these building blocks do and how they affect the results of Transact-SQL statements.

Volume 2, “*Commands*,” provides reference information about the Transact-SQL commands, which you use to create statements.

Volume 3, “*Procedures*” provides reference information about system procedures, catalog stored procedures, extended stored procedures, and dbcc stored procedures. All procedures are created using Transact-SQL statements.

Volume 4, “*System Tables*,” provides reference information about the system tables, which store information about your server, databases, users, and other details of your server. It provides information about the tables in the dbccdb and dbccalt databases.

Audience

The *Adaptive Server Reference Manual* is intended as a reference tool for Transact-SQL users of all levels.

How to use this book

- Chapter 1, “System Procedures,” lists the Adaptive Server system procedures in a table that provides the name and a brief description. Click on a procedure name in the table to go directly to the procedure.
- Chapter 2 through Chapter 7 provide manual pages for the individual system procedures.
- Chapter 25, “Catalog Stored Procedures,” contains reference pages for Adaptive Server catalog stored procedures.
- Chapter 26, “System Extended Stored Procedures,” contains reference pages for Adaptive Server system extended stored procedures.
- Chapter 27, “dbcc Stored Procedures,” contains reference pages for Adaptive Server dbcc stored procedures.

Related documents

The following documents comprise the Sybase Adaptive Server Enterprise documentation:

- The release bulletin for your platform – contains last-minute information that was too late to be included in the books.

A more recent version of the release bulletin may be available on the World Wide Web. To check for critical product or document information that was added after the release of the product CD, use the Sybase Technical Library.
- The *Installation Guide* for your platform – describes installation, upgrade, and configuration procedures for all Adaptive Server and related Sybase products.
- *Configuring Adaptive Server Enterprise* for your platform – provides instructions for performing specific configuration tasks for Adaptive Server.
- *What's New in Adaptive Server Enterprise?* – describes the new features in Adaptive Server version 12.5, the system changes added to support those features, and the changes that may affect your existing applications.
- *Transact-SQL User's Guide* – documents Transact-SQL, Sybase's enhanced version of the relational database language. This manual serves as a textbook for beginning users of the database management system. This manual also contains descriptions of the pubs2 and pubs3 sample databases.
- *System Administration Guide* – provides in-depth information about administering servers and databases. This manual includes instructions and guidelines for managing physical resources, security, user and system databases, and specifying character conversion, international language, and sort order settings.
- *Reference Manual* – contains detailed information about all Transact-SQL commands, functions, procedures, and datatypes. This manual also contains a list of the Transact-SQL reserved words and definitions of system tables.
- *Performance and Tuning Guide* – explains how to tune Adaptive Server for maximum performance. This manual includes information about database design issues that affect performance, query optimization, how to tune Adaptive Server for very large databases, disk and cache issues, and the effects of locking and cursors on performance.

- The *Utility Guide* – documents the Adaptive Server utility programs, such as `isql` and `bcp`, which are executed at the operating system level.
- The *Quick Reference Guide* – provides a comprehensive listing of the names and syntax for commands, functions, system procedures, extended system procedures, datatypes, and utilities in a pocket-sized book. Available only in print version.
- The *System Tables Diagram* – illustrates system tables and their entity relationships in a poster format. Available only in print version.
- *Error Messages and Troubleshooting Guide* – explains how to resolve frequently occurring error messages and describes solutions to system problems frequently encountered by users.
- *Component Integration Services User's Guide* – explains how to use the Adaptive Server Component Integration Services feature to connect remote Sybase and non-Sybase databases.
- *Java in Adaptive Server Enterprise* – describes how to install and use Java classes as datatypes, functions, and stored procedures in the Adaptive Server database.
- *Using Sybase Failover in a High Availability System* – provides instructions for using Sybase's Failover to configure an Adaptive Server as a companion server in a high availability system.
- *Using Adaptive Server Distributed Transaction Management Features* – explains how to configure, use, and troubleshoot Adaptive Server DTM features in distributed transaction processing environments.
- *EJB Server User's Guide* – explains how to use EJB Server to deploy and execute Enterprise JavaBeans in Adaptive Server.
- *XA Interface Integration Guide for CICS, Encina, and TUXEDO* – provides instructions for using Sybase's DTM XA interface with X/Open XA transaction managers.
- *Glossary* – defines technical terms used in the Adaptive Server documentation.
- *Sybase jConnect for JDBC Programmer's Reference* – describes the jConnect for JDBC product and explains how to use it to access data stored in relational database management systems.
- *Full-Text Search Specialty Data Store User's Guide* – describes how to use the Full-Text Search feature with Verity to search Adaptive Server Enterprise data.

- *Historical Server User's Guide* –describes how to use Historical Server to obtain performance information for SQL Server and Adaptive Server.
- *Monitor Server User's Guide* – describes how to use Monitor Server to obtain performance statistics from SQL Server and Adaptive Server.
- *Monitor Client Library Programmer's Guide* – describes how to write Monitor Client Library applications that access Adaptive Server performance data.

Other sources of information

Use the Sybase Technical Library CD and the Technical Library Product Manuals Web site to learn more about your product:

- Technical Library CD contains product manuals and is included with your software. The DynaText browser (downloadable from Product Manuals at <http://www.sybase.com/detail/1,3693,1010661,00.html>) allows you to access technical information about your product in an easy-to-use format.

Refer to the *Technical Library Installation Guide* in your documentation package for instructions on installing and starting the Technical Library.

- Technical Library Product Manuals Web site is an HTML version of the Technical Library CD that you can access using a standard Web browser. In addition to product manuals, you will find links to the Technical Documents Web site (formerly known as Tech Info Library), the Solved Cases page, and Sybase/Powersoft newsgroups.

To access the Technical Library Product Manuals Web site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

Conventions

The following sections describe conventions used in this manual.

SQL is a free-form language. There are no rules about the number of words you can put on a line or where you must break a line. However, for readability, all examples and most syntax statements in this manual are formatted so that each clause of a statement begins on a new line. Clauses that have more than one part extend to additional lines, which are indented. Complex commands are formatted using modified Backus Naur Form (BNF) notation.

Table 1 shows the conventions for syntax statements that appear in this manual:

Table 1: Font and syntax conventions for this manual

Element	Example
Command names, command options, utility names, utility options, and other keywords are bold.	select sp_configure
Database names, datatypes, file names and path names are in italics.	<i>master database</i>

Element	Example
Variables, or words that stand for values that you fill in, are in italics.	<code>select <i>column_name</i> from <i>table_name</i> where <i>search_conditions</i></code>
Type parentheses as part of the command.	<code>compute <i>row_aggregate</i> (<i>column_name</i>)</code>
Double colon, equals sign indicates that the syntax is written in BNF notation. Do not type this symbol. Indicates “is defined as”.	<code>::=</code>
Curly braces mean that you must choose at least one of the enclosed options. Do not type the braces.	<code>{<i>cash, check, credit</i>}</code>
Brackets mean that to choose one or more of the enclosed options is optional. Do not type the brackets.	<code>[<i>cash check credit</i>]</code>
The comma means you may choose as many of the options shown as you want. Separate your choices with commas as part of the command.	<code><i>cash, check, credit</i></code>
The pipe or vertical bar() means you may select only one of the options shown.	<code><i>cash check credit</i></code>
An ellipsis (...) means that you can <i>repeat</i> the last unit as many times as you like.	<code>buy thing = price [<i>cash check credit</i>] [, thing = price [<i>cash check credit</i>]]...</code> You must buy at least one thing and give its price. You may choose a method of payment: one of the items enclosed in square brackets. You may also choose to buy additional things: as many of them as you like. For each thing you buy, give its name, its price, and (optionally) a method of payment.

- Syntax statements (displaying the syntax and all options for a command) appear as follows:

```
sp_dropdevice [device_name]
```

or, for a command with more options:

```
select column_name  
from table_name  
where search_conditions
```

In syntax statements, keywords (commands) are in normal font and identifiers are in lowercase. Italic font shows user-supplied words.

- Examples showing the use of Transact-SQL commands are printed like this:

```
select * from publishers
```

-
- Examples of output from the computer appear as follows:

pub_id	pub_name	city	state
0736	New Age Books	Boston	MA
0877	Binnet & Hardley	Washington	DC
1389	Algodata Infosystems	Berkeley	CA

(3 rows affected)

In this manual, most of the examples are in lowercase. However, you can disregard case when typing Transact-SQL keywords. For example, **SELECT**, **Select**, and **select** are the same.

Adaptive Server's sensitivity to the case of database objects, such as table names, depends on the sort order installed on Adaptive Server. You can change case sensitivity for single-byte character sets by reconfiguring the Adaptive Server sort order. For more information, see the *System Administration Guide*.

If you need help

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.

System Procedures

This chapter describes the system procedures, which are Sybase-supplied stored procedures used for updating and getting reports from system tables. “List of system procedures” on page 4 lists the system procedures described in this volume.

Introduction to system procedures

System procedures are created by installmaster at installation. They are located in the sybssystemprocs database, and owned by the System Administrator.

Some system procedures can be run only in a specific database, but many of them can be run in any database. You can create your own system procedures that can be executed from any database. For more information, see the System Administration Guide.

All system procedures execute at isolation level 1.

All system procedures report a return status. For example:

```
return status = 0
```

means that the procedure executed successfully. The examples in this book do not include the return status.

Permissions on system procedures

Permissions for system procedures are set in the sybssystemprocs database.

Some system procedures can be run only by Database Owners. These procedures make sure that the user executing the procedure is the owner of the database from which they are being executed.

Other system procedures (for example, all the `sp_help` procedures) can be executed by any user who has been granted permission, provided that the permission was granted in `sybssystemprocs`. A user must have permission to execute a system procedure either in all databases or in none of them.

A user who is not listed in `sybssystemprocs..sysusers` is treated as a “guest” user in `sybssystemprocs` and is automatically granted permission on many of the system procedures.

To deny a user permission on a system procedure, the System Administrator must add the user to `sybssystemprocs..sysusers` and write a `revoke` statement that applies to that procedure. The owner of a user database cannot directly control permissions on the system procedures within his or her own database.

Executing system procedures

If a system procedure is executed in a database other than `sybssystemprocs`, it operates on the system tables in the database in which it was executed. For example, if the Database Owner of `pubs2` runs `sp_adduser` in `pubs2`, the new user is added to `pubs2..sysusers`.

To run a system procedure in a specific database, either:

- Open that database with the `use` command and execute the procedure, or
- Qualify the procedure name with the database name.

For example, the user-defined system procedure `sp_foo`, which executes the `db_name()` system function, returns the name of the database in which it is executed. When executed in the `pubs2` database, it returns the value “`pubs2`”:

```
exec pubs2..sp_foo
-----
pubs2
(1 row affected, return status = 0)
```

When executed in `sybssystemprocs`, it returns the value “`sybssystemprocs`”:

```
exec sybssystemprocs..sp_foo
-----
sybssystemprocs
(1 row affected, return status = 0)
```

Entering parameter values

If a parameter value for a system procedure contains punctuation or embedded blanks, or is a reserved word, you must enclose it in single or double quotes. If the parameter is an object name qualified by a database name or owner name, enclose the entire name in single or double quotes.

Note

Do not use delimited identifiers as system procedure parameters; they may produce unexpected results.

If a procedure has multiple optional parameters, you can supply parameters in the form:

```
@parametername = value
```

instead of supplying all the parameters. The parameter names in the syntax statements match the parameter names defined by the procedures.

For example, the syntax for `sp_addlogin` is:

```
sp_addlogin login_name, password [, defdb  
[, deflanguage [, fullname]]]
```

To use `sp_addlogin` to create a login for “susan” with a password of “wonderful”, a full name of Susan B. Anthony, and the server’s default database and language, you can use:

```
sp_addlogin susan, wonderful,  
@fullname="Susan B. Anthony"
```

This provides the same information as the command with all the parameters specified:

```
sp_addlogin susan, wonderful, public_db,  
us_english, "Susan B. Anthony"
```

You can also use “null” as a placeholder:

```
sp_addlogin susan, wonderful, null, null,  
"Susan B. Anthony"
```

Do not enclose “null” in quotes.

SQL has no rules about the number of words you can put on a line or where you must break a line. If you issue a system procedure followed by a command, Adaptive Server attempts to execute the system procedure, then the command. For example, if you execute the command:

```
sp_help checkpoint
```

Adaptive Server returns the output from `sp_help`, then runs the checkpoint command.

If you specify more parameters than the number of parameters expected by the system procedure, the extra parameters are ignored by Adaptive Server.

Messages

System procedures return informational and error messages, which are listed with each procedure in this book. System procedure error messages start at error number 17000.

Error messages from the functions and commands included in a procedure are documented in *Troubleshooting and Error Messages Guide*.

System procedure tables

Several *system procedure tables* in the master database, such as `spt_values`, `spt_committab`, `spt_monitor`, and `spt_limit_types`, are used by system procedures to convert internal system values (for example, status bits) into human-readable format.

`spt_values` is never updated. To see how it is used, execute `sp_helptext` to look at the text for one of the system procedures that references it.

In addition, some system procedures create and then drop temporary tables.

List of system procedures

Table 1-1 provides a brief description of each system procedure.

Table 1-1: System procedures

Procedure	Description
<code>sp_activeroles</code> on page 15	Displays all active roles granted to a user's login.

Procedure	Description
sp_addalias on page 17	Allows an Adaptive Server user to be known in a database as another user.
sp_addauditrecord on page 18	Allows users to enter user-defined audit records (comments) into the audit trail.
sp_addauditable on page 20	Adds another system audit table after auditing is installed.
sp_addengine on page 21	Adds an engine to an existing engine group or, if the group does not exist, creates an engine group and adds the engine.
sp_addexclass on page 22	Creates or updates a user-defined execution class that you can bind to client applications, logins, and stored procedures.
sp_addextendedproc on page 23	Creates an extended stored procedure (ESP) in the master database.
sp_addexternlogin on page 24	Creates an alternate login account and password to use when communicating with a remote server through Component Integration Services.
sp_addgroup on page 26	Adds a group to a database. Groups are used as collective names in granting and revoking privileges.
sp_addlanguage on page 27	Defines the names of the months and days, and the date format, for an alternate language.
sp_addlogin on page 30	Adds a new user account to Adaptive Server.
sp_addmessage on page 32	Adds user-defined messages to sysusermessages for use by stored procedure print and raiserror calls and by sp_bindmsg.
sp_addobjectdef on page 34	Specifies the mapping between a local table and an external storage location.
sp_add_qpgroup on page 37	Adds an abstract plan group.
sp_addremotelogin on page 38	Authorizes a new remote server user by adding an entry to master.dbo.sysremotelogins.
sp_add_resource_limit on page 40	Creates a limit on the amount of server resources that a login or application can use to execute a query, query batch, or transaction.
sp_addsegment on page 47	Defines a segment on a database device in the current database.
sp_addserver on page 48	Defines a remote server or defines the name of the local server.
sp_addthreshold on page 51	Creates a threshold to monitor space on a database segment. When free space on the segment falls below the specified level, Adaptive Server executes the associated stored procedure.
sp_add_time_range on page 55	Adds a named time range to Adaptive Server.
sp_addtype on page 58	Creates a user-defined datatype.
sp_addumpdevice on page 62	Adds a dump device to Adaptive Server.
sp_adduser on page 64	Adds a new user to the current database.
sp_altermessage on page 67	Enables and disables the logging of a specific system-defined or user-defined message in the Adaptive Server error log.
sp_audit on page 68	Allows a System Security Officer to configure auditing options.

Procedure	Description
sp_autoconnect on page 72	Defines a passthrough connection to a remote server for a specific user, which allows the named user to enter passthrough mode automatically at login.
sp_bindcache on page 75	Binds a database, table, index, text object, or image object to a data cache.
sp_bindefault on page 78	Binds a user-defined default to a column or user-defined datatype.
sp_bindexclass on page 80	Associates an execution class with a client application, login, or stored procedure.
sp_bindmsg on page 83	Binds a user message to a referential integrity constraint or check constraint.
sp_bindrule on page 84	Binds a rule to a column or user-defined datatype.
sp_cacheconfig on page 87	Creates, configures, reconfigures, drops, and provides information about data caches.
sp_cachestrategy on page 95	Enables or disables prefetching (large I/O) and MRU cache replacement strategy for a table, index, text object, or image object.
sp_changedbowner on page 98	Changes the owner of a database.
sp_changegroup on page 99	Changes a user's group.
sp_checknames on page 100	Checks the current database for names that contain characters not in the 7-bit ASCII set.
sp_checkreswords on page 101	Detects and displays identifiers that are Transact-SQL reserved words. Checks server names, device names, database names, segment names, user-defined datatypes, object names, column names, user names, login names, and remote login names.
sp_checksouce on page 114	Checks for the existence of the source text of the compiled object .
sp_chgattribute on page 116	Changes the max_rows_per_page value for future space allocations of a table or index.
sp_clearpsexex on page 120	Clears the execution attributes of the client application, login, or stored procedure that was set by sp_setpsexex.
sp_clearstats on page 121	Initiates a new accounting period for all server users or for a specified user. Prints statistics for the previous period by executing sp_reportstats.
sp_cmp_all_qplans on page 122	Compares all abstract plans in two abstract plan groups.
sp_cmp_qplans on page 124	Compares two abstract plans.
sp_commonkey on page 125	Defines a common key—columns that are frequently joined—between two tables or views.
sp_companion on page 127	Performs cluster operations such as configuring Adaptive Server as a secondary companion in a high availability system and moving a companion server from one failover mode to another
sp_configure on page 130	Displays or changes configuration parameters.
sp_copy_all_qplans on page 135	Copies all plans for one abstract plan group to another group.

Procedure	Description
sp_copy_all_qplans on page 135	Copies one abstract plan to an abstract plan group.
sp_countmetadata on page 136	Displays the number of indexes, objects, or databases in Adaptive Server.
sp_cursorinfo on page 138	Reports information about a specific cursor or all cursors that are active for your session.
sp_dboption on page 143	Displays or changes database options.
sp_dbrecovery_order on page 150	Specifies the order in which user databases are recovered and lists the user-defined recovery order of a database or all databases.
sp_dbremap on page 152	Forces Adaptive Server to recognize changes made by alter database. Run this procedure only when instructed to do so by an Adaptive Server message.
sp_depends on page 156	Displays information about database object dependencies—the view(s), trigger(s), and procedure(s) that depend on a specified table or view, and the table(s) and view(s) that the specified view, trigger, or procedure depends on.
sp_deviceattr on page 162	Changes the dsync setting of an existing database device file.
sp_diskdefault on page 163	Specifies whether or not a database device can be used for database storage if the user does not specify a database device or specifies default with the create database or alter database commands.
sp_displayaudit on page 164	Displays the status of audit options.
sp_displaylevel on page 168	Sets or shows which Adaptive Server configuration parameters appear in sp_configure output.
sp_displaylogin on page 169	Displays information about a login account.
sp_displayroles on page 171	Displays all roles granted to another role, or displays the entire hierarchy tree of roles in table format.
sp_dropalias on page 175	Removes the alias user name identity established with sp_addalias.
sp_drop_all_qplans on page 175	Deletes all abstract plans in an abstract plan group.
sp_dropdevice on page 176	Drops an Adaptive Server database device or dump device.
sp_dropengine on page 177	Drops an engine from a specified engine group or, if the engine is the last one in the group, drops the engine group.
sp_dropexeclss on page 177	Drops a user-defined execution class.
sp_dropextendedproc on page 178	Removes an ESP from the master database.
sp_droplockpromote on page 180	Removes lock promotion values from a table or database.
sp_dropgroup on page 180	Drops a group from a database.
sp_dropkey on page 181	Removes a key defined with sp_primarykey, sp_foreignkey, or sp_commonkey from the syskeys table.
sp_droplanguage on page 182	Drops an alternate language from the server and removes its row from master.dbo.syslanguages.
sp_droplogin on page 183	Drops an Adaptive Server user login by deleting the user's entry in master.dbo.syslogins.

Procedure	Description
sp_dropmessage on page 184	Drops user-defined messages from sysusermessages.
sp_dropobjectdef on page 187	<i>Component Integration Services only</i> – Deletes the external storage mapping provided for a local object
sp_drop_qpgroup on page 188	Drops an abstract plan group.
sp_drop_qplan on page 189	Drops an abstract plan.
sp_dropremotelogin on page 191	Drops a remote user login.
sp_drop_resource_limit on page 192	Removes one or more resource limits from Adaptive Server.
sp_dropsegment on page 196	Drops a segment from a database or unmaps a segment from a particular database device.
sp_dropserver on page 197	Drops a server from the list of known servers.
sp_dropthreshold on page 198	Removes a free-space threshold from a segment.
sp_drop_time_range on page 199	Removes a user-defined time range from Adaptive Server.
sp_droptype on page 200	Drops a user-defined datatype.
sp_dropuser on page 200	Drops a user from the current database.
sp_dumpoptimize on page 203	Specifies the amount of data dumped by Backup Server during the dump database operation.
sp_estspace on page 209	Estimates the amount of space required for a table and its indexes, and the time needed to create the index.
sp_export_qpgroup on page 212	Exports all plans for a specified user and abstract plan group to a user table.
sp_extendsegment on page 213	Extends the range of a segment to another database device.
sp_familylock on page 215	Reports information about all the locks held by a family (coordinating process and its worker processes) executing a statement in parallel.
sp_find_qplan on page 217	Finds an abstract plan, given a pattern from the query text or plan text.
sp_flushstats on page 219	Flushes statistics from in-memory storage to the systabstats system table.
sp_forceonline_db on page 219	Provides access to all the pages in a database that were previously taken offline by recovery.
sp_forceonline_page on page 222	Provides access to pages previously taken offline by recovery.
sp_foreignkey on page 224	Defines a foreign key on a table or view in the current database.
sp_freedll on page 225	Unloads a dynamic link library (DLL) that was previously loaded into XP Server memory to support the execution of an ESP.
sp_getmessage on page 226	Retrieves stored message strings from sysmessages and sysusermessages for print and raiserror statements.
sp_grantlogin on page 227	<i>Windows NT only</i> – When Integrated Security mode or Mixed mode (with Named Pipes) is active, assigns Adaptive Server roles or default permissions to Windows NT users and groups.

Procedure	Description
sp_ha_admin on page 229	Performs administrative tasks on Adaptive Servers configured with Sybase Failover in a high availability system. sp_ha_admin is installed with the <i>installhavss</i> script (<i>insthasv</i> on Windows NT).
sp_help on page 231	Reports information about a database object (any object listed in sysobjects) and about Adaptive Server-supplied or user-defined datatypes.
sp_helppartition on page 237	Lists the first page and the control page for each partition in a partitioned table.
sp_helpcache on page 239	Displays information about the objects that are bound to a data cache or the amount of overhead required for a specified cache size.
sp_helpdb on page 248	Reports information about a particular database or about all databases.
sp_helpdevice on page 251	Reports information about a particular device or about all Adaptive Server database devices and dump devices.
sp_helpextendedproc on page 252	Displays ESPs registered in the current database, along with their associated DLL files.
sp_helpexternlogin on page 253	<i>Component Integration Services only</i> – Reports information about external login names.
sp_helpgroup on page 254	Reports information about a particular group or about all groups in the current database.
sp_helpindex on page 255	Reports information about the indexes created on a table.
sp_helpjava on page 259	Displays information about Java classes and associated JARs that are installed in the database.
sp_helpjoins on page 261	Lists the columns in two tables or views that are likely join candidates.
sp_helpkey on page 262	Reports information about a primary, foreign, or common key of a particular table or view, or about all keys in the current database.
sp_helplanguage on page 264	Reports information about a particular alternate language or about all languages.
sp_helplog on page 265	Reports the name of the device that contains the first page of the transaction log.
sp_helpobjectdef on page 265	<i>Component Integration Services only</i> – Reports information about remote object definitions. Shows owners, objects, type, and definition.
sp_help_qpgroup on page 266	Reports information on an abstract plan group.
sp_help_qplan on page 268	Reports information about an abstract plan.
sp_helpremotelogin on page 270	Reports information about a particular remote server's logins or about all remote servers' logins.
sp_help_resource_limit on page 270	Reports information about all resource limits, limits for a given login or application, limits in effect at a given time or day of the week, or limits with a given scope or action.
sp_helpprotect on page 273	Reports information about permissions for database objects, users, groups, or roles.

Procedure	Description
sp_helpsegment on page 279	Reports information about a particular segment or about all segments in the current database.
sp_helpserver on page 281	Reports information about a particular remote server or about all remote servers.
sp_helpsort on page 282	Displays Adaptive Server's default sort order and character set.
sp_helptext on page 283	Prints the text of a system procedure, trigger, view, default, rule, or integrity check constraint.
sp_helpthreshold on page 285	Reports the segment, free-space value, status, and stored procedure associated with all thresholds in the current database or all thresholds for a particular segment.
sp_helpuser on page 285	Reports information about a particular user or about all users in the current database.
sp_hidetext on page 287	Hides the source text for the specified compiled object .
sp_import_qpgroup on page 289	Imports abstract plans from a user table into an abstract plan group.
sp_indsuspect on page 290	Checks user tables for indexes marked as suspect during recovery following a sort order change.
sp_listsuspect_db on page 291	Lists all databases that have offline pages because of corruption detected on recovery.
sp_listsuspect_object on page 291	Lists all indexes in a database that are currently offline because of corruption detected on recovery.
sp_listsuspect_page on page 292	Lists all pages that are currently offline because of corruption detected on recovery.
sp_lock on page 293	Reports information about processes that currently hold locks.
sp_locklogin on page 297	Locks an Adaptive Server account so that the user cannot log in, or displays a list of all locked accounts.
sp_logdevice on page 298	Moves the transaction log of a database with log and data on the same device to a separate database device.
sp_loginconfig on page 300	<i>Windows NT only</i> – Displays the value of one or all integrated security parameters.
sp_logininfo on page 301	<i>Windows NT only</i> – Displays all roles granted to Windows NT users and groups with sp_grantlogin.
sp_logiosize on page 302	Changes the log I/O size used by Adaptive Server to a different memory pool when it is doing I/O for the transaction log of the current database.
sp_modifylogin on page 307	Modifies the default database, default language, default role activation, or full name for an Adaptive Server login account.
sp_modify_resource_limit on page 310	Changes a resource limit by specifying a new limit value or the action to take when the limit is exceeded, or both.
sp_modify_time_range on page 312	Changes the start day, start time, end day, and/or end time associated with a named time range.

Procedure	Description
sp_modifythreshold on page 314	Modifies a threshold by associating it with a different threshold procedure, free-space level, or segment name. You <i>cannot</i> use sp_modifythreshold to change the amount of free space or the segment name for the last-chance threshold.
sp_monitor on page 318	Displays statistics about Adaptive Server.
sp_object_stats on page 324	Shows lock contention, lock wait-time, and deadlock statistics for tables and indexes.
sp_passthru on page 329	<i>Component Integration Services only</i> – Allows the user to pass a SQL command buffer to a remote server.
sp_password on page 331	Adds or changes a password for an Adaptive Server login account.
sp_placeobject on page 332	Puts future space allocations for a table or an index on a particular segment.
sp_plan_dbccdb on page 333	Recommends suitable sizes for new dbccdb and dbccalt databases, lists suitable devices for dbccdb and dbccalt, and suggests a cache size and a suitable number of worker processes for the target database.
sp_poolconfig on page 336	Creates, drops, resizes, and provides information about memory pools within data caches.
sp_primarykey on page 341	Defines a primary key on a table or view.
sp_processmail on page 342	<i>Windows NT only</i> – Reads, processes, sends, and deletes messages in the Adaptive Server message inbox.
sp_procqmode on page 344	Displays the query processing mode of a stored procedure, view, or trigger.
sp_procxmode on page 346	Displays or changes the transaction modes associated with stored procedures.
sp_recompile on page 349	Causes each stored procedure and trigger that uses the named table to be recompiled the next time it runs.
sp_remap on page 350	Remaps a stored procedure, trigger, rule, default, or view from releases later than 4.8 and earlier than 10.0 to be compatible with releases 10.0 and later. Use sp_remap on pre-release 11.0 objects that the release 11.0 upgrade procedure failed to remap.
sp_remotoption on page 351	Displays or changes remote login options.
sp_rename on page 355	Changes the name of a user-created object or user-defined datatype in the current database.
sp_renamedb on page 356	Changes the name of a database. You <i>cannot</i> rename system databases or databases with external referential integrity constraints.
sp_rename_qpgroup on page 359	Renames an abstract plan group.
sp_reportstats on page 359	Reports statistics on system usage.
sp_revokelogin on page 361	<i>Windows NT only</i> – When Integrated Security mode or Mixed mode (with Named Pipes) is active, revokes Adaptive Server roles and default permissions from Windows NT users and groups.

Procedure	Description
sp_role on page 361	Grants or revokes system roles to an Adaptive Server login account.
sp_sendmsg on page 363	Sends a message to a User Datagram Protocol (UDP) port.
sp_serveroption on page 365	Displays or changes remote server options.
sp_setlangalias on page 368	Assigns or changes the alias for an alternate language.
sp_setrowlockpromote on page 373	Sets or changes the lock promotion thresholds for a database, for a table, or for Adaptive Server.
sp_setpsexex on page 371	Sets custom execution attributes “on the fly” for an active client application, login, or stored procedure.
sp_set_qplan on page 372	Changes the text of the abstract plan of an existing plan without changing the associated query.
sp_setsuspect_granularity on page 376	Displays and sets the recovery fault isolation mode.
sp_setsuspect_threshold on page 378	On recovery, sets the maximum number of suspect pages that Adaptive Server will allow in the specified database before taking the entire database offline.
sp_showcontrolinfo on page 379	Displays information about engine group assignments, bound client applications, logins, and stored procedures.
sp_showexeclass on page 381	Displays the execution class attributes and the engines in any engine group associated with the specified execution class.
sp_showplan on page 382	Displays the query plan for any user connection for the current SQL statement (or a previous statement in the same batch). The query plan is displayed in showplan format.
sp_showpsexex on page 383	Displays execution class, current priority, and affinity for all processes running on Adaptive Server.
sp_spaceused on page 385	Displays estimates of the number of rows, the number of data pages, and the space used by one table or by all tables in the current database.
sp_ssladmin on page 387	Adds, deletes, or displays a list of server certificates for Adaptive Server.
sp_syntax on page 389	Displays the syntax of Transact-SQL statements, system procedures, utilities, and other routines, depending on which products and corresponding sp_syntax scripts exist on Adaptive Server. As of 12.0, this sproc is not supported. Retaining info here in case it becomes viable again in a future release. -- ewheeler 2Dec99
sp_sysmon on page 390	Displays performance information.
sp_thresholdaction on page 393	Executes automatically when the number of free pages on the log segment falls below the last-chance threshold, unless the threshold is associated with a different procedure. Sybase does not provide this procedure.
sp_transactions on page 395	Reports information about active transactions.
sp_unbindcache on page 403	Unbinds a database, table, index, text object, or image object from a data cache.

Procedure	Description
sp_unbindcache_all on page 405	Unbinds all objects that are bound to a cache.
sp_unbindefault on page 405	Unbinds a created default value from a column or from a user-defined datatype.
sp_unbindefault on page 405	Unbinds a database, table, index, text object, or image object from a data cache.
sp_unbindmsg on page 408	Unbinds a user-defined message from a constraint.
sp_unbindrule on page 409	Unbinds a rule from a column or from a user-defined datatype.
sp_volchanged on page 410	Notifies the Backup Server™ that the operator performed the requested volume handling during a dump or load.
sp_who on page 413	Reports information about all current Adaptive Server users and processes or about a particular user or process.

sp_activeroles

Description	Displays all active roles.
Syntax	<code>sp_activeroles [expand_down]</code>
Parameters	<code>expand_down</code> – shows the hierarchy tree of all active roles contained by your roles.

Examples

Example 1

```
sp_activeroles

Role Name
-----
sa_role
sso_role
oper_role
replication_role
```

Example 2

```
sp_activeroles expand_down
```

Role Name	Parent Role Name	Level
sa_role	NULL	1
doctor_role	NULL	1
oper_role	NULL	1

Usage	<ul style="list-style-type: none"> • <code>sp_activeroles</code> displays all your active roles and all roles contained by those roles. • For information about creating, managing, and using roles, see the <i>System Administration Guide</i>.
Permissions	Any user can execute <code>sp_activeroles</code> .
See also	<i>Commands</i> – <code>alter role</code> , <code>create role</code> , <code>drop role</code> , <code>grant</code> , <code>revoke</code> , <code>set</code> <i>Functions</i> – <code>mut_excl_roles</code> , <code>proc_role</code> , <code>role_contain</code> , <code>role_name</code> <i>System procedures</i> – <code>sp_displayroles</code>

Procedures: *sp_addalias* – *sp_add_resource_limit*

sp_addalias

Description	Allows an Adaptive Server user to be known in a database as another user.
Syntax	<code>sp_addalias loginame, name_in_db</code>
Parameters	<p><code>loginame</code> – is the master.dbo.syslogins name of the user who wants an alternate identity in the current database.</p> <p><code>name_in_db</code> – is the database user name to alias <i>loginame</i> to. The name must exist in both master.dbo.syslogins and in the sysusers table of the current database.</p>
Examples	<pre>sp_addalias victoria, albert</pre> <p>There is a user named “albert” in the database’s sysusers table and a login for a user named “victoria” in master.dbo.syslogins. This command allows “victoria” to use the current database by assuming the name “albert”.</p>
Usage	<ul style="list-style-type: none"> • Executing <code>sp_addalias</code> maps one user to another in the current database. The mapping is shown in <code>sysalternates</code>, where the two users’ <code>suids</code> (system user IDs) are connected. • A user can be aliased to only one database user at a time. • A report on any users mapped to a specified user can be generated with <code>sp_helpuser</code>, giving the specified user’s name as an argument. • When a user tries to use a database, Adaptive Server checks <code>sysusers</code> to confirm that the user is listed there. If the user is not listed there, Adaptive Server then checks <code>sysalternates</code>. If the user’s <code>suid</code> is listed in <code>sysalternates</code>, mapped to a database user’s <code>suid</code>, Adaptive Server treats the first user as the second user while using the database. <p>If the user named in <i>loginame</i> is in the database’s <code>sysusers</code> table, Adaptive Server does not use the user’s alias identity, because it checks <code>sysusers</code> and finds the <code>loginame</code> before checking <code>sysalternates</code>, where the alias is listed.</p>

Permissions	Only the Database Owner or a System Administrator can execute sp_addalias.
See also	<i>Command</i> – use <i>System procedures</i> – sp_addlogin, sp_adduser, sp_dropalias, sp_helpuser

sp_addauditrecord

Description	Allows users to enter user-defined audit records (comments) into the audit trail.
Syntax	sp_addauditrecord [text [, db_name [, obj_name [, owner_name [, dbid [, objid]]]]]]
Parameters	<i>text</i> – is the text of the message to add to the current audit table. The text is inserted into the extrainfo field of the table. <i>db_name</i> – is the name of the database referred to in the record. The name is inserted into the dbname field of the current audit table. <i>obj_name</i> – is the name of the object referred to in the record. The name is inserted into the objname field of the current audit table. <i>owner_name</i> – is the owner of the object referred to in the record. The name is inserted into the objowner field of the current audit table. <i>dbid</i> – is the database ID number of db_name. Do not enclose this integer value in quotes. <i>dbid</i> is inserted into the dbid field of the current audit table. <i>objid</i> – is the object ID number of obj_name. Do not enclose this integer value in quotes. <i>objid</i> is inserted into the objid field of the current audit table.

Examples	Example 1
----------	------------------

```
sp_addauditrecord "I gave A. Smith permission to view the payroll table in the corporate database. This permission was in effect from 3:10 to 3:30 pm"
```

```
on 9/22/92.", "corporate", "payroll", "dbo", 10,  
1004738270
```

Adds “I gave A. Smith permission to view the payroll table in the corporate database. This permission was in effect from 3:10 to 3:30 pm on 9/22/92.” to the `extrainfo` field; “corporate” to the `dbname` field; “payroll” to the `objname` field; “dbo” to the `objowner` field; “10” to the `dbid` field, and “1004738270” to the `objid` field of the current audit table.

Example 2

```
sp_addauditrecord @text="I am disabling auditing  
briefly while we reconfigure the system",  
@db_name="corporate"
```

Adds this record to the audit trail. This example uses parameter names with the `@` prefix, which allows you to leave some fields empty.

Usage

- Adaptive Server writes all audit records to the current audit table. The current audit table is determined by the value of the current audit table configuration parameter, set with `sp_configure`. An installation can have up to eight system audit tables, named `sysaudits_01`, `sysaudits_02`, and so forth, through `sysaudits_08`.

Note The records actually are first stored in the in-memory audit queue, and the audit process later writes the records from the audit queue to the current audit table. Therefore, you cannot count on an audit record being stored immediately in the audit table.

- You can use `sp_addauditrecord` if:
 - You have been granted execute permission on `sp_addauditrecord`. (No special role is required.)
 - Auditing is enabled. (A System Security Officer used `sp_configure` to turn on the auditing configuration parameter.)
 - The `adhoc` option of `sp_audit` is set to on.

Permissions

Only a System Security Officer can execute `sp_addauditrecord`. The Database Owner of `sybsecurity` (who must also be a System Security Officer) can grant execute permission to other users.

See also

System procedure – `sp_audit`

sp_addaudittable

Description	Adds another system audit table after auditing is installed.
Syntax	sp_addaudittable <i>devname</i>
Parameters	<p>devname</p> <ul style="list-style-type: none">– is the name of the device for the audit table. Specify a device name or specify “default”. If you specify “default”, Adaptive Server creates the audit table on the same device as the sybsecurity database. Otherwise, Adaptive Server creates the table on the device you specify.
Examples	<p>Example 1</p> <pre>sp_addaudittable auditdev2</pre> <p>Creates a system audit table on auditdev2. If only one system audit table (sysaudits_01) exists when you execute the procedure, Adaptive Server names the new audit table sysaudits_02 and places it on its own segment, called aud_seg_02, on auditdev2.</p> <p>Example 2</p> <pre>sp_addaudittable "default"</pre> <p>Creates a system audit table on the same device as the sybsecurity database. If two system audit tables (sysaudits_01 and sysaudits_02) exist when you execute the procedure, Adaptive Server names the new audit table sysaudits_03 and places it on its own segment, called aud_seg_03, on the same device as the sybsecurity database.</p>
Usage	<ul style="list-style-type: none">• Auditing must already be installed when you run sp_addaudittable. Follow this procedure to add a system audit table:<ol style="list-style-type: none">a Create the device for the audit table, using disk init. For example, run a command like this for UNIX:<pre>disk init name = "auditdev2", physname = "/dev/rxyla", size = "5K"</pre>b Add the device to the sybsecurity database with the alter database command. For example, to add <i>auditdev2</i> to the sybsecurity database, use this command:<pre>alter database sybsecurity on auditdev2</pre>c Execute sp_addaudittable to create the table.

- Adaptive Server names the new system audit table and the new segment according to how many audit tables are already defined. For example, if five audit tables are defined before you execute the procedure, Adaptive Server names the new audit table `sysaudits_06` and the new segment `aud_seg_06`. If you specify “default”, Adaptive Server places the segment on the same device as the `sybsecurity` database. Otherwise, Adaptive Server places the segment on the device you name.
- A maximum of eight audit tables is allowed. If you already have eight audit tables, and you attempt to execute `sp_addauditable` to add another one, Adaptive Server displays an error message.
- For information about how to install auditing, see the installation documentation for your platform. For a discussion about how to use auditing, see the *System Administration Guide*.

Permissions Only a user who is both a System Administrator and a System Security Officer can execute `sp_addauditable`.

See also *System procedure* – `sp_audit`

sp_addengine

Description Adds an engine to an existing engine group or, if the group does not exist, creates an engine group and adds the engine.

Syntax `sp_addengine engine_number, engine_group`

Parameters `engine_number`
– is the number of the engine you are adding to the group. Legal values are between 0 and a maximum equal to the number of configured online engines minus one.

`engine_group`
– is the name of the engine group to which you are adding the engine. If `engine_group` does not exist, Adaptive Server creates it and adds the engine to it. Engine group names must conform to the rules for identifiers. For details, see Chapter 7, “Expressions, Identifiers, and Wildcard Characters.”

Examples `sp_addengine 2, DS_GROUP`

If no engine group is called DS_GROUP, this statement establishes the group. If DS_GROUP already exists, this statement adds engine number 2 to that group.

Usage

- sp_addengine creates a new engine group if the value of engine_group does not already exist.
- The engine groups ANYENGINE and LASTONLINE are predefined. ANYENGINE includes all existing engines. LASTONLINE specifies the engine with highest engine number. A System Administrator can create additional engine groups. You cannot modify predefined engine groups.
- As soon as you use sp_bindexeclclass to bind applications or logins to an execution class associated with engine_group, the associated process may start running on engine_number.
- Prior to making engine affinity assignments, study the environment and consider the number of non-preferred applications and the number of Adaptive Server engines available. For more information about non-preferred applications, see the Performance and Tuning Guide.

Permissions

Only a System Administrator can execute sp_addengine.

See also

System procedures – sp_addexeclclass, sp_bindexeclclass, sp_clearpsex, sp_dropengine, sp_setpsex, sp_showcontrolinfo, sp_showexeclclass, sp_showpsex, sp_unbindexeclclass

sp_addexeclclass

Description

Creates or updates a user-defined execution class that you can bind to client applications, logins, and stored procedures.

Syntax

sp_addexeclclass *classname*, *priority*, *timeslice*, *engine_group*

Parameters

classname

– is the name of the new execution class.

priority

– is the priority value with which to run the client application, login, or stored procedure after it is associated with this execution class. Legal values are HIGH, LOW, and MEDIUM.

timeslice

- is the time unit assigned to processes associated with this class. Adaptive Server currently ignores this parameter.

engine_group

- identifies an existing group of engines on which processes associated with this class can run.

Examples

```
sp_addexeclass "DS", "LOW", 0, "DS_GROUP"
```

This statement defines a new execution class called DS with a *priority* value of LOW and associates it with the engine group DS_GROUP.

Usage

- *sp_addexeclass* creates or updates a user-defined execution class that you can bind to client applications, logins, and stored procedures. If the class already exists, the class attribute values are updated with the values supplied by the user.
- Use the predefined engine group parameter ANYENGINE if you do not want to restrict the execution object to an engine group.
- Use *sp_addengine* to define engine groups. Use *sp_showexeclass* to display execution class attributes and the engines in any engine group associated with the specified execution class. *sp_showcontrolinfo* lists the existing engine groups.

Permissions

Only a System Administrator can execute *sp_addexeclass*.

See also

System procedures – *sp_addengine*, *sp_bindexeclass*, *sp_clearpsex*, *sp_dropengine*, *sp_dropexeclass*, *sp_setpsex*, *sp_showcontrolinfo*, *sp_showexeclass*, *sp_showpsex*, *sp_unbindexeclass*

sp_addextendedproc

Description

Creates an extended stored procedure (ESP) in the master database.

Syntax

```
sp_addextendedproc esp_name, dll_name
```

Parameters

esp_name

- is the name of the extended stored procedure. This name must be identical to the name of the procedural language function that implements the ESP. *esp_name* must be a valid Adaptive Server identifier.

dll_name

– is the name of the dynamic link library (DLL) file containing the function specified by *esp_name*. The *dll_name* can be specified with no extension or with its platform-specific extension, such as *.dll* on Windows NT or *.so* on Sun Solaris. If an extension is specified, the *dll_name* must be enclosed in quotation marks.

Examples

```
sp_addextendedproc xp_echo, "sqlsrvdll.dll"
```

Registers an ESP for the function named `xp_echo`, which is in the `sqlsrvdll.dll` file. The name of the resulting ESP database object is also `xp_echo`.

Usage

- Execute `sp_addextendedproc` from the master database.
- The *esp_name* is case sensitive. It must match the name of the function in the DLL.
- The DLL represented by *dll_name* must reside on the server machine on which the ESP is being created and the DLL directory must be in the *\$PATH* on Windows NT, the *\$LD_LIBRARY_PATH* on Digital UNIX, or the *\$SH_LIBRARY_PATH* on HP. If the file is not found, the search mechanism also searches *\$\$SYBASE/dll* on Windows NT and *\$\$SYBASE/lib* on other platforms.
- On Windows NT, an ESP function should not call a C run-time signal routine. This can cause XP Server to fail, because Open Server™ does not support signal handling on Windows NT.

Permissions

Only a System Administrator can execute `sp_addextendedproc`.

See also

Commands – create procedure

System procedures – `sp_dropextendedproc`, `sp_helpextendedproc`

sp_addexternlogin

Description

Component Integration Services only – Creates an alternate login account and password to use when communicating with a remote server through Component Integration Services.

Syntax

```
sp_addexternlogin servername, loginname, externname  
[, externpassword] [rolename]
```


Parameters

server

– is the name of the remote server that has been added to the local server with *sp_addserver*.

loginname

– is the name of the Adaptive Server login account for which to create an alternate login account.

externname

– is the name of an account on the remote server *server*. This account is used when logging into *server*.

externpassword

– is the password for *externname*

rolename

the name of the role the user is assigned when they log into Adaptive Server.

Examples

Example 1

```
sp_addexternlogin JOBSERV, sa, system, sys_pass
```

Allows the local server to gain access to the remote server JOBSERV using the remote login “system” and the remote password “sys_pass” on behalf of user “sa”.

Example 2

```
sp_addexternlogin CIS1012, bobj, jordan, hitchpost
```

When the user “bobj” logs into the remote server CIS1012, he connects using the remote server login name “jordan” and the password “hitchpost”.

Example 3

```
sp_addexternlogin DB2, NULL, sybase, syb_password
```

All the users who connect to the DB2 database have the name sybase and the password syb_password.

Usage

- *sp_addexternlogin* assigns an alternate login name and password to be used when communicating with a remote server. It stores the password internally in encrypted form.

You can use *sp_addexternlogin* only when Component Integration Services is installed and configured.

- The login and password have a many to one mapping. That is, you can assign all the users who need to log into a remote server the same name and password.
- When you establish a connection to a remote server for a user that has more than one role active, each role is searched for an external login mapping and uses the first mapping it finds to establish the login. This is the same order as displayed by the stored procedure sp_activeroles.
- If you perform role mapping, and a user's role is changed (using set role), any connections made to remote servers that used role mapping must be disconnected. You cannot do this if a transaction is pending. You cannot use set role if a transaction is active and remote connections are present that used role mapping.
- You can assign external logins to Adaptive Server roles. You can assign anyone with a particular role corresponding login name and password for any given remote server.
- Before running sp_addexternlogin, add the remote server to Adaptive Server with sp_addserver.
- *externname* and *externpassword* must be a valid user and password combination on the node where the *server* runs.
- Sites with automatic password expiration need to plan for periodic updates of passwords for external logins.
- Use sp_dropexternlogin to remove the definition of the external login.
- sp_addexternlogin cannot be used from within a transaction.
- The “sa” account and the *loginname* account are the only users who can modify remote access for a given local user.

Permissions

Only the *loginname*, a System Administrator, and a System Security Officer can execute sp_addexternlogin.

See also

System procedures – sp_addserver, sp_dropexternlogin, sp_helpserver

sp_addgroup

Description

Adds a group to a database. Groups are used as collective names in granting and revoking privileges.

Syntax

sp_addgroup *grpname*

Parameters	<i>grpname</i> – is the name of the group. Group names must conform to the rules for identifiers.
Examples	<code>sp_addgroup accounting</code> Creates a group named <code>accounting</code> in the current database.
Usage	<ul style="list-style-type: none">• <code>sp_addgroup</code> adds the new group to a database’s <code>sysusers</code> table. Each group’s user ID (<code>uid</code>) is 16384 or larger (except “public,” which is always 0).• A group and a user cannot have the same name.• Once a group has been created, add new users with <code>sp_adduser</code>. To add an existing user to a group, use <code>sp_changegroup</code>.• Every database is created with a group named “public”. Every user is automatically a member of “public”. Each user can be a member of one additional group.
Permissions	Only the Database Owner, a System Administrator, or a System Security Officer can execute <code>sp_addgroup</code> .
See also	<i>Commands</i> – <code>grant</code> , <code>revoke</code> <i>System procedures</i> – <code>sp_adduser</code> , <code>sp_changegroup</code> , <code>sp_dropgroup</code> , <code>sp_helpgroup</code>

sp_addlanguage

Description	Defines the names of the months and days for an alternate language and its date format.
Syntax	<code>sp_addlanguage language, alias, months, shortmons, days, datefmt, datefirst</code>
Parameters	<i>language</i> – is the official language name for the language, entered in 7-bit ASCII characters only. <i>alias</i> – substitutes for the alternate language’s official name. Enter either “null”, to make the alias the same as the official language name, or a name you prefer. You can use 8-bit ASCII characters in an alias—”français”, for example—if your terminal supports them.

months

– is a list of the full names of the 12 months, ordered from January through December, separated only by commas (no spaces allowed). Month names can be up to 20 characters long and can contain 8-bit ASCII characters.

shortmons

– is a list of the abbreviated names of the 12 months, ordered from January through December, separated only by commas (no spaces allowed). Month abbreviations can be up to 9 characters long and can contain 8-bit ASCII characters.

days

– is a list of the full names of the seven days, ordered from Monday through Sunday, separated only by commas (no spaces allowed). Day names can be up to 30 characters long and can contain 8-bit ASCII characters.

datefmt

– is the date order of the date parts *month/day/year* for entering *datetime* or *smalldatetime* data. Valid arguments are *mdy*, *dmy*, *ymd*, *ydm*, *myd*, or *dym*. “*dmy*” indicates that dates are in day/month/year order.

datefirst

– sets the number of the first weekday for date calculations. For example, Monday is 1, Tuesday is 2, and so on.

Examples

```
sp_addlanguage french, null,  
"janvier,fevrier,mars,avril,mai,juin,juillet,  
aout,septembre,octobre,novembre,decembre",  
"jan,fev,mars,avr,mai,juin,juil,aout,sept,oct,  
nov,dec",  
"lundi,mardi,mercredi,jeudi,vendredi,samedi,  
dimanche",  
dmy, 1
```

This stored procedure adds French to the languages available on the server. “*null*” makes the alias the same as the official name, “*french*”. Date order is “*dmy*”—day/month/year. “*1*” specifies that *lundi*, the first item in the *days* list, is the first weekday. Because the French do not capitalize the names of the days and months except when they appear at the beginning of a sentence, this example shows them being added in lowercase.

Usage

- Usually, you add alternate languages from one of Adaptive Server's Language Modules using the *langinstall* utility or the Adaptive Server installation program. A Language Module supplies the names of the dates and translated error messages for that language. However, if a Language Module is not provided with your server, use *sp_addlanguage* to define the date names and format.
- Use *sp_modifylogin* to change a user's default language. If you set a user's default language to a language added with *sp_addlanguage*, and there are no localization files for the language, the users receive an informational message when they log in, indicating that their client software could not open the localization files.

System Table Changes

- *sp_addlanguage* creates an entry in *master.dbo.syslanguages*, inserting a unique numeric value in the *langid* column for each alternate language. *langid* 0 is reserved for U.S. English.
- The *language* parameter becomes the official language name, stored in the *name* column of *master.dbo.syslanguages*. Language names must be unique. Use *sp_helplanguage* to display a list of the alternate languages available on Adaptive Server.
- *sp_addlanguage* sets the *alias* column in *master.dbo.syslanguages* to the official language name if NULL is entered for *alias*, but System Administrators can change the value of *syslanguage.alias* with *sp_setlangalias*.
- *sp_addlanguage* sets the *upgrade* column in *master.dbo.syslanguages* to 0.

Dates for Languages added with sp_addlanguage

- For alternate languages added with Language Modules, Adaptive Server sends date values to clients as *datetime* datatype, and the clients use localization files to display the dates in the user's current language. For date strings added with *sp_addlanguage*, use the *convert* function to convert the dates to character data in the server:

```
select convert(char, pubdate) from table
```

where *pubdate* is *datetime* data and *table* is any table.

- When users perform data entry on date values and need to use date names created with *sp_addlanguage*, the client must have these values input as character data, and sent to the server as character data.

Permissions

Only a System Administrator can execute *sp_addlanguage*.

See also

Commands – set

System procedures – sp_droplanguage, sp_helplanguage, sp_modifylogin, sp_setlangalias

sp_addlogin

Description

Adds a new user account to Adaptive Server; specifies the password expiration interval, the minimum password length, and the maximum number of failed logins allowed for a specified login at creation.

Syntax

```
sp_addlogin loginame, passwd [, defdb]  
           [, deflanguage] [, fullname] [, passwdexp]  
           [, minpwrlen] [, maxfailedlogins]
```

Parameters

loginame

– is the user’s login name. Login names must conform to the rules for identifiers.

passwd

– is the user’s password. Passwords must be at least 6 characters long. If you specify a shorter password, sp_addlogin returns an error message and exits. Enclose passwords that include characters besides A-Z, a-z, or 0-9 in quotation marks. Also enclose passwords that *begin* with 0-9 in quotation marks.

defdb

– is the name of the default database assigned when a user logs into Adaptive Server. If you do not specify *defdb*, the default, master, is used.

deflanguage

– is the official name of the default language assigned when a user logs into Adaptive Server. The Adaptive Server default language, defined by the default language id configuration parameter, is used if you do not specify *deflanguage*.

fullname

– is the full name of the user who owns the login account. This can be used for documentation and identification purposes.

passwdexp

– specifies the password expiration interval in days. It can be any value between 0 and 32767, inclusive.

minpwrlen

– specifies the minimum password length required for that login. The values range between 0 and 30 characters.

maxfailedlogins

– is the number of allowable failed login attempts. It can be any whole number between 0 and 32767.

Examples

Example 1

```
sp_addlogin albert, longer1, corporate
```

Creates an Adaptive Server login for “albert” with the password “longer1” and the default database corporate.

Example 2

```
sp_addlogin claire, bleurouge, public_db, french
```

Creates an Adaptive Server login for “claire”. Her password is “bleurouge”, her default database is public_db, and her default language is French.

Example 3

```
sp_addlogin robertw, terrible2, public_db, null,  
"Robert Willis"
```

Creates an Adaptive Server login for “robertw”. His password is “terrible2”, his default database is public_db, and his full name is “Robert Willis”. Do not enclose null in quotes.

Example 4

```
sp_addlogin susan, wonderful, null, null, "Susan B.  
Anthony"
```

Creates a login for “susan” with a password of “wonderful”, a full name of “Susan B. Anthony”, and the server’s default database and language. Do not enclose null in quotes.

Example 5

```
sp_addlogin susan, wonderful,  
@fullname="Susan B. Anthony"
```

An alternative way of creating the login shown in example 4.

Usage	<ul style="list-style-type: none">• For ease of management, it is strongly recommended that all users' Adaptive Server login names be the same as their operating system login names. This makes it easier to correlate audit data between the operating system and Adaptive Server. Otherwise, keep a record of the correspondence between operating system and server login names.• After assigning a default database to a user with sp_addlogin, the Database Owner or System Administrator must provide access to the database by executing sp_adduser or sp_addalias.• Although a user can use sp_modifylogin to change his or her own default database at any time, a database cannot be used without permission from the Database Owner.• A user can use sp_password at any time to change his or her own password. A System Security Officer can use sp_password to change any user's password.• A user can use sp_modifylogin to change his or her own default language. A System Administrator can use sp_modifylogin to change any user's default language.• A user can use sp_modifylogin to change his or her own <i>fullname</i>. A System Administrator can use sp_modifylogin to change any user's <i>fullname</i>.
Permissions	Only a System Administrator or a System Security Officer can execute sp_addlogin.
See also	<i>System procedures</i> – sp_addalias, sp_adduser, sp_droplogin, sp_locklogin, sp_modifylogin, sp_password, sp_role

sp_addmessage

Description	Adds user-defined messages to sysusermessages for use by stored procedure print and raiserror calls and by sp_bindmsg.
Syntax	sp_addmessage <i>message_num</i> , <i>message_text</i> [, <i>language</i> [, <i>with_log</i> [, <i>replace</i>]]]
Parameters	<i>message_num</i> – is the message number of the message to add. The message number for a user-defined message must be 20000 or greater.

message_text

– is the text of the message to add. The maximum length is 1024 bytes.

language

– is the language of the message to add. This must be a valid language name in the *syslanguages* table. If this parameter is missing, Adaptive Server assumes that messages are in the default session language indicated by *@@langid*.

with_log

– specifies whether the message is logged in the Adaptive Server error log as well as in the Windows NT Event Log on Windows NT servers, if logging is enabled. If *with_log* is TRUE, the message is logged, regardless of the severity of the error. If *with_log* is FALSE, the message may or may not be logged, depending on the severity of the error. If you do not specify a value for *with_log*, the default is FALSE.

replace

– specifies whether to overwrite an existing message of the same number and *languid*. If *replace* is specified, the existing message is overwritten; if *replace* is omitted, it is not. If you do not specify a value for *replace*, the default, FALSE, is used.

Examples

Example 1

```
sp_addmessage 20001, "The table '%1!' is not owned  
by the user '%2!'."
```

Adds a message with the number 20001 to *sysusermessages*.

Example 2

```
sp_addmessage 20002, "The procedure '%1!' is not  
owned by the user '%2!'." , NULL, TRUE, "replace"
```

Adds a message with the number 20002 to *sysusermessages*. This message is logged in the Adaptive Server error log, as well as in the Windows NT Event Log on Windows NT servers, if event logging is enabled. If a message numbered 20002 exists in the default session language, this message overwrites the old message.

Usage

- *sp_addmessage* does not overwrite an existing message of the same number and *languid* unless you specify *@replace = "replace"*.

- print and raiserror recognize placeholders in the message text to print out. A single message can contain up to 20 unique placeholders in any order. These placeholders are replaced with the formatted contents of any arguments that follow the message when the text of the message is sent to the client.

The placeholders are numbered to allow reordering of the arguments when Adaptive Server is translating a message to a language with a different grammatical structure. A placeholder for an argument appears as “%nn!”, a percent sign (%), followed by an integer from 1 to 20, followed by an exclamation point (!). The integer represents the argument number in the string in the argument list. “%1!” is the first argument in the original version, “%2!” is the second argument, and so on.

Permissions	Any user can execute sp_addmessage.
See also	<i>Commands</i> – print, raiserror <i>System procedures</i> – sp_altermessage, sp_dropmessage, sp_getmessage

sp_addobjectdef

Description	<i>Component Integration Services only</i> – Specifies the mapping between a local table and an external storage location.
Syntax	sp_addobjectdef <i>tablename</i> , "objectdef" [,"objecttype"]
Parameters	<p><i>tablename</i></p> <ul style="list-style-type: none"> – is the name of the object as it is defined in a local table. The <i>tablename</i> can be in any of the following forms: <ul style="list-style-type: none"> • <i>dbname.owner.object</i> • <i>dbname..object</i> • <i>owner.object</i> • <i>object</i> <p><i>dbname</i> and <i>owner</i> are optional. <i>object</i> is required. If you do not specify an <i>owner</i>, the default (current user name) is used. If you specify a <i>dbname</i>, it must be the current database name, and you must specify <i>owner</i> or mark the owner with a placeholder in the format <i>dbname..object</i>. Enclose any multipart <i>tablename</i> values in quotes.</p>

objectdef

– is a string naming the external storage location of the object. The *objecttype* at *objectdef* can be a table, view, or read-only remote procedure call (RPC) result set accessible to a remote server. A table, view, or RPC uses the following format for *objectdef*:

server_name.dbname.owner.object

server_name and *object* are required. *dbname* and *owner* are optional, but if they are not supplied, a placeholder in the format *dbname..object*, is required.

For more information, see “Server Classes” in the Component Integration Services User’s Guide.

objecttype

– is one of the values that specify the format of the object named by *objectdef*. Table 3-1 describes the valid values. Enclose the *objecttype* value in quotes.

Table 3-1: Allowable values for *objecttype*

Value	Description
table	Indicates that the object named by <i>objectdef</i> is a table accessible to a remote server. This value is the default for <i>objecttype</i> .
view	Indicates that the object named by <i>objectdef</i> is a view managed by a remote server and processed as a table.
rpc	Indicates that the object named by <i>objectdef</i> is an RPC managed by a remote server. Adaptive Server processes the result set from the RPC as a read-only table.

Table 3-2 summarizes how each *objecttype* is used.

Table 3-2: Summary of *objecttype* uses

<i>objecttype</i>	<i>create table</i>	<i>create existing table</i>	Write to Table	Read from Table
table	Yes	Yes	Yes	Yes
view	No	Yes	Yes	Yes
rpc	No	Yes	No	Yes

Examples

Example 1

```
sp_addobjectdef "finance.dbo.accounts",
"SYBASE.pubs.dbo.accounts", "table"
```

Maps the local table accounts in the database finance to the remote object pubs.dbo.accounts in the remote server named SYBASE. The current database must be finance. A subsequent create table creates a table in the pubs database. If pubs.dbo.accounts is an existing table, a create existing table statement populates the table finance.dbo.accounts with information about the remote table.

Example 2

```
sp_addobjectdef stockcheck,  
"NEWYORK.wallstreet.kelly.stockcheck", "rpc"
```

Maps the local table stockcheck to an RPC named stockcheck on remote server NEWYORK in the database wallstreet with owner “kelly”. The result set from RPC stockcheck is seen as a read-only table. Typically, the next operation would be a create existing table statement for the object stockcheck.

Usage

- sp_addobjectdef specifies the mapping between a local table and an external storage location. It identifies the format of the object at that location.

You can use sp_addobjectdef only when Component Integration Services is installed and configured.

- sp_addobjectdef replaces the sp_addtabledef command. sp_addtabledef allows existing scripts to run without modification. Internally, sp_addtabledef invokes sp_addobjectdef.
- Only the System Administrator can provide the name of another user as a table owner.
- When *objecttype* is table, view, or rpc, the *objectdef* parameter takes the following form:

```
"server_name.database.owner.tablename"
```

- *server_name* represents a server that has already been added to sys.servers by sp_addserver.
- *database* may not be required. Some server classes do not support it.
- *owner* should always be provided, to avoid ambiguity. If you do not specify *owner*, the remote object referenced may vary, depending on whether or not the external login corresponds to the remote object owner.
- *tablename* is the name of a remote server table.

- Use *sp_addobjectdef* before issuing any create table or create existing table commands. create table is valid only for the *objecttype* values table and file. When either create table or create existing table is used, Adaptive Server checks sysattributes to determine whether any table mapping has been specified for the object. Follow the *objecttype* values view and rpc with create existing table statements.
- After the table has been created, all future references to the local table name (by select, insert, delete and update) are mapped to the correct location.
- For information about RMS, see the Component Integration Services User's Guide.

Permissions

Any user can execute *sp_addobjectdef*.

See also

Commands – create existing table, create table, drop table

System procedures – *sp_addlogin*, *sp_addserver*, *sp_defaultloc*, *sp_dropobjectdef*, *sp_helpserver*

sp_add_qpgroup

Description

Adds an abstract plan group.

Syntax

sp_add_qpgroup new_name

Parameters

new_name

– is the name of the new abstract plan group. Group names must be valid identifiers.

Examples

```
sp_add_qpgroup dev_plans
```

Creates a new abstract plan group named dev_plans.

Usage

- Use *sp_add_qpgroup* to add abstract plan groups for use in capturing or creating abstract plans. The abstract plan group must exist before you can create, save, or copy plans into a group.
- *sp_add_qpgroup* cannot be run in a transaction.

Permissions

Only a System Administrator or Database Owner can execute *sp_add_qpgroup*.

See also

Commands – set

System procedures – *sp_help_qpgroup*

sp_addremotelogin

Description	Authorizes a new remote server user by adding an entry to master.dbo.sysremotelogins.
Syntax	sp_addremotelogin <i>remoteserver</i> [, <i>loginame</i> [, <i>remotename</i>]]
Parameters	<i>remoteserver</i> – is the name of the remote server to which the remote login applies. This server must be known to the local server by an entry in the master.dbo.sysservers table, which was created with sp_addserver.

Note This manual page uses the term "local server" to refer to the server that is executing the remote procedures run from a "remote server".

loginame

– is the login name of the user on the local server. *loginame* must already exist in the master.dbo.syslogins table.

remotename

– is the name used by the remote server when logging into the local server. All *remotenames* that are not explicitly matched to a local *loginame* are automatically matched to a local name. In example 1 , the local name is the remote name that is used to log in. In example 2 , the local name is “albert”.

Examples

Example 1

```
sp_addremotelogin GATEWAY
```

Creates an entry in the sysremotelogins table for the remote server GATEWAY, for purposes of login validation. This is a simple way to map remote names to local names when the local and remote servers have the same users.

This example results in a value of -1 for the suid column and a value of NULL for the remoteusername in a row of sysremotelogins.

Example 2

```
sp_addremotelogin GATEWAY, albert
```

Creates an entry that maps all logins from the remote server GATEWAY to the local user name “albert”. Adaptive Server adds a row to sysremotelogins with Albert’s server user ID in the suid column and a null value for the remoteusername.

For these logins to be able to run RPCs on the local server, they must specify a password for the RPC connection when they log into the local server, or they must be “trusted” on the local server. To define these logins as “trusted”, use *sp_remotoption*.

Example 3

```
sp_addremotelogin GATEWAY, ralph, pogo
```

Maps a remote login from the remote user “pogo” on the remote server GATEWAY to the local user “ralph”. Adaptive Server adds a row to *sysremotelogins* with Ralph’s server user ID in the *suid* column and “pogo” in the *remoteusername* column.

Usage

- When a remote login is received, the local server tries to map the remote user to a local user in three different ways:
 - First, the local server looks for a row in *sysremotelogins* that matches the remote server name and the remote user name. If the local server finds a matching row, the local server user ID for that row is used to log in the remote user. This applies to mappings from a specified remote user.
 - If no matching row is found, the local server searches for a row that has a null remote name and a local server user ID other than -1. If such a row is found, the remote user is mapped to the local server user ID in that row. This applies to mappings from any remote user from the remote server to a specific local name.
 - Finally, if the previous attempts failed, the local server checks the *sysremotelogins* table for an entry that has a null remote name and a local server user ID of -1. If such a row is found, the local server uses the remote name supplied by the remote server to look for a local server user ID in the *syslogins* table. This applies when login names from the remote server and the local server are the same.
- The name of the local user may be different on the remote server.
- If you use *sp_addremotelogin* to map all users from a remote server to the same local name, use *sp_remotoption* to specify the “trusted” option for those users. For example, if all users from the server GOODSRV that are mapped to “albert” are to be “trusted”, use *sp_remotoption* as follows:

```
sp_remotoption GOODSRV, albert, NULL, trusted,  
true
```

Logins that are not specified as “trusted” cannot execute RPCs on the local server unless they specify passwords for the local server when they log into the remote server. In Open Client™ Client-Library™, the user can use the ct_remote_pwd routine to specify a password for server-to-server connections. isql and bcp do not permit users to specify a password for RPC connections.

If users are logged into the remote server using “unified login”, these logins are already authenticated by a security mechanism. These logins must also be trusted on the local server, or the users must specify passwords for the server when they log into the remote server.

- For more information about setting up servers for remote procedure calls and for using “unified login”, see the System Administration Guide.
- Every remote login entry has a status. The default status for the trusted option is false (not trusted). This means that when a remote login comes in using that entry, the password is checked. If you do not want the password to be checked, change the status of the trusted option to true with sp_remotoption.

Permissions

Only a System Administrator can execute sp_addremotelogin.

See also

System procedures – sp_addlogin, sp_addserver, sp_dropremotelogin, sp_helpremotelogin, sp_helprotect, sp_helpserver, sp_remotoption

Utility – isql

sp_add_resource_limit

Description

Creates a limit on the number of server resources that can be used by an Adaptive Server login and/or an application to execute a query, query batch, or transaction.

Syntax

```
sp_add_resource_limit name, appname, rangename, limittype, limitvalue  
[, enforced [, action [, scope ]]]
```

Parameters

name

– is the Adaptive Server login to which the limit applies. You must specify either a *name* or an *appname* or both. To create a limit that applies to all users of a particular application, specify a *name* of NULL.

appname

– is the name of the application to which the limit applies. You must specify either a *name* or an *appname* or both. To create a limit that applies to all applications used by an Adaptive Server login, specify an *appname* of null. To create a limit that applies to a particular application, specify the application name that the client program passes to the Adaptive Server in the login packet.

rangename

– is the time range during which the limit is enforced. The time range must exist in the *systimeranges* system table of the master database at the time you create the limit.

limittype

– is the type of resource to limit. This must be one of the following:

Limit Type	Description
<i>row_count</i>	Limits the number of rows a query can return
<i>elapsed_time</i>	Limits the number of seconds, in wall-clock time, that a query batch or transaction can run
<i>io_cost</i>	Limits either the actual cost or the optimizer's cost estimate for processing a query

limitvalue

– is the maximum amount of the server resource (I/O cost, elapsed time in seconds, or row count) that can be used by the login or application before Adaptive Server enforces the limit. This must be a positive, nonzero integer that is less than or equal to 2^{31} . The following table indicates what value to specify for each limit type:

Limit Type	Limit Value
<i>row_count</i>	The maximum number of rows that can be returned by a query before the limit is enforced.
<i>elapsed_time</i>	The number of seconds, in wall-clock time, that a query batch or transaction can run before the limit is enforced.
<i>io_cost</i>	A unitless measure derived from the optimizer's costing formula.

enforced

– determines whether the limit is enforced prior to or during query execution. The following table lists the valid values for each limit type:

enforced Code	Description	Limit Type
1	Action is taken when the estimated I/O cost of execution exceeds the specified limit.	io_cost
2	Action is taken when the actual row count, elapsed time, or I/O cost of execution exceeds the specified limit.	row_count elapsed_time io_cost
3	Action is taken when either the estimated cost or the actual cost exceeds the specified limit.	io_cost

If you specify an *enforced* value of 3, Adaptive Server performs a logical “or” of 1 and 2. For example, assume *enforced* is set to 3. If you run a query whose *io_cost* exceeds the estimated cost, the specified *action* is executed. If the query is within the limits specified for estimated cost but exceeds the actual cost, the specified *action* is also executed.

If you do not specify an *enforced* value, Adaptive Server enforces limit 2 for *row_count* and *elapsed_time* and limit 3 for *io_cost*. In other words, if the limit type is *io_cost*, the specified action is executed if the query exceeds either the estimated or actual cost.

action

– is the action to take when the limit is exceeded. The following action codes are valid for all limit types:

action Code	Description
1	Issues a warning
2	Aborts the query batch
3	Aborts the transaction
4	Kills the session

If you do not specify an *action* value, Adaptive Server uses a default value of 2 (abort the query batch).

scope

– is the scope of the limit. Specify one of the following codes appropriate to the type of limit:

scope Code	Description	Limit Type
1	Query	io_cost row_count
2	Query batch (one or more SQL statements sent by the client to the server)	elapsed_time
4	Transaction	elapsed_time
6	Query batch <i>and</i> transaction	elapsed_time

If you do not specify a *scope* value, the limit applies to all possible scopes for the limit type.

Examples

Example 1

```
sp_add_resource_limit NULL, payroll, early_morning,
elapsed_time, 120, 2, 1, 2
```

Creates a resource limit that applies to all users of the payroll application during the early_morning time range. If the query batch takes more than 120 seconds to execute, Adaptive Server issues a warning.

Example 2

```
sp_add_resource_limit joe_user, NULL, midday,
row_count, 5000, 2, 3, 1
```

Creates a resource limit that applies to all ad hoc queries and applications run by “joe_user” during the midday time range. When a query returns more than 5000 rows, Adaptive Server aborts the transaction.

Example 3

```
sp_add_resource_limit joe_user, NULL, midday,
io_cost, 650, 1, 3, 1
```

Creates a resource limit that applies to all ad hoc queries and applications run by “joe_user” during the midday time range. When the optimizer estimates that the I/O cost would exceed 650, Adaptive Server aborts the transaction.

Usage

- You must enable `sp_configure` “allow resource limits” for resource limits to take effect.

- Multiple resource limits can exist for a given user, application, limit type, scope, and enforcement time, as long as their time ranges do not overlap.
- All limits for the currently active named time ranges and the “at all times” range for a login and/or application name are bound to the user’s session at login time. Therefore, if a user logs into Adaptive Server independently of a given application, resource limits that restrict the user in combination with that application do not apply. To guarantee restrictions on that user, create a resource limit that is specific to the user and independent of any application.
- Since either the user login name or application name, or both, are used to identify a resource limit, Adaptive Server observes a predefined search precedence while scanning the sysresourcelimits table for applicable limits for a login session. The following table describes the precedence of matching ordered pairs of login name and application name:

Level	Login Name	Application Name
1	“joe_user”	payroll
2	NULL	payroll
3	“joe_user”	NULL

If one or more matches are found for a given precedence level, no further levels are searched. This prevents conflicts regarding similar limits for different login/application combinations.

If no match is found at any level, no limit is imposed on the session.

- When you add, delete, or modify resource limits, Adaptive Server rebinds the limits for each session for that login and/or application at the beginning of the next query batch for that session.
- When you change the currently active time ranges, Adaptive Server rebinds limits for the session. This rebinding occurs at the beginning of the next query batch.
- You cannot associate the limits for a particular login, application, or login/application combination with named time ranges that overlap (except for limits that share the same time range).

For example, if a user is limited to retrieving 50 rows between 9:00 a.m. and 1:00 p.m., you cannot create a second resource limit for the same user that limits him to retrieving 100 rows between 10:00 a.m. and 12:00 noon. However, you can create a resource hierarchy by assigning the 100-row limit to the *user* between 10:00 a.m. and 12:00 noon and assigning the 50-row limit to an *application*, like *isql*, between 9:00 a.m. and 1:00 p.m.

- For more information on resource limits, see the *System Administration Guide*.

Permissions

Only a System Administrator can execute *sp_add_resource_limit*.

See also

System procedures – *sp_configure*, *sp_drop_resource_limit*,
sp_help_resource_limit, *sp_modify_resource_limit*

Utility – *isql*

Procedures: *sp_addsegment* – *sp_adduser*

sp_addsegment

Description	Defines a segment on a database device in a database.
Syntax	<code>sp_addsegment <i>segname</i>, <i>dbname</i>, <i>devname</i></code>
Parameters	<p><i>segname</i></p> <ul style="list-style-type: none"> – is the name of the new segment to add to the <code>syssegments</code> table of the database. Segment names are unique in each database. <p><i>dbname</i></p> <ul style="list-style-type: none"> – specifies the name of the database in which to define the segment. <i>dbname</i> must be the name of the current database or match the database name qualifying <code>sp_addsegment</code>. <p><i>devname</i></p> <ul style="list-style-type: none"> – is the name of the database device in which to locate <i>segname</i>. A database device can have more than one segment associated with it.

Examples

Example 1

```
sp_addsegment indexes, pubs2, dev1
```

Creates a segment named `indexes` for the database `pubs2` on the database device named `dev1`.

Example 2

```
disk init
  name = "pubs2_dev",
  physname = "/dev/pubs_2_dev",
  vdevno = 9, size = 5120
go
alter database pubs2 on pubs2_dev = 2
go
pubs2..sp_addsegment indexes, pubs2, dev1
```

Creates a segment named `indexes` for the database `pubs2` on the database device named `dev1`.

Usage

- sp_addsegment defines segment names for database devices created with disk init and assigned to a specific database with an alter database or create database command.
- After defining a segment, use it in create table and create index commands and in the sp_placeobject procedure to place a table or index on the segment.

When a table or index is created on a particular segment, all subsequent data for the table or index is located on the segment.

- Use the system procedure sp_extendsegment to extend the range of a segment to another database device used by the same database.
- If a database is extended with alter database on a device used by that database, the segments mapped to that device are also extended.
- The system and default segments are mapped to each database device included in a create database or alter database command. The logsegment is also mapped to each device, unless you place it on a separate device with the log on extension to create database or with sp_logdevice. For more information, see the System Administration Guide.
- If you attempt to use sp_addsegment in a database that has both data and the log on the same device, Adaptive Server returns an error message.

Permissions

Only the Database Owner or a System Administrator can execute sp_addsegment.

See also

Commands – alter database, create index, create table, disk init
System procedures – sp_dropsegment, sp_extendsegment, sp_helppdb, sp_helpdevice, sp_placeobject

sp_addserver

Description

Defines a remote server, or defines the name of the local server.

Syntax

sp_addserver *lname* [, *class* [, *pname*]]

Parameters

lname

– is the name used to address the server on your system. *sp_addserver* adds a row to the *sys.servers* table if there is no entry already present for *lname*. Server names must be unique and must conform to the rules for identifiers.

class

– identifies the category of server being added. Table 4-1 lists allowable values for the *class* parameter:

Table 4-1: Allowable values for *server_class* parameter

class Parameter Value	Description
access_server	Server coded to the DirectConnect™ specification (Component Integration Services only)
db2	Server accessible by Net-Gateway™ or MDI™ Database Gateway (Component Integration Services only)
direct_connect	Functionally the same as <i>access_server</i> (Component Integration Services only)
generic	Server coded to the Generic Access Module specification (Component Integration Services only)
local	Local server (there can be only one) used only once after start-up, or after restarting Adaptive Server, to identify the local server name so that it can appear in messages printed by Adaptive Server
null	Remote server with no category defined
sql_server	Another Adaptive Server or Omni server (this is the default value)

pname

– is the name in the interfaces file for the server named *lname*. This enables you to establish local aliases for other Adaptive Servers or Backup Servers that you may need to communicate with. If you do not specify a *pname*, *lname* is used.

Examples

Example 1

```
sp_addserver GATEWAY
```

Adds an entry for a remote server named GATEWAY in *master.dbo.sys.servers*. The *pname* is also GATEWAY.

Example 2

```
sp_addserver GATEWAY, null, VIOLET
```

Adds an entry for a remote server named GATEWAY in master.dbo.syssservers. The *pname* is VIOLET. If there is already a syssservers entry for GATEWAY with a different *pname*, the *pname* of server GATEWAY changes to VIOLET.

Example 3

```
sp_addserver PRODUCTION, local
```

Adds an entry for the local server named PRODUCTION.

Example 4

```
sp_addserver SQLSRV10, sql_server, SS_MOSS
```

Adds an entry for a remote server known to the local server as SQLSRV10. The remote server is of server class sql_server. The *network_name* for SQLSRV10 is SS_MOSS.

Example 5

```
sp_addserver RDBAM_ALPHA, generic, rdbam_alpha
```

Adds an entry for a remote server known to the local server as *RDBAM_ALPHA*. The remote server *RDBAM_ALPHA* is written to the Generic Access Module specification, which requires server class generic.

Usage

- The syssservers table identifies the name of the local server and its options, and any remote servers that the local server can communicate with.

To execute a remote procedure call on a remote server, the remote server must exist in the syssservers table.
- If *lname* already exists as a server name in the syssservers table, sp_addserver changes the remote server's srvnetname to the name specified by *pname*. When it does this, sp_addserver reports which server it changed, what the old network name was, and what the new network name is.
- The installation or upgrade process for your server adds an entry in syssservers for a Backup Server. If you remove this entry, you cannot back up your databases.
- Adaptive Server requires that the Backup Server have an *lname* of SYB_BACKUP. If you do not want to use that as the name of your Backup Server, or if you have more than one Backup Server running on your system, modify the *pname* for server SYB_BACKUP with sp_addserver so that Adaptive Server can communicate with Backup Server for database dumps and loads.

- If you specify an *lname*, *pname* and *class* that already exist in *syssservers*, *sp_addserver* prints an error message and does not update *syssservers*.
- Use *sp_serveroption* to set or clear server options.
- For information on using Component Integration Services, see the Component Integration Services User’s Guide.

Permissions

Only a System Security Officer can execute *sp_addserver*.

See also

System procedures – *sp_addremotelogin*, *sp_dropremotelogin*, *sp_dropserver*, *sp_helpremotelogin*, *sp_helpserver*, *sp_serveroption*

sp_addthreshold

Description

Creates a threshold to monitor space on a database segment. When free space on the segment falls below the specified level, Adaptive Server executes the associated stored procedure.

Syntax

sp_addthreshold dbname, segname, free_space, proc_name

Parameters

dbname

– is the database for which to add the threshold. This must be the name of the current database.

segname

– is the segment for which to monitor free space. Use quotes when specifying the “default” segment.

free_space

– is the number of free pages at which the threshold is crossed. When free space in the segment falls below this level, Adaptive Server executes the associated stored procedure.

proc_name

– is the stored procedure to be executed when the amount of free space on *segname* drops below *free_space*. The procedure can be located in any database on the current Adaptive Server or on an Open Server. Thresholds cannot execute procedures on remote Adaptive Servers.

Examples

Example 1

```
sp_addthreshold mydb, segment1, 200, pr_warning
```

Creates a threshold for segment1. When the free space on segment1 drops below 200 pages, Adaptive Server executes the procedure pr_warning.

Example 2

```
sp_addthreshold userdb, user_data, 100,  
"o_server...mail_me"
```

Creates a threshold for the user_data segment. When the free space on user_data falls below 100 pages, Adaptive Server executes a remote procedure call to the Open Server mail_me procedure.

Example 3

```
pubs2..sp_addthreshold pubs2, indexes, 100,  
pr_warning
```

Creates a threshold on the indexes segment of the pubs2 database. You can issue this command from any database.

Usage

- For more information about using thresholds, see the System Administration Guide.

Crossing a threshold

- When a threshold is crossed, Adaptive Server executes the associated stored procedure. Adaptive Server uses the following search path for the threshold procedure:
 - If the procedure name does not specify a database, Adaptive Server looks in the database in which the threshold was crossed.
 - If the procedure is not found in this database, and the procedure name begins with “sp_”, Adaptive Server looks in the subsystemprocs database.

If the procedure is not found in either database, Adaptive Server sends an error message to the error log.

- Adaptive Server uses a *hysteresis value*, the global variable @@thresh_hysteresis, to determine how sensitive thresholds are to variations in free space. Once a threshold executes its procedure, it is deactivated. The threshold remains inactive until the amount of free space in the segment rises to @@thresh_hysteresis pages above the threshold. This prevents thresholds from executing their procedures repeatedly in response to minor fluctuations in free space.

The last-chance threshold

- By default, Adaptive Server monitors the free space on the segment where the log resides and executes *sp_thresholdaction* when the amount of free space is less than that required to permit a successful dump of the transaction log. This amount of free space, called the *last-chance threshold*, is calculated by Adaptive Server and cannot be changed by users.
- If the last-chance threshold is crossed before a transaction is logged, Adaptive Server suspends the transaction until log space is freed. Use *sp_dboption* to change this behavior for a particular database. *sp_dboption* "abort tran on log full", true causes Adaptive Server to roll back all transactions that have not yet been logged when the last-chance threshold is crossed.
- Only databases that store their logs on a separate segment can have a last-chance threshold. Use *sp_logdevice* to move the transaction log to a separate device.

Creating additional thresholds

- Each database can have up to 256 thresholds, including the last-chance threshold.
- When you add a threshold, it must be at least 2 times @@thresh_hysteresis pages from the closest threshold.

Creating threshold procedures

- Any user with create procedure permission can create a threshold procedure in a database. Usually, a System Administrator creates *sp_thresholdaction* in the *sybsystemprocs* database, and the Database Owners create threshold procedures in user databases.
- *sp_addthreshold* does not verify that the specified procedure exists. It is possible to add a threshold before creating the procedure it executes.
- *sp_addthreshold* checks to ensure that the user adding the threshold procedure has been directly granted the "sa_role". All system roles active when the threshold procedure is created are entered in *systhresholds* as valid roles for the user writing the procedure. However, only directly granted system roles are activated when the threshold fires. Indirectly granted system roles and user-defined roles are not activated.
- Adaptive Server passes four parameters to a threshold procedure:

- @dbname, varchar(30), which identifies the database
- @segmentname, varchar(30), which identifies the segment
- @space_left, int, which indicates the number of free pages associated with the threshold
- @status, int, which has a value of 1 for last-chance thresholds and 0 for other thresholds

These parameters are passed by position rather than by name; your threshold procedure can use other names for them, but it must declare them in the order shown and with the correct datatypes.

- It is not necessary to create a different procedure for each threshold. To minimize maintenance, you can create a single threshold procedure in the sybsystemprocs database that is executed for all thresholds in Adaptive Server.
- Include print and raiserror statements in the threshold procedure to send output to the error log.

Executing threshold procedures

- Tasks initiated when a threshold is crossed execute as background tasks. These tasks do not have an associated terminal or user session. If you execute sp_who while these tasks are running, the status column shows “background”.
- Adaptive Server executes the threshold procedure with the permissions the user had at the time he or she added the threshold, minus any permissions that have since been revoked.
- Each threshold procedure uses one user connection, for as long as it takes for the procedure to execute.

Changing or deleting thresholds

- Use sp_helpthreshold for information about existing thresholds.
- Use sp_modifythreshold to associate a threshold with a new threshold procedure, free-space value, or segment. (You cannot change the free-space value or segment name associated with the last-chance threshold.)

Each time a user modifies a threshold, that user becomes the threshold owner. When the threshold is crossed, Adaptive Server executes the threshold with the permissions the owner had at the time he or she modified the threshold, minus any permissions that have since been revoked.

- Use *sp_droptreshold* to drop a threshold from a segment.

Disabling free-space accounting

Warning! System procedures cannot provide accurate information about space allocation when free-space accounting is disabled.

- Use the *no free space acctg* option of *sp_dboption* to disable free-space accounting on non-log segments.
- You cannot disable free-space accounting on log segments.

Permissions

Only the Database Owner or a System Administrator can execute *sp_addthreshold*.

See also

Commands – create procedure, dump transaction

Functions – *lct_admin*

System procedures – *sp_dboption*, *sp_droptreshold*, *sp_helpthreshold*, *sp_modifythreshold*, *sp_thresholdaction*

sp_add_time_range

Description

Adds a named time range to an Adaptive Server.

Syntax

sp_add_time_range *name*, *startday*, *endday*,
starttime, *endtime*

Parameters

name

– is the name of the time range. Time range names must be 30 characters or fewer. The name cannot already exist in the *systimeranges* system table of the master database.

startday

– is the day of the week on which the time range begins. This must be the full weekday name for the default server language, as stored in the *syslanguages* system table of the master database.

endday

– is the day of the week on which the time range ends. This must be the full weekday name for the default server language, as stored in the *syslanguages* system table of the master database. The *endday* can fall either earlier or later in the week than the *startday* or can be the same day as the *startday*.

starttime

– is the time of day when the time range begins. Specify the *starttime* in terms of a 24-hour clock, with a value between “00:00” (midnight) and “23:59” (11:59 p.m.). Use the following form:

"HH:MM"

endtime

– is the time of day when the time range ends. Specify the *endtime* in terms of a 24-hour clock, with a value between “00:00” (midnight) and “23:59” (11:59 p.m.). Use the following form:

"HH:MM"

Note To create a time range that spans the entire day, specify both a start time and an end time of “00:00”.

The *endtime* must occur later in the day than the *starttime*, unless *endtime* is “00:00”.

Examples

Example 1

```
sp_add_time_range business_hours, monday, Friday,  
"09:00", "17:00"
```

Creates the *business_hours* time range, which is active Monday through Friday, from 9:00 a.m. to 5:00 p.m.

Example 2

```
sp_add_time_range before_hours, Monday, Friday,  
"00:00", "09:00"
```

```
sp_add_time_range after_hours, Monday, Friday,  
"18:00", "00:00"
```

Creates two time ranges, *before_hours* and *after_hours*, that, together, span all non-business hours Monday through Friday. The *before_hours* time range covers the period from 12:00 midnight to 9:00 a.m., Monday through Friday. The *after_hours* time range covers the period from 6:00 p.m. through 12:00 midnight, Monday through Friday.

Example 3

```
sp_add_time_range weekends, Saturday, Sunday,  
"00:00", "00:00"
```


Creates the weekends time range, which is 12:00 midnight Saturday to 12:00 midnight Sunday.

Example 4

```
sp_add_time_range Fri_thru_Mon, Friday, Monday,  
"09:00", "17:00"
```

Creates the *Fri_thru_Mon* time range, which is 9:00 a.m. to 5:00 p.m., Friday, Saturday, Sunday, and Monday.

Example 5

```
sp_add_time_range Wednesday_night, Wednesday,  
Wednesday, "17:00", "00:00"
```

Creates the *Wednesday_night* time range, which is Wednesday from 5:00 p.m. to 12:00 midnight.

Usage

- Adaptive Server includes one named time range, the “at all times” time range. This time range covers all times, from the first day through the last of the week, from 00:00 through 23:59. It cannot be modified or deleted.
- Adaptive Server generates a unique ID number for each named time range and inserts it into the *systemranges* system table,
- When storing a time range in the *systemranges* system table, Adaptive Server converts its *startday* and *endday* values into integers. For servers with a default language of *us_english*, the week begins on Monday (day 1) and ends on Sunday (day 7).
- It is possible to create a time range that overlaps with one or more other time ranges.
- Range days are contiguous, so the days of the week can wrap around the end to the beginning of the week. In other words, Sunday and Monday are contiguous days, as are Tuesday and Wednesday.
- The active time ranges are bound to a session at the beginning of each query batch. A change in the server’s active time ranges due to a change in actual time has no effect on a session during the processing of a query batch. In other words, if a resource limit restricts a query batch during a given time range but a query batch begins before that time range becomes active, the query batch that is already running is not affected by the resource limit.

- The addition, modification, and deletion of time ranges using the system procedures does not affect the active time ranges for sessions currently in progress.
- If a resource limit has a transaction as its scope, and a change occurs in the server's active time ranges while a transaction is running, the newly active time range does not affect the transaction currently in progress.
- Changes to a resource limit that has a transaction as its scope does not affect any transactions currently in progress.
- For more information on time ranges, see the System Administration Guide.

Permissions

Only a System Administrator can execute sp_add_time_range.

See also

System procedures – sp_add_resource_limit, sp_drop_time_range, sp_modify_time_range

sp_addtype

Description

Creates a user-defined datatype.

Syntax

```
sp_addtype typename,  
          phystype [(length) | (precision [, scale])]  
          [, "identity" | nulltype]
```

Parameters

typename
– is the name of the user-defined datatype. Type names must conform to the rules for identifiers and must be unique in each database.

phystype

– is the physical or Adaptive Server-supplied datatype on which to base the user-defined datatype. You can specify any Adaptive Server datatype except timestamp.

The char, varchar, unichar, univarchar, nchar, nvarchar, binary, and varbinary datatypes expect a *length* in parentheses. If you do not supply one, Adaptive Server uses the default length of 1 character.

The float datatype expects a binary *precision* in parentheses. If you do not supply one, Adaptive Server uses the default precision for your platform.

The numeric and decimal datatypes expect a decimal *precision* and *scale*, in parentheses and separated by a comma. If you do not supply them, Adaptive Server uses a default precision of 18 and a scale of 0.

Enclose physical types that include punctuation, such as parentheses or commas, within single or double quotes.

identity

– indicates that the user-defined datatype has the IDENTITY property. Enclose the identity keyword within single or double quotes. You can specify the IDENTITY property only for numeric datatypes with a scale of 0.

IDENTITY columns store sequential numbers, such as invoice numbers or employee numbers, that are generated by Adaptive Server. The value of the IDENTITY column uniquely identifies each row in a table. IDENTITY columns are not updatable and do not allow null values.

nulltype

– indicates how the user-defined datatype handles null value entries. Acceptable values for this parameter are null, NULL, nonull, NONULL, "not null", and "NOT NULL". Any *nulltype* that includes a blank space must be enclosed in single or double quotes.

If you omit both the IDENTITY property and the *nulltype*, Adaptive Server creates the datatype using the null mode defined for the database. By default, datatypes for which no *nulltype* is specified are created NOT NULL (that is, null values are not allowed and explicit entries are required). For compliance to the SQL standards, use the *sp_dboption* system procedure to set the allow nulls by default option to true. This changes the database's null mode to NULL.

Examples

Example 1

```
sp_addtype ssn, "varchar(11)"
```

Creates a user-defined datatype called `ssn` to be used for columns that hold social security numbers. Since the *nulltype* parameter is not specified, Adaptive Server creates the datatype using the database's default null mode. Notice that `varchar(11)` is enclosed in quotation marks, because it contains punctuation (parentheses).

Example 2

```
sp_addtype birthday, "datetime", null
```

Creates a user-defined datatype called `birthday` that allows null values.

Example 3

```
sp_addtype temp52, "numeric(5,2)"
```

Creates a user-defined datatype called `temp52` used to store temperatures of up to 5 significant digits with 2 places to the right of the decimal point.

Example 4

```
sp_addtype "row_id", "numeric(10,0)", "identity"
```

Creates a user-defined datatype called `row_id` with the `IDENTITY` property, to be used as a unique row identifier. Columns created with this datatype store system-generated values of up to 10 digits in length.

Example 5

```
sp_addtype systype, sysname
```

Creates a user-defined datatype with an underlying type of `sysname`. Although you cannot use the `sysname` datatype in a `create table`, `alter table`, or `create procedure` statement, you can use a user-defined datatype that is based on `sysname`.

Usage

- `sp_addtype` creates a user-defined datatype and adds it to the `systypes` system table. Once a user-defined datatype is created, you can use it in `create table` and `alter table` statements and bind defaults and rules to it.
- Build each user-defined datatype in terms of one of the Adaptive Server-supplied datatypes, specifying the length or the precision and scale, as appropriate. You cannot override the length, precision, or scale in a `create table` or `alter table` statement.

- A user-defined datatype name must be unique in the database, but user-defined datatypes with different names can have the same definitions.
- If `nchar` or `nvarchar` is specified as the *phystype*, the maximum length of columns created with the new type is the length specified in `sp_addtype` multiplied by the value of `@@ncharsize` at the time the type was added.
- If `unichar` or `univarchar` is specified as the *phystype*, the maximum length of columns created with the new type is the length specified in `sp_addtype` multiplied by the value of `2` at the time the type was added.
- Each system type has a *hierarchy*, stored in the `systypes` system table. User-defined datatypes have the same datatype hierarchy as the physical types on which they are based. In a mixed-mode expression, all types are converted to a common type, the type with the lowest hierarchy.

Use the following query to list the hierarchy for each system-supplied and user-defined type in your database:

```
select name, hierarchy
from systypes
order by hierarchy
```

Datatypes with the `IDENTITY` property

- If a user-defined datatype is defined with the `IDENTITY` property, all columns created from it are `IDENTITY` columns. You can specify `IDENTITY`, `NOT NULL`, or neither in the create or alter table statement. Following are three different ways to create an `IDENTITY` column from a user-defined datatype with the `IDENTITY` property:

```
create table new_table (id_col IdentType)
create table new_table (id_col IdentType
identity)
create table new_table (id_col IdentType not
null)
```

- When you create a column with the create table or alter table statement, you can override the null type specified with the `sp_addtype` system procedure:
 - Types specified as `NOT NULL` can be used to create `NULL` or `IDENTITY` columns.

- Types specified as NULL can be used to create NOT NULL columns, but not to create IDENTITY columns.

Note If you try to create a null column from an IDENTITY type, the create or alter table statement fails.

Permissions

Any user can execute sp_addtype.

See also

Commands – create default, create rule, create table

Datatypes – User-defined datatypes

System procedures – sp_bindefault, sp_bindrule, sp_dboption, sp_droptype, sp_rename, sp_unbindefault, sp_unbindrule

sp_addumpdevice

Description

Adds a dump device to Adaptive Server.

Syntax

```
sp_addumpdevice {"tape" | "disk"}, logicalname,  
physicalname [, tapesize]
```

Parameters

"tape"

– for tape drives. Enclose tape in quotes.

"disk"

– is for a disk or a file device. Enclose disk in quotes.

logicalname

– is the “logical” dump device name. It must be a valid identifier. Once you add a dump device to sysdevices, you can specify its logical name in the load and dump commands.

physicalname

– is the physical name of the device. You can specify either an absolute path name or a relative path name. During dumps and loads, the Backup Server resolves relative path names by looking in Adaptive Server’s current working directory. Enclose names containing non-alphanumeric characters in quotation marks. For UNIX platforms, specify a non-rewinding tape device name.

tapesize

– is the capacity of the tape dump device, specified in megabytes. OpenVMS systems ignore the *tapesize* parameter if specified. Other platforms require this parameter for tape devices but ignore it for disk devices. The *tapesize* should be at least five database pages (each page requires 2048 bytes). Sybase recommends that you specify a capacity that is slightly below the rated capacity for your device.

Examples

```
sp_addumpdevice "tape", mytapedump, "/dev/nrmt8", 40
```

Adds a 40MB tape device. Dump and load commands can reference the device by its physical name, */dev/nrmt8*, or its logical name, *mytapedump*.

```
sp_addumpdevice "disk", mydiskdump,  
"/dev/rxyl1d/dump.dat"
```

Adds a disk device named *mydiskdump*. Specify an absolute or relative path name and a file name.

Usage

- *sp_addumpdevice* adds a dump device to the *master.dbo.sysdevices* table. Tape devices are assigned a *cntrltpe* of 3; disk devices are assigned a *cntrltpe* of 2.
- To use an operating system file as a dump device, specify a device of type *disk* and an absolute or relative path name for the *physicalname*. Omit the *tapesize* parameter. If you specify a relative path name, dumps are made to—or loaded from—the current Adaptive Server working directory at the time the dump or load command executes.
- Ownership and permission problems can interfere with the use of disk or file dump devices. *sp_addumpdevice* adds the device to the *sysdevices* table, but does not guarantee that you can create a file as a dump device or that users can dump to a particular device.
- The *with capacity = megabytes* clause of the dump database and dump transaction commands can override the *tapesize* specified with *sp_addumpdevice*. On platforms that do not reliably detect the end-of-tape marker, the Backup Server issues a volume change request after the specified number of megabytes have been dumped.
- When a dump device fails, use *sp_dropdevice* to drop it from *sysdevices*. After replacing the device, use *sp_addumpdevice* to associate the logical device name with the new physical device. This avoids updating backup scripts and threshold procedures each time a dump device fails.
- To add database devices to *sysdevices*, use the *disk init* command.

Permissions Only a System Administrator can execute sp_addumpdevice.
See also *Commands* – disk init, dump database, dump transaction, load database, load transaction
System procedures – sp_dropdevice, sp_helpdevice

sp_adduser

Description Adds a new user to the current database.
Syntax `sp_adduser loginname [, name_in_db [, grpname]]`
Parameters
loginname
– is the user’s name in master.dbo.syslogins.
name_in_db
– is a new name for the user in the current database.
grpname
– adds the user to an existing group in the database.

Examples **Example 1**
`sp_adduser margaret`
Adds “margaret” to the database. Her database user name is the same as her Adaptive Server login name, and she belongs to the default group, “public”.

Example 2
`sp_adduser haroldq, harold, fort_mudge`
Adds “haroldq” to the database. When “haroldq” uses the current database, his name is “harold.” He belongs to the fort_mudge group, as well as to the default group “public”.

Usage

- The Database Owner executes sp_adduser to add a user name to the sysusers table of the current database, enabling the user to access the current database under his or her own name.
- Specifying a *name_in_db* parameter gives the new user a name in the database that is different from his or her login name in Adaptive Server. The ability to assign a user a different name is provided as a convenience. It is not an alias, as provided by sp_addalias, since it is not mapped to the identity and privileges of another user.

- A user and a group cannot have the same name.
- A user can be a member of only one group other than the default group, “public”. Every user is a member of the default group, “public”. Use *sp_changegroup* to change a user’s group.
- In order to access a database, a user must either be listed in *sysusers* (with *sp_adduser*) or mapped to another user in *sysalternates* (with *sp_addalias*), or there must be a “guest” entry in *sysusers*.

Permissions

Only the Database Owner, a System Administrator, or a System Security Officer can execute *sp_adduser*.

See also

Commands – grant, revoke, use

System procedures – *sp_addalias*, *sp_addgroup*, *sp_changegroup*, *sp_dropalias*, *sp_dropgroup*, *sp_helpuser*

Procedures: *sp_altermessage* – *sp_autoconnect*

sp_altermessage

Description	Enables and disables the logging of a system-defined or user-defined message in the Adaptive Server error log.
Syntax	<code>sp_altermessage message_id, parameter, parameter_value</code>
Parameters	<p><i>message_id</i> – is the message number of the message to be altered. This is the number of the message as it is recorded in the error column in the <code>sysmessages</code> or <code>sysusermessages</code> system table.</p> <p><i>parameter</i> – is the message parameter to be altered. The maximum length is 30 bytes. The only valid parameter is <code>with_log</code>.</p> <p><i>parameter_value</i> – is the new value for the parameter specified in <i>parameter</i>. The maximum length is 5 bytes. Values are <code>true</code> and <code>false</code>.</p>
Examples	<pre>sp_altermessage 2000, 'with_log', 'TRUE'</pre> <p>Specifies that message number 2000 in <code>sysmessages</code> should be logged in the Adaptive Server error log and also in the Windows NT Event Log (if logging is enabled).</p>
Usage	<ul style="list-style-type: none"> • If the <i>parameter_value</i> is <code>true</code>, the specified message is always logged. If it is <code>false</code>, the default logging behavior is used; the message may or may not be logged, depending on the severity of the error and other factors. Setting the <i>parameter_value</i> to <code>false</code> produces the same behavior that would occur if <code>sp_altermessage</code> had not been called. • On Windows NT servers, <code>sp_altermessage</code> also enables and disables logging in the Windows NT Event Log.
Permissions	Only the Database Owner or a System Administrator can execute <code>sp_altermessage</code> .
See also	<i>System procedures</i> – <code>sp_addmessage</code> , <code>sp_dropmessage</code>

sp_audit

Description Allows a System Security Officer to configure auditing options.

Syntax sp_audit *option*, *login_name*, *object_name* [, *setting*]

Parameters *option*
 – is the name of the auditing option to set. Table 5-1 lists the valid auditing options.

Table 5-1: Auditing options

Option	Description
adhoc	Allows users to use sp_addauditrecord to add their own user-defined audit records to the audit trail
all	Audits all actions performed by a particular user or by users with a particular role Note Auditing all actions does not affect whether users can add ad hoc audit records.
alter	Audits the execution of the alter table or alter database commands
bcp	Audits the execution of the bcp in utility
bind	Audits the execution of sp_bindefault, sp_bindmsg, and sp_bindrule system procedures
cmdtext	Audits all actions of a particular user, or by users with a particular role.
create	Audits the creation of database objects
dbaccess	Audits access to the current database from another database
dbcc	Audits the execution of any dbcc command
delete	Audits the deletion of rows from a table or view
disk	Audits the execution of disk init, disk refit, disk reinit, disk mirror, disk unmirror, and disk remirror
drop	Audits the dropping of database objects
dump	Audits the execution of dump database or dump transaction commands
errors	Audits errors, whether fatal or not
exec_procedure	Audits the execution of a stored procedure
exec_trigger	Audits the execution of a trigger
func_dbaccess	Audits access to a database via a Transact-SQL function
func_obj_access	Audits access to a database object via a Transact-SQL function
grant	Audits the execution of the grant command
insert	Audits the insertion of rows into a table or view
load	Audits the execution of the load database or load transaction commands
login	Audits all login attempts into Adaptive Server
logout	Audits all logout attempts from Adaptive Server
reference	Audits references between tables.
revoke	Audits the execution of the revoke command

Option	Description
<i>rpc</i>	Audits the execution of remote procedure calls
<i>security</i>	Audits the following security-relevant events: <ul style="list-style-type: none"> • Starting up or shutting down the server • Activating or deactivating a role • Issuing any of the following commands: <ul style="list-style-type: none"> • <i>connect</i> • <i>kill</i> • <i>online database</i> • <i>set proxy</i> • <i>set session authorization</i> • <i>sp_configure</i> • Using any of the following functions: <ul style="list-style-type: none"> • <i>valid_user</i> • <i>proc_role</i> (from within a system procedure) • Regenerating the SSO passwords
<i>select</i>	Audits the execution of the <i>select</i> command
<i>setuser</i>	Audits the execution of the <i>setuser</i> command
<i>table_access</i>	Audits access to any table by a specific user
<i>truncate</i>	Audits the execution of the <i>truncate table</i> command
<i>unbind</i>	Audits the execution of the <i>sp_unbindrule</i> , <i>sp_unbindmsg</i> , and <i>sp_unbindefault</i> system procedures
<i>update</i>	Audits updates to rows in a table or view
<i>view_access</i>	Audits access to any view by a specific user

login_name

– is the name of a specific login to be audited. To audit all logins, specify *all* for the *option* parameter. If you specify *all*, you can set the *login_name* parameter to a specific system role to audit all actions by users with that system role active. You cannot specify a user-defined role for *login_name*. For more information on the *login_name* values that are valid with each *option* value, see the System Administration Guide.

object_name

– is the name of the object to be audited. Valid values, depending on the value you specified for *option*, are:

- The object name, including the owner’s name if you do not own the object. For example, to audit a table named inventory that is owned by Joe, you would specify joe.inventory for *object_name*.
- all for all objects.
- default table, default view, default procedure, or default trigger to audit access to any new table, view, procedure, or trigger.

default table and default view are valid values for *object_name* when you specify delete, insert, select, or update for the *option* parameter. default procedure is valid when you specify the *exec_procedure* option. default trigger is valid when you specify the *exec_trigger* option.

For more information about the *object_name* values that are valid with each *option* value, see the System Administration Guide.

setting

– is the level of auditing. If you do not specify a value for *setting*, Adaptive Server displays the current auditing setting for the option. Valid values for the *setting* parameter are described in the following table:

setting value	Description
on	Activates auditing for the specified option. Adaptive Server generates audit records for events controlled by this option, whether the event passes or fails permission checks.
off	Deactivates auditing for the specified option.
pass	Activates auditing for events that pass permission checks.
fail	Deactivates auditing for events that fail permission checks.

If you specify pass for an option and later specify fail for the same option, or vice versa, the result is equivalent to specifying on. Adaptive Server generates audit records regardless of whether events pass or fail permission checks. Settings of on or off apply to all auditing options. Settings of pass and fail apply to all options except errors and adhoc. For these options, only on or off applies. The initial, default value of all options is off.

Examples

Example 1

```
sp_audit "security", "all", "all", "on"
```

Initiates auditing for security-relevant events. Both successful and failed events are audited.

Example 2

```
sp_audit "security", "all", "all"
```

Displays the setting of the security auditing option.

Example 3

```
sp_audit "create", "all", master, "on"
```

Initiates auditing for the creation of objects in the masterdatabase, including create database.

Example 4

```
sp_audit "create", "all", db1, "on"
```

Initiates auditing for the creation of all objects in the db1database.

Example 5

```
sp_audit "all", "sa_role", "all", "fail"
```

Initiates auditing for all failed executions by a System Administrator.

Example 6

```
sp_audit "update", "all", "default table", "on"
```

Initiates auditing for all updates to future tables in the current database. For example, if the current database is utility, all new tables created in utility will be audited for updates. The auditing for existing tables is not affected.

Usage

- `sp_audit` determines what will be audited when auditing is enabled. No actual auditing takes place until you use `sp_configure` to set the auditing parameter to on. Then, all auditing options that have been configured with `sp_audit` take effect. For more information, see `sp_configure`.
- If you are not the owner of the object being specified, qualify the *object_name* parameter value with the owner's name, in the following format:

```
"ownername.objname"
```

- You cannot activate default auditing for the following options in the tempdb database:
 - delete
 - insert

- select
 - update
 - exec_procedure
 - exec_trigger
- Table 5-2 lists the configuration parameters that control auditing.

Table 5-2: Configuration parameters that control auditing

Configuration Parameter	Effect
auditing	Enables or disables auditing for the server.
audit_queue_size	Establishes the size of the audit queue.
current_audit_table	Sets the current audit table. Adaptive Server writes all audit records to that table.
suspend_auditing_when_full	Controls the behavior of the audit process when an audit device becomes full.

The auditing, current_audit_table, and suspend_auditing_when_full configuration parameters are dynamic and take effect immediately. Because audit_queue_size affects memory allocation, the parameter is static and does not take effect until Adaptive Server is restarted.

- For more information about configuring Adaptive Server for auditing, see sp_configure in the System Administration Guide.

Permissions

Only a System Security Officer can execute sp_audit.

See also

System procedures – sp_addauditrecord, sp_configure

Utility commands – bcp

sp_autoconnect

Description

Component Integration Services only – Defines a passthrough connection to a remote server for a specific user, which allows the named user to enter passthrough mode automatically at login.

Syntax

```
sp_autoconnect server, {true|false}  
[, loginame]
```


Parameters	<p><i>server</i></p> <p>– is the name of a server to which an automatic passthrough connection is made. <i>server</i> must be the name of a remote server already added by <i>sp_addserver</i>. This server cannot be the local server.</p> <p>true false</p> <p>– determines whether the automatic passthrough connection is enabled or disabled for <i>server</i>. true enables the automatic connection. false disables it.</p> <p><i>loginame</i></p> <p>– specifies the name of the user for which automatic connection is required. If no <i>loginame</i> is supplied, the autoconnect status is modified for the current user.</p>
Examples	<p>Example 1</p> <pre>sp_autoconnect SYBASE, true</pre> <p>The current user is automatically connected to the server SYBASE the next time that user logs in. The user’s connection is placed in passthrough mode.</p> <p>Example 2</p> <pre>sp_autoconnect SYBASE, false, steve</pre> <p>Disables the autoconnect feature for the user “steve”.</p>
Usage	<ul style="list-style-type: none">• <i>sp_autoconnect</i> defines a passthrough connection to a remote server for a specific user, which allows the named user to enter passthrough mode automatically at login.• The System Administrator must grant connect to permission to the login prior to executing <i>sp_autoconnect</i>.• Use <i>sp_autoconnect</i> only when Component Integration Services is installed and configured.• Do not change the autoconnect status of the “sa” login account.• Changing the autoconnect status does not occur immediately for users who are currently connected. They must disconnect from the local server, then reconnect before the change is made.• Use <i>disconnect</i> to exit passthrough mode.
Permissions	Only a System Administrator can execute <i>sp_autoconnect</i> .

See also

Commands – connect to...disconnect, grant

System procedures – sp_addlogin, sp_addserver, sp_passthru,
sp_remotesql

Procedures: sp_bindcache – sp_bindrule

sp_bindcache

Description	Binds a database, table, index, text object, or image object to a data cache.
Syntax	<code>sp_bindcache <i>cachename</i>, <i>dbname</i> [, [<i>ownername</i>.]<i>tablename</i> [, <i>indexname</i> "text only"]]</code>
Parameters	<p><i>cachename</i> – is the name of an active data cache.</p> <p><i>dbname</i> – is the name of the database to be bound to the cache or the name of the database containing the table, index, text or image object to be bound to the cache.</p> <p><i>ownername</i> – is the name of the table’s owner. If the table is owned by “dbo”, the owner name is optional.</p> <p><i>tablename</i> – is the name of the table to be bound to the cache, or the name of the table whose index, text object, or image object is to be bound to the cache.</p> <p><i>indexname</i> – is the name of the index to be bound to the cache.</p> <p>text only – binds text or image objects to a cache. When this parameter is used, you cannot give an index name at the same time.</p>

Examples

Example 1

```
sp_bindcache pub_cache, pubs2, titles
```

Binds the titles table to the cache named pub_cache.

Example 2

```
sp_bindcache pub_ix_cache, pubs2, titles,
```

```
title_id_cix
```

Binds the clustered index titles.title_id_cix to the pub_ix_cache.

Example 3

```
sp_bindcache tempdb_cache, tempdb
```

Binds tempdb to the tempdb_cache.

Example 4

```
sp_bindcache logcache, pubs2, syslogs
```

Binds the pubs2 transaction log, syslogs, to the cache named logcache.

Example 5

```
sp_bindcache pub_cache, pubs2, au_pix, "text only"
```

Binds the image chain for the au_pix table to the cache named pub_cache.

Usage

- A database or database object can be bound to only one cache. You can bind a database to one cache and bind individual tables, indexes, text objects, or image objects in the database to other caches. The database binding serves as the default binding for all objects in the database that have no other binding. The data cache hierarchy for a table or index is as follows:
 - If the object is bound to a cache, the object binding is used.
 - If the object is not bound to a cache, but the object's database is bound to a cache, the database binding is used.
 - If neither the object nor its database is bound to a cache, the default data cache is used.
- The cache and the object or database being bound to it must exist before you can execute sp_bindcache. Create a cache with sp_cacheconfig and restart Adaptive Server before binding objects to the cache.
- Cache bindings take effect immediately, and do not require a restart of the server. When you bind an object to a data cache:
 - Any pages for the object that are currently in memory are cleared.
 - When the object is used in queries, its pages are read into the bound cache.

- You can bind an index to a different cache than the table it references. If you bind a clustered index to a cache, the binding affects only the root and intermediate pages of the index. It does not affect the data pages (which are, by definition, the leaf pages of the index).
- To bind a database, you must be using the master database. To bind tables, indexes, text objects, or image objects, you must be using the database where the objects are stored.
- To bind any system tables in a database, you must be using the database and the database must be in single-user mode. Use the command:

```
sp_dboption db_name, "single user", true
```

For more information, see *sp_dboption*.

- You do not have to unbind objects or databases in order to bind them to a different cache. Issuing *sp_bindcache* on an object that is already bound drops the old binding and creates the new one.
- *sp_bindcache* needs to acquire an exclusive table lock when you are binding a table or its indexes to a cache so that no pages can be read while the binding is taking place. If a user holds locks on a table, and you issue *sp_bindcache* on that object, the task doing the binding sleeps until the locks are released.
- When you bind or unbind an object, all stored procedures that reference the object are recompiled the next time they are executed. When you change the binding for a database, all stored procedures that reference objects in the bound database are recompiled the next time they are executed.
- When you drop a table, index, or database, all associated cache bindings are dropped. If you re-create the table, index, or database, you must use *sp_bindcache* again if you want it bound to a cache.
- If a database or a database object is bound to a cache, and the cache is dropped, the cache bindings are marked invalid, but remain stored in the *sysattributes* system table(s). Warnings are printed in the error log when Adaptive Server is restarted. If a cache of the same name is created, the bindings become valid when Adaptive Server is restarted.
- The following procedures provide information about the bindings for their respective objects: *sp_helpdb* for databases, *sp_help* for tables, and *sp_helpindex* for indexes. *sp_helpcache* provides information about all objects bound to a particular cache.

- Use `sp_spaceused` to see the current size of tables and indexes, and `sp_estspace` to estimate the size of tables that you expect to grow. Use `sp_cacheconfig` to see information about cache size and status, and to configure and reconfigure caches.

Restrictions

- The master database, the system tables in master, and the indexes on the system tables in master cannot be bound to a cache. You can bind non-system tables from master, and their indexes, to caches.
- You cannot bind a database or an object to a cache if:
 - Isolation level 0 reads are active on the table
 - The task doing the binding currently has a cursor open on the table
- If a cache has the type log only, you can bind a syslogs table only to that cache. Use `sp_cacheconfig` to see a cache's type.

Permissions

Only a System Administrator can execute `sp_bindcache`.

See also

System procedures – `sp_cacheconfig`, `sp_configure`, `sp_help`, `sp_helpcache`, `sp_helppdb`, `sp_helpindex`, `sp_poolconfig`, `sp_unbindcache`, `sp_unbindcache_all`

sp_bindefault

Description

Binds a user-defined default to a column or user-defined datatype.

Syntax

`sp_bindefault defname, objname [, futureonly]`

Parameters

defname

– is the name of a default created with `create default` statements to bind to specific columns or user-defined datatypes.

objname

– is the name of the table and column, or user-defined datatype, to which the default is to be bound. If the *objname* parameter is not of the form “*table.column*”, it is assumed to be a user-defined datatype. If the object name includes embedded blanks or punctuation, or is a reserved word, enclose it in quotation marks.

Existing columns of the user-defined datatype inherit the default *defname*, unless you specify `futureonly`.

futureonly

- prevents existing columns of a user-defined datatype from acquiring the new default. This parameter is optional when you are binding a default to a user-defined datatype. It is never used to bind a default to a column.

Examples

Example 1

```
sp_bindefault today, "employees.startdate"
```

Assuming that a default named `today` has been defined in the current database with `create default`, this command binds it to the `startdate` column of the `employees` table. Each new row added to the `employees` table has the value of the `today` default in the `startdate` column, unless another value is supplied.

Example 2

```
sp_bindefault def_ssn, ssn
```

Assuming that a default named `def_ssn` and a user-defined datatype named `ssn` exist, this command binds `def_ssn` to `ssn`. The default is inherited by all columns that are assigned the user-defined datatype `ssn` when a table is created. Existing columns of type `ssn` also inherit the default `def_ssn`, unless you specify `futureonly` (which prevents existing columns of that user-defined datatype from inheriting the default), or unless the column's default has previously been changed (in which case the changed default is maintained).

Example 3

```
sp_bindefault def_ssn, ssn, futureonly
```

Binds the default `def_ssn` to the user-defined datatype `ssn`. Because the `futureonly` parameter is included, no existing columns of type `ssn` are affected.

Usage

- You can create column defaults in two ways: by declaring the default as a column constraint in the `create table` or `alter table` statement or by creating the default using the `create default` statement and binding it to a column using `sp_bindefault`. Using `create default`, you can bind that default to more than one column in the database.
- You cannot bind a default to an Adaptive Server-supplied datatype.
- You cannot bind a default to a system table.

- Defaults bound to a column or user-defined datatype with the IDENTITY property have no effect on column values. Each time you insert a row into the table, Adaptive Server assigns the next sequential number to the IDENTITY column.
- If binding a default to a column, give the *objname* argument in the form “*table.column*”. Any other format is assumed to be the name of a user-defined datatype.
- If a default already exists on a column, you must remove it before binding a new default. Use *sp_unbinddefault* to remove defaults created with *sp_bindefault*. To remove defaults created with *create table* or *alter table*, use *alter table* to replace the default with NULL.
- Existing columns of the user-defined datatype inherit the new default unless you specify *futureonly*. New columns of the user-defined datatype always inherit the default. Binding a default to a user-defined datatype overrides defaults bound to columns of that type; to restore column bindings, *unbind* and *rebind* the column default.
- Statements that use a default cannot be in the same batch as their *sp_bindefault* statement.

Permissions

Only the object owner can execute *sp_bindefault*.

See also

Commands – *create default*, *create table*, *drop default*

System procedures – *sp_unbinddefault*

sp_bindexclass

Description

Associates an execution class with a client application, login, or stored procedure.

Syntax

sp_bindexclass "*object_name*", "*object_type*", "*scope*", "*classname*"

Parameters

object_name

– is the name of the client application, login, or stored procedure to be associated with the execution class, *classname*.

object_type

– identifies the type of *object_name*. Use *ap* for application, *lg* for login, or *pr* for stored procedure.

scope

– is the name of a client application or login, or it can be NULL for ap and lg objects. It is the name of the stored procedure owner (user name) for objects. When the object with object_name interacts with the application or login, classname attributes apply for the scope you set.

classname

– specifies the type of class to associate with object_name. Values are:

- EC1, EC2, or EC3
- The name of a user-defined execution class
- ANYENGINE

Examples

Example 1

```
sp_bindexeclass 'isql', 'ap', NULL, 'EC3'
```

This statement specifies that Transact-SQL applications will execute with EC3 attributes for any login or application process (because the value of *scope* is NULL) that invokes isql, unless the login or application is bound to a higher execution class.

Example 2

```
sp_bindexeclass 'sa', 'lg', 'isql', 'EC1'
```

This statement specifies that when a login with the System Administrator role executes Transact-SQL applications, the login process executes with EC1 attributes. If you have already executed the statement in the first example, then any other login or client application that invokes isql will execute with EC3 attributes.

Example 3

```
sp_bindexeclass 'my_proc', 'PR', 'kundu', 'EC3'
```

This statement assigns EC3 attributes to the stored procedure named my_proc owned by user kundu.

Usage

- *sp_bindexeclass* associates an execution class with a client application, login, or stored procedure. Create execution classes with *sp_addexeclass*.

- When scope is NULL, object_name has no scope. classname's execution attributes apply to all of its interactions. For example, if object_name is an application name, the attributes apply to any login process that invokes the application. If object_name is a login name, the attributes apply to a particular login process for any application invoked by the login process.
- When binding a stored procedure to an execution class, you must use the name of the stored procedure owner (user name) for the scope parameter. This narrows the identity of a stored procedure when there are multiple invocations of it in the same database.
- Due to precedence and scoping rules, the execution class being bound may or may not have been in effect for the object called object_name. The object automatically binds itself to another execution class, depending on other binding specifications, precedence, and scoping rules. If no other binding is applicable, the object binds to the default execution class, EC2.
- Binding fails when you attempt to bind an active process to an engine group with no online engines.
- Adaptive Server creates a row in the sysattributes table containing the object ID and user ID in the row that stores data for the binding.
- A stored procedure must exist before it can be bound.
- Stored procedure bindings must be done in the database in which the stored procedure resides. Therefore, when binding system procedures, execute sp_bindexeclass from within the sybsystemprocs database.
- Only the “priority attribute” of the execution class is used when you bind the class to a stored procedure.
- The name of the owner of a stored procedure must be supplied as the scope parameter when you are binding a stored procedure to an execution class. This helps to uniquely identify a stored procedure when multiple stored procedures with the same name (but different owners) exist in the database.

Permissions

Only a System Administrator can execute sp_bindexeclass.

See also

System procedures – sp_addexeclass, sp_showexeclass, sp_unbindexeclass

Utility – isql

sp_bindmsg

Description	Binds a user message to a referential integrity constraint or check constraint.
Syntax	<code>sp_bindmsg <i>constrname</i>, <i>msgid</i></code>
Parameters	<p><i>constrname</i></p> <ul style="list-style-type: none"> – is the name of the integrity constraint to which you are binding a message. Use the constraint clause of the create table command, or the add constraint clause of the alter table command to create and name constraints. <p><i>msgid</i></p> <ul style="list-style-type: none"> – is the number of the user message to be bound to an integrity constraint. The message must exist in the sysusermessages table in the local database prior to calling <code>sp_bindmsg</code>.
Examples	<pre>sp_bindmsg positive_balance, 20100</pre> <p>Binds user message number 20100 to the <code>positive_balance</code> constraint.</p>
Usage	<ul style="list-style-type: none"> • <code>sp_bindmsg</code> binds a user message to an integrity constraint by adding the message number to the constraint row in the <code>sysconstraints</code> table. • Only one message can be bound to a constraint. To change the message for a constraint, just bind a new message. The new message number replaces the old message number in the <code>sysconstraints</code> table. • You cannot bind a message to a unique constraint because a unique constraint does not have a constraint row in <code>sysconstraints</code> (a unique constraint is a unique index). • Use the <code>sp_addmessage</code> procedure to insert user messages into the <code>sysusermessages</code> table. • The <code>sp_getmessage</code> procedure retrieves message text from the <code>sysusermessages</code> table. • <code>sp_help <i>tablename</i></code> displays all constraint names declared on <i>tablename</i>.
Permissions	Only the object owner can execute <code>sp_bindmsg</code> .
See also	<p><i>Commands</i> – alter table, create table</p> <p><i>System procedures</i> – <code>sp_addmessage</code>, <code>sp_getmessage</code>, <code>sp_unbindmsg</code></p>

sp_bindrule

Description	Binds a rule to a column or user-defined datatype.
Syntax	sp_bindrule <i>rulename</i> , <i>objname</i> [, futureonly]
Parameters	<p><i>rulename</i></p> <ul style="list-style-type: none">– is the name of a rule. Create rules with create rule statements and bind rules to specific columns or user-defined datatypes with sp_bindrule. <p><i>objname</i></p> <ul style="list-style-type: none">– is the name of the table and column, or user-defined datatype, to which the rule is to be bound. If <i>objname</i> is not of the form “<i>table.column</i>”, it is assumed to be a user-defined datatype. If the object name has embedded blanks or punctuation, or is a reserved word, enclose it in quotation marks. <p>futureonly</p> <ul style="list-style-type: none">– prevents existing columns of a user-defined datatype from inheriting the new rule. This parameter is optional when you bind a rule to a user-defined datatype. It is meaningless when you bind a rule to a column.

Examples

Example 1

```
sp_bindrule today, "employees.startdate"
```

Assuming that a rule named `today` has been created in the current database with `create rule`, this command binds it to the `startdate` column of the `employees` table. When a row is added to `employees`, the data for the `startdate` column is checked against the rule `today`.

Example 2

```
sp_bindrule rule_ssn, ssn
```

Assuming the existence of a rule named `rule_ssn` and a user-defined datatype named `ssn`, this command binds `rule_ssn` to `ssn`. In a `create table` statement, columns of type `ssn` inherit the rule `rule_ssn`. Existing columns of type `ssn` also inherit the rule `rule_ssn`, unless `ssn`'s rule was previously changed (in which case the changed rule is maintained in the future only).

Example 3

```
sp_bindrule rule_ssn, ssn, futureonly
```

The rule `rule_ssn` is bound to the user-defined datatype `ssn`, but no existing columns of type `ssn` are affected. `futureonly` prevents existing columns of type `ssn` from inheriting the rule.

Usage

- Create a rule using the create rule statement. Then execute *sp_bindrule* to bind it to a column or user-defined datatype in the current database.
- Rules are enforced when an insert is attempted, not when *sp_bindrule* is executed. You can bind a character rule to a column with an exact or approximate numeric datatype, even though such an insert is illegal.
- You cannot use *sp_bindrule* to bind a check constraint for a column in a create table statement.
- You cannot bind a rule to an Adaptive Server-supplied datatype or to a text or an image column.
- You cannot bind a rule to a system table.
- If you are binding to a column, the *objname* argument must be of the form “*table.column*”. Any other format is assumed to be the name of a user-defined datatype.
- Statements that use a rule cannot be in the same batch as their *sp_bindrule* statement.
- You can bind a rule to a column or user-defined datatype without unbinding an existing rule. Rules bound to columns always take precedence over rules bound to user-defined datatypes. Binding a rule to a column will replace a rule bound to the user-defined datatype of that column, but binding a rule to a datatype will not replace a rule bound to a column of that user-defined datatype. Table 6-1 indicates the precedence when binding rules to columns and user-defined datatypes where rules already exist:

Table 6-1: Precedence of new and old bound rules

New Rule Bound To:	Old Rule Bound To:	
	User-Defined Datatype	Column
user-defined datatype	Replaces old rule	No change
column	Replaces old rule	Replaces old rule

- Existing columns of the user-defined datatype inherit the new rule unless their rule was previously changed, or the value of the optional third parameter is *futureonly*. New columns of the user-defined datatype always inherit the rule.

Permissions

Only the object owner can execute *sp_bindrule*.

See also

Commands – create rule, drop rule

System procedures – *sp_unbindrule*

Procedures: *sp_cacheconfig* – *sp_cursorinfo*

sp_cacheconfig

Description	Creates, configures, reconfigures, and drops data caches, and provides information about them.
Syntax	<code>sp_cacheconfig [cachename [,"cache_size[P K M G]"] [,logonly mixed] [,strict relaxed] [, "cache_partition=[1 2 4 8 16 32 64]"]</code>
Parameters	<p><i>cachename</i></p> <ul style="list-style-type: none"> – is the name of the data cache to be created or configured. Cache names must be unique, and can be up to 30 characters long. A cache name does not have to be a valid Adaptive Server identifier, that is, it can contain spaces and other special characters. <p><i>cache_size</i></p> <ul style="list-style-type: none"> – is the size of the data cache to be created or, if the cache already exists, the new size of the data cache. The minimum size of a cache is one fourth the server's logical page size. Size units can be specified with P for pages, K for kilobytes, M for megabytes, or G for gigabytes. The default is K. For megabytes and gigabytes, you can specify floating-point values. The cache size is in multiples of the logical page size, up to four times a logical page size. <p>logonly mixed</p> <ul style="list-style-type: none"> – specifies the type of cache. <p>strict relaxed</p> <ul style="list-style-type: none"> – specifies the cache replacement policy. <p>cache_partition</p> <ul style="list-style-type: none"> – specifies the number of partitions to create in the cache. Each pool in the cache must be at least one fourth the logical page size.
Examples	<p>Example 1</p> <pre>sp_cacheconfig pub_cache, "10M"</pre>

Creates the data cache `pub_cache` with 10MB of space. All space is in the default logical page size memory pool.

Example 2

```
sp_cacheconfig pub_cache
```

Reports the current configuration of `pub_cache` and any memory pools in the cache.

Example 3

```
sp_cacheconfig pub_cache, "0"
```

Drops `pub_cache` at the next start of Adaptive Server.

Example 4

```
sp_cacheconfig pub_log_cache, "2000K", logonly
```

Creates `pub_log_cache` and sets its type to `logonly` in a single step.

Example 5

```
sp_cacheconfig pub_log_cache, "2000K"  
sp_cacheconfig pub_log_cache, logonly
```

The first command creates the cache `pub_log_cache` with the default type `mixed`. The second command changes its status to `logonly`. The resulting configuration is the same as that in example 4.

Example 6

```
sp_cacheconfig 'newcache', '50M', mixed, strict,  
"cache_partition=2"
```


Creates a cache and sets the size, type, replacement policy and number of cache partitions.

Usage

- Creating data caches divides Adaptive Server's single default data cache into smaller caches. You can then configure pools within a data cache to allow Adaptive Server to perform large I/O using `sp_poolconfig`. You can bind tables, indexes, databases, and text or image chains to a specific cache using `sp_bindcache`.
- The minimum cache size is 256 times the logical page size. For example, a 4K server would have a minimum cache size of 1024K.
- When you first create a data cache:
 - All space is allocated to the logical page size memory pool.
 - The default type is `mixed`.

- Figure 7-1 shows a data cache for a 2K server with two user-defined data caches configured and the following pools:
 - The default data cache with a 2K pool and a 16K pool
 - A user cache with a 2K pool and a 16K pool
 - A log cache with a 2K pool and a 4K pool

Figure 7-1: Data cache with default and user-defined caches



- Creating, dropping, and changing the replacement policy or number of partitions require a restart of Adaptive Server for the configuration to take effect. You cannot configure pools or bind objects to caches until the cache is active, that is, until the server has been restarted.

Other changes to data caches take effect without a restart, including changing the type, creating, dropping, and resizing memory pools with *sp_poolconfig*, changing the wash percentage of the pools, and binding and unbinding objects.

- The default data cache must always have the type default, and no other cache can have the type default.

- The Adaptive Server housekeeper task does not do any buffer washing in caches with a type of logonly or in caches with a relaxed LRU replacement policy.
- The following commands perform only 2K I/O: disk init, some dbcc commands, and drop table. The dbcc checkdb and dbcc checktable commands can perform large I/O for tables, but perform 2K I/O on indexes. Table 7-1 shows cache usage, depending on the binding of the database or object.

Table 7-1: Cache usage for Transact-SQL commands

Command	Database Bound	Table or Index Is Bound	Database or Object Not Bound
create index	Bound cache	N/A	Default data cache
disk init	N/A	N/A	Default data cache
dbcc checkdb	Bound cache	N/A	Default data cache
dbcc checktable, indexalloc, tablealloc	Bound cache	Bound cache	Default data cache
drop table	Bound cache	Bound cache	Default data cache

- Recovery uses only the logical page size pool of the default data cache. All pages for all transactions that must be rolled back or rolled forward are read into and changed in this pool. Be sure that your default logical page size pool is large enough for these transactions.
- When you use sp_cacheconfig with no parameters, it reports information about all of the caches on the server. If you specify only a cache name, it reports information about only the specified cache. If you use a fragment of a cache name, it reports information for all names matching “%fragment%”.

All reports include a block of information that reports information about caches, and a separate block of data for each cache that provides information about the pools within the cache.

The output below, from a server using 2K, shows the configuration for:

- The default data cache with two pools: a 2K pool and a 16K pool. The default data cache has 2 partitions.
- pubs_cache with two pools: 2K and 16K
- pubs_log, with the type set to logonly and cache replacement policy set to relaxed, with a 2K pool and a 4K pool

```

Cache Name           Status      Type      Config Value  Run Value
-----
default data cache   Active     Default   0.00 Mb      26.09 Mb
pubs_cache           Active     Mixed     10.00 Mb     10.00 Mb
pubs_log             Active     Log Only  2.40 Mb      2.40 Mb
-----
Total                12.40 Mb   38.49 Mb
=====
Cache: default data cache,  Status: Active,  Type: Default
      Config Size: 0.00 Mb,  Run Size: 26.09 Mb
      Config Replacement: strict LRU,  Run Replacement: strict LRU
      Config Partition:      2,  Run Partition:      2
IO Size  Wash Size Config Size  Run Size      APF Percent
-----
      2 Kb   3704 Kb    0.00 Mb    18.09 Mb    10
      16 Kb  1632 Kb    8.00 Mb    8.00 Mb    10
=====
Cache: pubs_cache,  Status: Active,  Type: Mixed
      Config Size: 10.00 Mb,  Run Size: 10.00 Mb
      Config Replacement: strict LRU,  Run Replacement: strict LRU
      Config Partition:      1,  Run Partition:      1
IO Size  Wash Size Config Size  Run Size      APF Percent
-----
      2 Kb   1228 Kb    0.00 Mb    6.00 Mb    10
      16 Kb   816 Kb    4.00 Mb    4.00 Mb    10
=====
Cache: pubs_log,  Status: Active,  Type: Log Only
      Config Size: 2.40 Mb,  Run Size: 2.40 Mb
      Config Replacement: relaxed LRU,  Run Replacement: relaxed LRU
      Config Partition:      1,  Run Partition:      1
IO Size  Wash Size Config Size  Run Size      APF Percent
-----
      2 Kb   206 Kb    0.00 Mb    1.01 Mb    10
      16 Kb   272 Kb    1.40 Mb    1.39 Mb    10

```

Table 7-2 lists the meaning of the columns in the output:

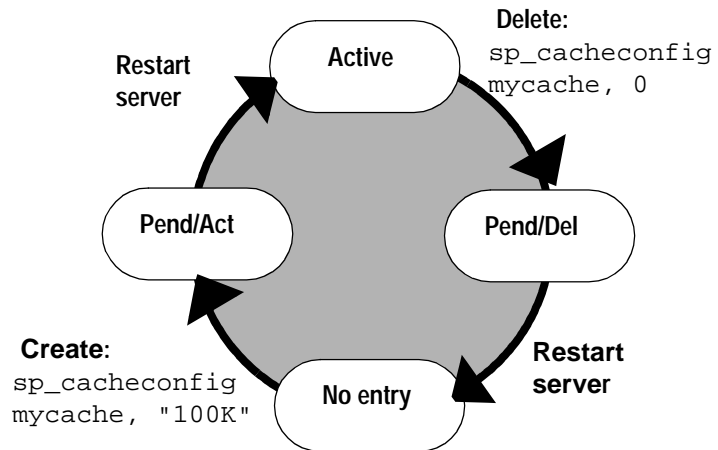
Table 7-2: *sp_cacheconfig* output

Column	Meaning
Cache Name	The name of the cache.

Column	Meaning
Status	One of the following: <ul style="list-style-type: none"> • “Active” • “Pend/Act” • “Pend/Del” These are explained following this table.
Type	“Mixed” or “Log Only” for user-defined caches, “Default” for the default data cache.
I/O Size	The size of I/O for a memory pool. This column is blank on the line that shows that cache configuration.
Wash Size	The size of the wash area for the pool. As pages enter the wash area of the cache, they are written to disk. This column is blank on the line that shows the cache configuration.
Config Value or Config Size	The size that the cache or pool will have after the next time Adaptive Server is restarted. These are the values that take effect the next time Adaptive Server is restarted. If the value is 0, the size has not been explicitly configured, and a default value will be used.
Run Value or Run Size	The size of the cache or pool now in use on Adaptive Server.
Config/ Run Replacement	The cache policy (strict or relaxed) that will be used for the cache after the next restart, and the current replacement policy. These will be different only if the policy has been changed since the last reboot.
Config/Run Partition	The number of cache partitions that will be used for the cache after the next restart, and the current number of partitions. These will be different if sp_cacheconfig has been used to change the number of partitions since the last reboot.
APF Percent	The percentage of buffers in the pool that can hold buffers that have been fetched by asynchronous prefetch, but have not been used.
Total	The total size of data cache, if the report covers all caches, or the current size of the particular cache, if you specify a cache name.

The status “Pend” is short for pending. It always occurs in combination with either “Act” for Active or “Del” for Delete. It indicates that a configuration action has taken place, but that the server must be restarted in order for the changes to take effect.

When you first create a new cache, but have not yet restarted Adaptive Server, the status is “Pend/Act”, meaning that the cache has just been configured and will be active after a restart. If you set the size of a cache to 0 to delete it, the status changes from “Active” to “Pend/Del”, meaning that the cache still exists, and still functions, but that it will be deleted at the next restart.

Figure 7-2: Effects of restarts and `sp_cacheconfig` on cache status

- You can also configure caches and pools by editing the configuration file. For more information, see the System Administration Guide.

Data cache memory

- When Adaptive Server is first installed, all data cache memory is assigned to the logical page size pool of the cache named default data cache. The default data cache is used by all objects that are not explicitly bound to a data cache with `sp_bindcache` or whose databases are not bound to a cache.
- When you create data caches, the memory allocation comes from the default data cache. Memory for caches is allocated out of the memory allocated to Adaptive Server with the `total_logical_memory` configuration parameter. To increase the amount of space available for caches, increase total logical memory, or decrease other configuration settings that use memory. If you need to decrease the size of total logical memory, the space must be available in the default data cache.

You cannot reduce the size of the default data cache to less than one fourth a logical page. In most cases, the default cache should be much larger than the minimum. This cache is used for all objects, including system tables, that are not bound to another cache, and is the only cache used during recovery. For more information, see the *System Administration Guide*.

- A data cache requires a small percentage of overhead for structures that manage the cache. All cache overhead is taken from the default data cache. To see the amount of overhead required for a specific size of cache, use `sp_helpcache`, giving the size:

```
sp_helpcache "200M"  
10.38Mb of overhead memory will be needed to  
manage a cache of size 200M
```

Changing existing caches

- To change the size of an existing cache, specify the cache's name and the new size.
 - If you increase the size of an existing cache, all of the added space is placed in the logical page size pool.
 - To reduce the size of an existing cache, all of the space must be available in the logical page size pool. You may need to use `sp_poolconfig` to move space from other pools to this pool.
- If you have a database or any nonlog objects bound to a cache, you cannot change its type to logonly.

Using cache partitions

- Cache partitions can be used to reduce cache spinlock contention without needing to create separate caches and bind database objects to them. For more information on monitoring cache spinlock contention, see the *Performance and Tuning Guide*.
- You can set the default number of cache partitions for all caches with the configuration parameter `global cache partition number`. See the *System Administration Guide*.

Dropping caches

- To drop or delete a data cache, change its size to 0, as shown in example 3. When you set a cache's size to 0, the cache is marked for deletion, but it is not dropped until the next restart of the server. The cache remains active, and all objects that are bound to that cache continue to use it.

You cannot drop the default data cache.

- If you drop a cache that has objects bound to it, all of the object bindings for the cache are marked invalid the next time you restart Adaptive Server. A message is printed to the error log on restart, giving the database ID, object ID and index ID:

```
00:95/11/05 18:20:39.42 server Cache binding for
database '6', object '8', index '0' is being
marked invalid in Sysattributes.
```

If you subsequently create a cache of the same name, bindings are marked valid when the cache is activated.

- You cannot run *sp_cacheconfig* within a transaction.

Permissions

Only a System Administrator can execute *sp_cacheconfig* to change cache configurations. Any user can execute *sp_cacheconfig* to view cache configurations.

See also

System procedures – *sp_bindcache*, *sp_helpcache*, *sp_unbindcache*, *sp_unbindcache_all*

sp_cachestrategy

Description

Enables or disables prefetching (large I/O) and MRU cache replacement strategy for a table, index, text object, or image object.

Syntax

```
sp_cachestrategy dbname, [ownername.]tablename
[, indexname | "text only" | "table only"
[, { prefetch | mru }, { "on" | "off"}]]
```

Parameters

dbname

– is the name of the database where the object is stored.

ownername

– is the name of the table's owner. If the table is owned by "dbo", the owner name is optional.

tablename

– is the name of the table.

indexname

– is the name of the index on the table.

text only

– changes the cache strategy for a text or image object.

table only

– changes the cache strategy for a table.

prefetch

| mru – is prefetch or mru, and specifies which setting to change.

on | off

– specifies the setting, "on" or "off", enclosed in quotes.

Examples

Example 1

```
sp_cachestrategy pubs2, titles
object name          index name      large IO MRU
-----
dbo.titles           titleidind    ON          ON
```

Displays information about cache strategies for the titles table.

Example 2

```
sp_cachestrategy pubs2, titles, titleind
```

Displays information about cache strategies for the titleind index.

Example 3

```
sp_cachestrategy pubs2, titles, titleind, prefetch,
"off"
```

Disables prefetch on the titleind index of the titles table.

Example 4

```
sp_cachestrategy pubs2, authors, "table only", mru,
"on"
```

Reenables MRU replacement strategy on the authors table.

Example 5

```
sp_cachestrategy pubs2, blurbs, "text only",
prefetch, "on"
```

Reenables prefetching on the text pages of the blurbs table.

Usage

- If memory pools for large I/O are configured for the cache used by a table or an index, the optimizer can choose to prefetch data or index pages by performing large I/Os of up to eight data pages at a time. This prefetch strategy can be used on the data pages of a table or on the leaf-level pages of a nonclustered index. By default, prefetching is enabled for all tables, indexes, and text or image objects. Setting the prefetch option to "off" disables prefetch for the specified object.

- The optimizer can choose to use **MRU replacement strategy** to fetch and discard buffers in cache for table scans and index scans for I/O of any size. By default, this strategy is enabled for all objects. Setting `mru` to "off" disables this strategy. If you turn `mru` off for an object, all pages are read into the MRU/LRU chain in cache, and they remain in the cache until they are flushed by additional I/O. For more information on cache strategies, see the Performance and Tuning Guide.
- You can change the cache strategy only for objects in the current database.
- When you use `sp_cachestrategy` without specifying the strategy and setting, it reports the current settings for the object, as shown in example 1.
- To see the size, status and I/O size of all data caches on the server, use `sp_cacheconfig`.
- Setting `prefetch` "on" has no effect on tables or indexes that are read into a cache that allows only 2K I/O. The `mru` strategy can be used in all caches, regardless of available I/O size.

Overrides

- If prefetching is turned on for a table or an index, you can override the prefetching for a session with `set prefetch "off"`. If prefetching is turned off for an object, you cannot override that setting.
- The `prefetch`, `lru`, and `mru` options to the `select`, `delete` and `update` commands suggest the I/O size and cache strategy for individual statements. If prefetching or MRU strategy is enabled for a table or an index, you can override it for a query by specifying I/O the size of the logical page size for `prefetch`, and by specifying `lru` strategy. For example, the following command forces LRU strategy, logical page size I/O, and a table scan of the `titles` table:

```
select avg(advance)
from titles (index titles prefetch 2 lru)
```

If you request a `prefetch` size, and the object's cache is not configured for I/O of the requested size, the optimizer chooses the best available I/O size.

- If prefetching is enabled for an object with `sp_cachestrategy`, using a prefetch specification of the logical page size in a select, update or delete command overrides an earlier set prefetch "on" statement. Specifying a larger I/O size in a select, update or delete command does not override a set prefetch "off" command.

Permissions Only a System Administrator or the object owner can execute `sp_cachestrategy`.

See also *Commands* – delete, select, set, update

Stored procedures – `sp_cacheconfig`, `sp_poolconfig`

sp_changedbowner

Description Changes the owner of a user database.

Syntax `sp_changedbowner loginame [, true]`

Parameters *loginame*
– is the login name of the new owner of the current database.

`true`
– transfers aliases and their permissions to the new database owner. Values are “true” and “TRUE”.

Examples `sp_changedbowner albert`

Makes the user “albert” the owner of the current database.

Usage

- The new owner must not already be known as either a user or alias (that is, the new owner must not already be listed in `sysusers` or `sysalternates`). Executing `sp_changedbowner` with the single parameter *loginame* changes the database ownership to *loginame* and drops aliases of users who could act as the old “dbo.”
- After executing `sp_changedbowner`, the new owner is known as the Database Owner inside the database.
- `sp_changedbowner` cannot transfer ownership of the system databases.
- The new owner must already have a login name in Adaptive Server, but must *not* have a database user name or alias name in the database. To assign database ownership to such a user, drop the user name or alias entry before executing `sp_changedbowner`.

- To grant permissions to the new owner, a System Administrator must grant them to the Database Owner, since the user is no longer known inside the database under any other name.

Permissions

Only a System Administrator can execute *sp_changedbowner*.

See also

Commands – create database

System procedures – *sp_addlogin*, *sp_dropalias*, *sp_dropuser*, *sp_helpdb*

sp_changegroup

Description

Changes a user's group.

Syntax

sp_changegroup grpname, username

Parameters

grpname

– is the name of the group. The group must already exist in the current database. If you use “public” as the *grpname*, enclose it in quotes, because it is a keyword.

username

– is the name of the user to be added to the group. The user must already exist in the current database.

Examples

Example 1

```
sp_changegroup fort_mudge, albert
```

The user “albert” is now a member of the “fort_mudge” group. It doesn't matter what group “albert” belonged to before.

Example 2

```
sp_changegroup "public", albert
```

Removes “albert” from the group he belonged to without making him a member of a new group (all users are always members of “public”).

Usage

- Executing *sp_changegroup* adds the specified user to the specified group. The user is dropped from the group he or she previously belonged to and is added to the one specified by *grpname*.
- New database users can be added to groups at the same time they are given access to the database with *sp_adduser*.

- Groups are used as a collective name for granting and revoking privileges. Every user is always a member of the default group, “public”, and can belong to only one other group.
- To remove someone from a group without making that user a member of a new group, use sp_changegroup to change the user’s group to “public”, as shown above in example 2.
- When a user changes from one group to another, the user loses all permissions that he or she had as a result of belonging to the old group and gains the permissions granted to the new group.

Permissions Only the Database Owner, a System Administrator, or a System Security Officer can execute sp_changegroup.

See also *Commands* – grant, revoke
System procedures – sp_addgroup, sp_adduser, sp_dropgroup, sp_helpgroup

sp_checknames

Description Checks the current database for names that contain characters not in the 7-bit ASCII set.

Syntax sp_checknames

Parameters None.

Examples sp_checknames

```
Looking for non 7-bit ASCII characters in the system tables
of database:
"master"
```

```
=====
Table.Column name: "syslogins.password"
```

The following logins have passwords that contain non 7-bit ASCII characters. If you wish to change them use "sp_password"; Remember, only the sa and the login itself may examine or change the syslogins.password column:

```

suid  name
-----
1  sa
2  probe
3  bogususer

```

Usage	<ul style="list-style-type: none"> • <i>sp_checknames</i> examines the names of all objects, columns, indexes, user names, group names, and other elements in the current database for characters outside of the 7-bit ASCII set. It reports illegal names and gives instructions to make them compatible with the 7-bit ASCII set. • Run <i>sp_checknames</i> in every database on your server after upgrading from a SQL Server of release 4.0.x or 4.2.x, and after using a default character set that was not 7-bit ASCII. • Follow the instructions in the <i>sp_checknames</i> report to correct all non-ASCII names.
Permissions	Any user can execute <i>sp_checknames</i> .
See also	<p><i>Commands</i> – update</p> <p><i>System procedures</i> – <i>sp_password</i>, <i>sp_rename</i>, <i>sp_renamedb</i></p>

sp_checkreswords

Description	Detects and displays identifiers that are Transact-SQL reserved words. Checks server names, device names, database names, segment names, user-defined datatypes, object names, column names, user names, login names, and remote login names.
Syntax	<i>sp_checkreswords</i> [<i>user_name_param</i>]
Parameters	<p><i>user_name_param</i></p> <p>– is the name of a user in the current database. If you supply <i>user_name_param</i>, <i>sp_checkreswords</i> checks only for objects that are owned by the specified user.</p>
Examples	<p>Example 1</p> <pre> 1> /* executed in the master database */ 2> sp_checkreswords </pre>

Reserved Words Used as Database Object Names for Database master

Upgrade renames sysobjects.schema to sysobjects.schemact.

Owner

dbo

Table Reserved Word Column Names

authorization cascade

Object Type Reserved Word Object Names

rule constraint
stored procedure check
user table arith_overflow
user table authorization

Owner

lemur

Table Reserved Word Column Names

key close

Table Reserved Word Index Names

key isolation

Object Type Reserved Word Object Names

default isolation
rule level
stored procedure mirror
user table key

Reserved Word Datatype Names

identity

```
-----
-----
Database-wide Objects
-----
```

```
Reserved Word User Names
-----
```

```
at
identity
```

```
Reserved Word Login Names
-----
```

```
at
identity
```

```
Reserved Word as Database Names
-----
```

```
work
```

```
Reserved Word as Language Names
-----
```

```
national
```

```
Reserved Word as Server Names
-----
```

```
mirror
primary
```

```
Reserved Word ServerNetNames
-----
```

```
mirror
primary
```

This example shows the results if *sp_checkreswords* is executed in the master database.

Example 2

```
1> /* executed in the user database, user_db */
2> sp_checkreswords
```

Reserved Words Used as Database Object Names for Database user_db

Upgrade renames sysobjects schema to sysobjects.schemact.

Owner

tamarin

Table	Reserved Word Column Names
cursor	current
endtran	current
key	identity
key	varying
schema	primary
schema	references
schema	role
schema	some
schema	user
schema	work

Table	Reserved Word Index Names
key	double

Object Type	Reserved Word Object Names
default	escape
rule	fetch
stored procedure	foreign
user table	cursor
user table	key
user table	schema
view	endtran

Database-wide Objects

Found no reserved words used as names for database-wide objects.

This example shows the results if sp_checkreswords is executed in the user database user_db.

Usage

- *sp_checkreswords* reports the names of existing objects that are reserved words. Transact-SQL does not allow words that are part of any command syntax to be used as identifiers, unless you are using delimited identifiers. Reserved words are pieces of SQL syntax, and they have special meaning when you use them as part of a command. For example, in pre-release 10.0 SQL Server, you could have a table called *work*, and select data from it with this query:

```
select * from work
```

work was a new reserved word in SQL Server release 10.0, part of the command *commit work*. Issuing the same *select* statement in release 10.0 or later causes a syntax error. *sp_checkreswords* finds identifiers that would cause these problems.

- *sp_checkreswords* also finds reserved words, used as identifiers, that were created using the *set quoted_identifier* option.
- Use *sp_checkreswords* before or immediately after upgrading to a new release of Adaptive Server. For information on installing and running this procedure before performing the upgrade, see the installation documentation for your platform.

Run *sp_checkreswords* in the master database and in each user database. Also run it in *model* and *sysystemprocs*, if you have added users or objects to those databases.

- The return status indicates the number of items found.
- If you supply a user name, *sp_checkreswords* checks for all of the objects that can be owned by a user tables, indexes, views, procedures, triggers, rules, defaults, and user-defined datatypes. It reports all identifiers that are reserved words.
- If your current database is not the master database, and you do not provide a user name, *sp_checkreswords* checks for all of the objects above, with a separate section in the report for each user name. It also checks *sysusers* and *syssegments* for user names and segment names that are reserved words. You only need to check *model* and *sysystemprocs* if you have added objects, users, or user-defined datatypes.

- If your current database is master, and you do not provide a user name, sp_checkreswords performs all of the checks above and also checks sysdatabases, syslogins, syscharsets, syssservers, sysremotelogins, sysdevices, and syslanguages for reserved words used as the names of databases, local or remote logins, local and remote servers, character sets, and languages.

Handling reported instances of reserved words

- If sp_checkreswords reports that reserved words are used as identifiers, you have two options:
 - Use sp_rename, sp_renamedb, or update the system tables to change the name of the identifier.
 - Use set quoted_identifier on if the reserved word is a table name, view name, or column name. If most of your applications use stored procedures, you can drop and re-create these procedures with set quoted_identifier on, and quote all identifiers. All users will be able to run the procedures, without having to use set quoted_identifier on for their session. You can use set quoted_identifier on, create views that give alternative names to tables or columns, and change your applications to reference the view instead.

The following example provides alternatives for the new reserved words “key”, “level”, and “work”:

```
create view keyview
as
select lvl = "level", wrk = "work"
from "key"
```

The syntax for the set command is:

```
set quoted_identifier on
```

- If you do not either change the identifiers or use delimited identifiers, any query that uses the reserved words as identifiers reports an error, usually a syntax error. For example:

```
select level, work from key
Msg 156, Level 15, State 1:
Server 'rosie', Line 1:
```

Incorrect syntax near the keyword 'level'.

Note The quoted identifier option is a SQL92 option and may not be supported by many client products that support other Adaptive Server features. For example, you cannot use `bcp` on tables whose names are reserved words.

Before choosing the quoted identifier option, perform a test on various objects using all the tools you will use to access Adaptive Server. Use `set quoted_identifier on`, create a table with a reserved word for a name and reserved words for column names. If the client product generates SQL code, it must enclose identifiers in double quotes (if they are reserved words) and character constants in single quotes.

- Procedures, triggers, and views that depend on objects whose names have been changed may work after the name change, but will stop working when the query plan is recompiled. Recompilation takes place for many reasons, without notification to the user. To avoid unsuspected loss of functionality, change the names of objects in procedures, triggers, and views immediately after you change the object name.
- Whether you change the object names or use delimited identifiers, you must change all stored procedures, views, triggers, and applications that include the reserved word. If you change object names, you must change identifiers; if you use delimited identifiers, you must add the `set quoted_identifier` option and quotation marks.
- If you do not have the text of your procedures, triggers, views, rules, and defaults saved in operating system files, you can use `defncopy` to copy the definitions from the server to files. See `defncopy` in .

Changing identifiers

- If you change the names of the items reported by `sp_checkreswords`, you must change the names in all procedures, triggers, views, and applications that reference the object using the reserved word.
- Dump your database before changing identifier names. After you change the identifier names, run `dbcc` to determine that there are no problems, and dump the database again.
- If you are changing identifiers on an active production database:
 - Perform the changes when the system is least busy, so that you will disrupt as few users as possible.

- Prepare carefully by finding all Open Client DB-Library™ programs, windowing applications, stored procedures, triggers, and scripts that use a particular identifier. This way, you can make the edits needed in the source code, then change the identifiers and replace the procedures and code as quickly as possible.
- The procedure `sp_depends` can help find procedures, views, and triggers that use table and view names.

Using `sp_rename` to change identifiers

- The system procedure `sp_rename` renames tables, indexes, views, procedures, triggers, rule, defaults, user-defined datatypes, and columns. Use `sp_renamedb` to rename databases.
- Table 7-3 shows the types of identifiers that you can change with `sp_rename` and lists other changes that may have to be made on the server and in your application programs.

Table 7-3: `sp_rename` and changing identifiers

Identifier	Remember To
Table name	<ul style="list-style-type: none"> • Drop all procedures, triggers and views that reference the table, and re-create them with the new name. Use <code>sp_depends</code> to find the objects that depend on the table. • Change all applications or SQL source scripts that reference the table to use the new table name. • Change <code>dbcc</code> scripts that perform table-level checks using table names.
Index name	<ul style="list-style-type: none"> • Drop any stored procedures that create or drop the index, and re-create them with the new name. • Change all applications or SQL source scripts that create or drop the index. • Change <code>dbcc</code> scripts that perform index-level checks using index names.
View name	<ul style="list-style-type: none"> • Drop all procedures, triggers, and views that reference the view, and re-create them with the new name. Use <code>sp_depends</code> to find the objects that depend on the view. • Change all applications or SQL source scripts that reference the view to use the new view name.
Procedure name	<ul style="list-style-type: none"> • Drop and re-create with the new procedure name all procedures and triggers that reference the procedure. • Change all applications or SQL source scripts that execute the procedure to use the new name. • If another server remotely calls the procedure, change applications on the remote server to use the new procedure name.
Trigger name	<ul style="list-style-type: none"> • Change any SQL source scripts that create the trigger.
Rule name	<ul style="list-style-type: none"> • Change any SQL source scripts that create the rule.
Default name	<ul style="list-style-type: none"> • Change any SQL source scripts that create the default.

Identifier	Remember To
User-defined datatype name	<ul style="list-style-type: none"> Drop all procedures that create tables with user-defined datatypes, and re-create them with the new name. Change any applications that create tables with user-defined datatypes.
Column name	<ul style="list-style-type: none"> Drop all procedures, triggers and views that reference the column, and re-create them with the new column name. <i>sp_depends</i> cannot find column name references. The following query displays the names of procedures, triggers, and views that reference a column named “key”: <pre>select distinct sysobjects.name from sysobjects, syscomments where sysobjects.id = syscomments.id and syscomments.text like "%key%"</pre> Change all applications and SQL source scripts that reference the column by name.

The following command changes the name of the view isolation to isolated:

```
sp_rename "isolation", isolated
```

The following command changes the name of a column in the renamed view isolated:

```
sp_rename "isolated.key", keyname
```

- Use *sp_depends* to get a list of all views, procedures, and triggers that reference a view, procedure, or table that will be renamed. To use *sp_depends* after renaming an object, give the new name. For example:

```
sp_depends new_name
```

Renaming databases with *sp_renamedb*

- To change the name of a database, use *sp_renamedb*. The database must be in single-user mode. Drop and re-create any procedures, triggers, and views that explicitly reference the database name. For more information, see *sp_renamedb*.

Changing other identifiers

- To change user names, login names, device names, remote server names, remote server user names, segment names, and character set and language names, first determine if you can drop the object or user, then add or create it again. If you cannot do that, use the command:

```
sp_configure "allow updates to system tables", 1
```

to allow direct updates to system tables. Only a System Security Officer can set the allow updates to system tables configuration parameter.

Errors during direct updates to system tables can create severe problems in Adaptive Server. To determine whether you can drop the objects or user, then re-create them, see Table 7-4.

Table 7-6: Considerations when changing identifiers on shows possible dependencies on this set of identifiers. See this table for possible dependencies, whether you choose to upgrade by dropping and recreating objects, by using delimited identifiers, or by performing direct updates to system tables.

Table 7-4: Alternatives to direct system tables updates when changing identifiers

Identifier Type	Suggested Actions to Avoid Updates to System Tables
User names and login names	To change the name of a user with no objects, first use <code>sp_helprotect username</code> in each database to record the user’s permissions. Then, drop the user from all of the databases (<code>sp_dropuser</code>), and drop the login (<code>sp_droplogin</code>). Finally, add the new login name (<code>sp_addlogin</code>), add the new user name to the databases (<code>sp_adduser</code>), and restore the user’s permissions with <code>grant</code> .
Device names	If this device is completely allocated, you will not need to use its name in a <code>create database</code> command, so you can leave the name unchanged.
Remote server names	Unless there are large numbers of remote login names from the remote server, drop the remote server (<code>sp_dropserver</code>) and add it with a new name (<code>sp_addserver</code>).
Remote server logins	Drop the remote login with <code>sp_dropremotelogin</code> , add it with a new name using <code>sp_addremotelogin</code> , and restore the user’s permission to execute procedures with <code>grant</code> .
Segment names	These are rarely used, once objects have been created on the segments.
Character set and language names	Languages and character sets have reserved words as identifiers only if a System Administrator has created alternative languages with <code>sp_addlanguage</code> . Drop the language with <code>sp_droplanguage</code> , and add it with a new name.

Warning! Direct updates to system tables can be very dangerous. You can make mistakes that make it impossible for Adaptive Server to run or make it impossible to access objects in your databases. Undertake this effort when you are calm and collected, and when little or no production activity is taking place on the server. If possible, use the alternative methods described Table 7-4.

- The following example shows a “safe” procedure for updating a user name, with all data modification preceded by a `begin transaction` command:

The System Security Officer executes the following command:

```
sp_configure "allow updates to system tables", 1
```

Then you can execute the following:

```
begin transaction
update sysusers
set name = "workerbee"
where name = "work"
```

At this point, run the query, and check to be sure that the command affected only the row that you intended to change. The only identifier change that affects more than one row is changing the language name in `syslogins`.

- If the query affected only the correct row, use `commit transaction`.
- If the query affected more than one row, or the incorrect row, use `rollback transaction`, determine the source of the problem, and execute the command correctly.

When you are finished, the System Security Officer turns off the allow updates to system tables configuration parameter with this command:

```
sp_configure "allow updates to system tables", 0
```

Warning! Only update system tables in a single database in each user defined transaction. Do not issue a `begin transaction` command and then update tables in several databases. Such actions can make recovery extremely difficult.

Table 7-5 shows the system tables and columns that you should update to change reserved words. The tables preceded by “`master.dbo.`” occur only in the master database. All other tables occur in master and in user databases. Be certain you are using the correct database before you attempt the update. You can check for the current database name with this command:

```
select db_name()
```

Table 7-5: System table columns to update when changing identifiers

Type of Identifier	Table to Update	Column Name
User name	sysusers	name
Login names	master.dbo.syslogins	name
Segment names	syssegments	name
Device name	sysdevices	name
Remote server name	syssservers	srvname
Remote server network name	syssservers	srvnetname
Character set names	master.dbo.syscharsets	name
Language name	master.dbo.syslanguages master.dbo.syslogins	name language

Table 7-6 shows other changes that may have to be made on the server and in your application programs:

Table 7-6: Considerations when changing identifiers

Identifier	Remember To
Login name	Change the user name in each database where this person is a user.
User name	Drop, edit, and re-create all procedures, triggers, and views that use qualified (<i>owner_name.object_name</i>) references to objects owned by this user. Change all applications and SQL source scripts that use qualified object names to use the new user name. You do not have to drop the objects themselves; sysusers is linked to sysobjects by the column that stores the user's ID, not the user's name.
Device name	Change any SQL source scripts or applications that reference the device name to use the new name.
Remote server name	Change the name on the remote server. If the name that sp_checkreswords reports is the name of the local server, you must restart the server before you can issue or receive remote procedure calls.
Remote server network name	Change the server's name in the interfaces files.
Remote server login name	Change the name on the remote server.
Segment name	Drop and re-create all procedures that create tables or indexes on the segment name. Change all applications that create objects on segments to use the new segment name.
Character set name	None.
Language name	Change both master.dbo.syslanguages and master.dbo.syslogins. The update to syslogins may involve many rows. Also, change the names of your localization files.

Using delimited identifiers

- You can use delimited identifiers for table names, column names, and view names. You cannot use delimited identifiers for other object names.
- If you choose to use delimited identifiers, use `set quoted_identifier on`, and drop and re-create all the procedures, triggers, and views that use the identifier. Edit the text for those objects, enclosing the reserved words in double quotes and enclosing all character strings in single quotes.

The following example shows the changes to make to queries in order to use delimited identifiers. This example updates a table named `work`, with columns named `key` and `level`. Here is the pre-release 10.0 query, which encloses character literals in double quotes, and the edited version of the query for use with delimited identifiers:

```
/* pre-release 10.0 version of query */
update work set level = "novice"
    where key = "19-732"
/* 10.0 or later version of query, using
** the quoted identifier option
*/
update "work" set "level" = 'novice'
    where "key" = '19-732'
```

- All applications that use the reserved word as an identifier must be changed as follows:
 - The application must set the quoted identifier option on.
 - All uses of the reserved word must be enclosed in double quotes.
 - All character literals used by the application while the quoted identifier option is turned on must be enclosed in single quotes. Otherwise, Adaptive Server attempts to interpret them as object names.

For example, the following query results in an error message:

```
set quoted_identifier on
select * from titles where title_id like "BU%"
```

Here is the correct query:

```
select * from titles where title_id like 'BU%'
```

- Stored procedures that you create while the delimited identifiers are in effect can be run without turning on the option. (The allow updates to system tables option also works this way.) This means that you can turn on quoted identifier mode, drop a stored procedure, edit it to insert quotation marks around reserved words used as identifiers, and re-create the procedure. All users can execute the procedure without using set quoted_identifier.

Permissions Only a System Administrator can execute sp_checksourc.

See also *Commands* – set
System procedures – sp_configure, sp_depends, sp_rename, sp_renamedb
Utilities – defncopy

sp_checksourc

Description Checks for the existence of the **source text** of the **compiled object**.

Syntax sp_checksourc [objname [, tablename [, username]]]

Parameters

objname
– is the compiled object to be checked for the existence of its source text.

tablename
– is the name of the table or view to be checked for the existence of all check constraints, defaults, and triggers defined on it.

username
– is the name of the user who owns the compiled objects to be checked for the existence of the source text.

Examples

Example 1

```
sp_checksourc
```

Checks for the existence of the source text of all compiled objects in the current database.

Example 2

```
sp_checksourc titleview
```

Checks for the existence of the source text of the view named titleview.

Example 3

```
sp_checksource title_vu, @username = Mary
```

Checks for the existence of the source text of the view named `titls_vu` that is owned by `Mary`.

Example 4

```
sp_checksource list_phone_proc
```

Checks for the existence of the source text of the custom stored procedure `list_phone_proc`.

Example 5

```
sp_checksource @tabname = "my_tab"
```

Checks for the existence of the source text of all the check constraints, triggers, and declarative defaults defined on the table named `my_tab`.

Example 6

```
sp_checksource @objname = "my_vu", @tabname =  
"my_tab"
```

Checks for the existence of the source text of the view `my_vu` and all check constraints, triggers, and defaults defined on the table `my_tab`.

Example 7

```
sp_checksource @username = "Tom"
```

Checks for the existence of the source text of all compiled objects owned by `Tom`.

Usage

- `sp_checksource` checks for the existence of the source text of the specified compiled object. If the source text exists for the specified object, `sp_checksource` returns 0. If the source text does not exist for the specified object, `sp_checksource` returns 1.
- If you do not provide any parameters, `sp_checksource` checks the existence of the source text for all compiled objects in the current database.
- To use `sp_checksource` with no parameters, you must be the Database Owner or System Administrator.

Permissions	Only a Database Owner or System Administrator can execute <code>sp_checksourc</code> to check for the existence of the source text of compiled objects that are owned by another user. Any user can execute <code>sp_checksourc</code> to check for the existence of the source text for his or her own compiled objects.
See also	<i>System procedures</i> – <code>sp_hidetext</code>

sp_chgattribute

Description	Changes the <code>max_rows_per_page</code> , <code>fillfactor</code> , <code>reservepagegap</code> , or <code>exp_row_size</code> value for future space allocations of a table or an index; sets the <code>concurrency_opt_threshold</code> for a table.
Syntax	<code>sp_chgattribute objname, {"max_rows_per_page" "fillfactor" "reservepagegap" "exp_row_size" "concurrency_opt_threshold"}, optvalue</code> <code>sp_chgattribute "table_name", "identity_gap", set_number</code>
Parameters	<p><i>objname</i></p> <ul style="list-style-type: none">– is the name of the table or index for which you want to change attributes. <p><code>max_rows_per_page</code></p> <ul style="list-style-type: none">– specifies the row size. Use this option for tables with variable-length columns. <p><code>fillfactor</code></p> <ul style="list-style-type: none">– specifies how full Adaptive Server will make each page when it is re-creating an index or copying table pages as a result of a <code>reorg rebuild</code> command or an <code>alter table</code> command to change the locking scheme. The <code>fillfactor</code> percentage is relevant only at the time the index is rebuilt. Valid values are 0–100. <p><code>reservepagegap</code></p> <ul style="list-style-type: none">– specifies the ratio of filled pages to empty pages that are to be left during extent I/O allocation operations. For each specified <i>num_pages</i>, an empty page is left for future expansion of the table. Valid values are 0–255. The default value is 0.

exp_row_size

– reserves a specified amount of space for the rows in data-only locked tables. Use this option to reduce the number of rows being forwarded, which can be expensive during updates. Valid values are 0, 1, and any value between the minimum and maximum row length for the table. 0 means a server-wide setting is applied, and 1 means to fully pack the rows on the data pages.

concurrency_opt_threshold

– specifies the table size, in pages, at which access to a data-only-locked table should begin optimizing for reducing I/O, rather than for concurrency. If the table is smaller than the number of pages specified by *concurrency_opt_threshold*, the query is optimized for concurrency by always using available indexes; if the table is larger than the number of pages specified by *concurrency_opt_threshold*, the query is optimized for I/O instead. Valid values are -1 to 32767. Setting the value to 0 disables concurrency optimization. Use -1 to enforce concurrency optimization for tables larger than 32767 pages. The default is 15 pages.

optvalue

– is the new value. Valid values and default values depend on which parameter is specified.

table_name

– is the name of the table for which you want to change the identity gap.

identity_gap

– indicates that you want to change the identity gap.

set_number

– is the new size of the identity gap.

Examples**Example 1**

```
sp_chgattribute authors, "max_rows_per_page", 1
```

Sets the *max_rows_per_page* to 1 for the *authors* table for all future space allocations.

Example 2

```
sp_chgattribute "titles.titleidind",  
"max_rows_per_page", 4
```

Sets the *max_rows_per_page* to 4 for the *titleidind* index for all future space allocations.

Example 3

```
sp_chgattribute "titles.title_ix", "fillfactor", 90
```

Specifies a fillfactor of 90 percent for pages in title_ix.

Example 4

```
sp_chgattribute authors, "exp_row_size", 120
```

Sets the exp_row_size to 120 for the authors table for all future space allocations.

Example 5

```
sp_chgattribute "titles.titleidind",  
"reservepagegap", 16
```

Sets the reservepagegap to 16 for the titleidind index for all future space allocations.

Example 6

```
sp_chgattribute "titles",  
concurrency_opt_threshold, 0
```

Turns off concurrency optimization for the titles table.

Example 7

```
sp_chgattribute "mytable", "identity_gap", 20
```

Sets the identity gap for mytable to 20.

Example 8

```
sp_chgattribute "mytable", "identity_gap", 0
```

Changes mytable to use the identity burning set factor setting instead of the identity_gap setting.

Usage

- sp_chgattribute changes the max_rows_per_page, fillfactor, reservepagegap, or exp_row_size value for future space allocations or data modifications of the table or index. It does not affect the space allocations of existing data pages. You can change these values for an object only in the current database.
- Use sp_help to see the stored space management values for a table. Use sp_helpindex to see the stored space management values for an index.

- Setting `max_rows_per_page` to 0 tells Adaptive Server to fill the data or index pages and not to limit the number of rows (this is the default behavior of Adaptive Server if `max_rows_per_page` is not set).
- Low values for *optvalue* may cause page splits. Page splits occur when new data or index rows need to be added to a page, and there is not enough room for the new row. Usually, the data on the existing page is split fairly evenly between the newly allocated page and the existing page.
- To approximate the maximum value for a nonclustered index, subtract 32 from the page size and divide the resulting number by the index key size. The following statement calculates the maximum value of `max_rows_per_page` for the nonclustered index `titleind`:

```
select
    (select @@pagesize - 32) / minlen
    from sysindexes where name = "titleind"
-----
                288
```

If you specify too high a value for *optvalue*, Adaptive Server returns an error message specifying the highest value allowed.

- If you specify an incorrect value for `max_rows_per_page`, `fillfactor`, `reservepagegap`, or `exp_row_size`, `sp_chgattribute` returns an error message specifying the valid values.
- For more information on `max_rows_per_page`, `fillfactor`, `reservepagegap`, `exp_row_size`, and `concurrency_opt_threshold`, see the Performance and Tuning Guide.
- For more information about identity gaps, see the section “Managing Identity Gaps in Tables” in Chapter 7, “Creating Databases and Tables” in the *Transact-SQL User’s Guide*.

Permissions

Only the object owner can execute `sp_chgattribute`.

See also

Commands – alter table, create index, create table

System procedures – `sp_helpindex`

sp_clearpsexex

Description	Clears the execution attributes of an Adaptive Server session that was set by sp_setpsexex.
Syntax	sp_clearpsexex <i>spid</i> , <i>exeattr</i>
Parameters	<i>spid</i> – is the process ID of the session for which execution attributes are to be cleared. <i>exeattr</i> – identifies the execution attributes to be cleared. Values for exeattr are "priority" and "enginegroup".
Examples	<pre>sp_clearpsexex 12, 'enginegroup'</pre> <p>Drops the engine group entry for process 12.</p>
Usage	<ul style="list-style-type: none">• sp_clearpsexex clears the execution attributes of the session that was set by sp_setpsexex. For more information, see the Performance and Tuning Guide.• If the execution attributes are not cleared during the lifetime of the session, they are cleared when the session exits or terminates abnormally.• sp_clearpsexex fails if there are no online engines in the associated engine group.• When you drop an engine group entry, the session executes on an engine group determined by a class definition or by the default class.• Use sp_who to list process IDs (spids).
Permissions	Only a System Administrator can execute sp_clearpsexex to clear priority attributes for all users. Any user can execute sp_clearpsexex to clear the priority attributes of tasks owned by that user.
See also	<i>System procedures</i> – sp_addexexclass, sp_bindexexclass, sp_dropexexclass, sp_showexexclass, sp_unbindexexclass

sp_clearstats

Description Initiates a new accounting period for all server users or for a specified user. Prints statistics for the previous period by executing *sp_reportstats*.

Syntax *sp_clearstats* [*loginame*]

Parameters *loginame*
– is the user’s login name.

Examples

Example 1

```

sp_clearstats

Name      Since          CPU  Percent CPU    I/O    Percent I/O
-----
probe    Jun 19 1990      0      0%             0      0%
julie    Jun 19 1990    10000    24.9962%      5000    24.325%
jason    Jun 19 1990    10002    25.0013%      5321    25.8866%
ken      Jun 19 1990    10001    24.9987%      5123    24.9234%
kathy    Jun 19 1990    10003    25.0038%      5111    24.865%
(5 rows affected)

Total CPU    Total I/O
-----
40006        20555

5 login accounts cleared.

```

Initiates a new accounting period for all users.

Example 2

```

sp_clearstats kathy

Name      Since          CPU  Percent CPU    I/O    Percent I/O
-----
KATHY    Jul 24 1990     498    49.8998%     483924  9.1829%
(1 row affected)

Total CPU    Total I/O
-----
998          98392

1 login account cleared.

```

Initiates a new accounting period for the user “kathy.”

Usage

- *sp_clearstats* creates an accounting period and should be run only at the end of a period.
- Because *sp_clearstats* clears out the accounting statistics, you must record the statistics *before* running the procedure.

- sp_clearstats updates the syslogins field accdate and clears the syslogins fields totcpu and totio.

Permissions Only a System Administrator can execute sp_clearstats.
See also *System procedures – sp_reportstats*

sp_cmp_all_qplans

Description Compares all abstract plans in two abstract plan groups.
Syntax sp_cmp_all_qplans *group1, group2* [, *mode*]
Parameters *group1, group2*
– are the names of the 2 abstract plan groups.
mode –
is the display option, one of: counts, brief, same, diff, first, second, offending and full. The default mode is counts.

Examples **Example 1**

```
sp_cmp_all_qplans dev_plans, prod_plans
```

If the two query plans groups are large, this might take some time.

Query plans that are the same

```
count
-----
49
```

Different query plans that have the same association key

```
count
-----
1
```

Query plans present only in group 'dev_plans' :

```
count
-----
1
```

Query plans present only in group 'prod_plans' :

```
count
-----
0
```

Generate a default report on 2 abstract plan groups.

Example 2

`sp_cmp_all_qplans dev_plans, prod_plans, brief`

Generates a report using the brief mode.

Usage

- Use `sp_cmp_all_qplans` to check for differences in abstract plans in two groups of plans.
- `sp_cmp_all_qplans` matches pairs of plans where the plans in each group have the same user ID and query text. The plans are classified as follows:
 - Plans that are the same
 - Plans that have the same association key in both groups, but have different abstract plans. The association key is the group ID, user ID and query text.
 - Plans that exist in one group, but do not exist in the other group

Table 7-7 shows the report modes and what type of information is reported for each mode.

Table 7-7: Report modes for *sp_cmp_all_qplans*

Mode	Reported Information
counts	The counts of: plans that are the same, plans that have the same association key, but different groups, and plans that exist in one group, but not the other. This is the default report mode.
brief	The information provided by counts, plus the IDs of the abstract plans in each group where the plans are different, but the association key is the same, and the IDs of plans that are in one group, but not in the other.
same	All counts, plus the IDs, queries, and plans for all abstract plans where the queries and plans match.
diff	All counts, plus the IDs, queries, and plans for all abstract plans where the queries and plans are different.
first	All counts, plus the IDs, queries, and plans for all abstract plans that are in the first plan group, but not in the second plan group.
second	All counts, plus the IDs, queries, and plans for all abstract plans that are in the second plan group, but not in the first plan group.
offending	All counts, plus the IDs, queries, and plans for all abstract plans that have different association keys or that do not exist in both groups. This is the combination of the diff, first and second modes
full	All counts, plus the IDs, queries, and plans for all abstract plans. This is the combination of same and offending modes.

- To compare two individual abstract plans, use `sp_cmp_qplans`. To see the names of abstract plan groups, use `sp_help_qpgroup`.

- When a System Administrator or Database Owner runs `sp_cmp_all_qplans`, it reports on all plans in the two groups. When another user executes `sp_cmp_all_qplans`, it reports only on plans that have the user's ID.

Permissions Any user can execute `sp_cmp_all_qplans`.

See also System procedures – `sp_cmp_qplans`

sp_cmp_qplans

Description Compares two abstract plans.

Syntax `sp_cmp_qplans id1, id2`

Parameters *id1, id2*
– are the IDs of two abstract plans.

Examples **Example 1**

```
sp_cmp_qplans 411252620, 1383780087
```

The queries are the same.

The query plans are the same.

Example 2

```
sp_cmp_qplans 2091258605, 647777465
```

The queries are the same.

The query plans are different.

Usage

- `sp_cmp_qplans` compares the queries, abstract plans, and hash keys of two abstract plans, and reports whether the queries are the same, and whether the plans are the same. It prints one of these messages for the query:

- The queries are the same.
- The queries are different.
- The queries are different but have the same hash key.

It prints one of these messages for the abstract plan:

- The query plans are the same.
- The query plans are different.

- *sp_cmp_qplans* also prints a return status showing the results of the comparison. The status values 1, 2 and 10 are additive. The status values are show in Table 7-8

Table 7-8: Return status values for *sp_cmp_qplans*

Return value	Meaning
0	The query text and abstract plans are the same.
+1	The queries and hash keys are different.
+2	The queries are different, but the hash keys are the same.
+10	The abstract plans are different.
100	One or both of the plan IDs does not exist.

- To find the ID of a plan, use *sp_help_qpgroup* or *sp_find_qplan*. Plan IDs are also returned by *create plan* and are included in *showplan* output.

Permissions

Any user can execute *sp_cmp_qplans* to compare plans that he or she owns. Only a System Administrator or the Database Owner can compare plans owned by another user.

See also

System procedures – *sp_cmp_all_qplans*, *sp_help_qpgroup*

sp_commonkey

Description

Defines a common key—columns that are frequently joined—between two tables or views.

Syntax

```
sp_commonkey tabaname, tabbname, col1a, col1b
[, col2a, col2b, ..., col8a, col8b]
```

Parameters

tabaname

– is the name of the first table or view to be joined.

tabbname

– is the name of the second table or view to be joined.

col1a

– is the name of the first column in the table or view *tabaname* that makes up the common key. Specify at least one pair of columns (one column from the first table or view and one from the second table or view).

collb

– is the name of the partner column in the table or view *tabbname* that is joined with *colla* in the table or view *tabaname*.

Examples

Example 1

```
sp_commonkey titles, titleauthor, title_id, title_id
```

Defines a common key on titles.titleid and titleauthor.titleid.

Example 2

```
sp_commonkey projects, departments, empid, empid
```

Assumes two tables, projects and departments, each with a column named empid. This statement defines a frequently used join on the two columns.

Usage

- Common keys are created in order to make explicit a logical relationship that is implicit in your database design. The information can be used by an application. sp_commonkey does not enforce referential integrity constraints; use the primary key and foreign key clauses of the create table or alter table command to enforce key relationships.
- Executing sp_commonkey adds the key to the syskeys system table. To display a report on the common keys that have been defined, use sp_helpkey.
- You must be the owner of at least one of the two tables or views in order to define a common key between them.
- The number of columns from the first table or view must be the same as the number of columns from the second table or view. Up to eight columns from each table or view can participate in the common key. The datatypes of the common columns must also agree. For columns that take a length specification, the lengths can differ. The null types of the common columns need not agree.
- The installation process runs sp_commonkey on appropriate columns of the system tables.

Permissions

Only the owner of *tabaname* or *tabbname* can execute sp_commonkey.

See also

Commands – alter table, create table, create trigger

System procedures – sp_dropkey, sp_foreignkey, sp_helpjoins, sp_helpkey, sp_primarykey

sp_companion

Description	Performs cluster operations such as configuring Adaptive Server as a secondary companion in a high availability system and moving a companion server from one failover mode to another. <i>sp_companion</i> is run from the secondary companion.
Syntax	<pre> sp_companion [server_name {,configure [{,with_proxydb NULL}] [,srvlogin] [,server_password] [,cluster_login] [,cluspassword]} drop suspend resume prepare_failback do_advisory} {, all help group attribute_name base attribute_name)} </pre>
Parameters	<p><i>server_name</i></p> <ul style="list-style-type: none"> – Is the name of the Adaptive Server on which you are performing a cluster operation. <p><i>configure</i></p> <ul style="list-style-type: none"> – Configures the server specified by <i>server_name</i> as the primary companion in a failover configuration. <p><i>drop</i></p> <ul style="list-style-type: none"> – Permanently drops a companion from failover configuration. After the command has completed, the servers are in single-server mode. <p><i>suspend</i></p> <ul style="list-style-type: none"> – Temporarily removes the companions from a failover configuration. After the command is completed, the companions are in suspended mode. <p><i>resume</i></p> <ul style="list-style-type: none"> – Reverses the suspend command and resumes normal companion mode between the companions. <p><i>prepare_failback</i></p> <ul style="list-style-type: none"> – Prepare the secondary companion to relinquish the primary companion's resources so it can failback.

do_advisory

– Verifies that the secondary companion is compatible for successfully performing the primary companion’s functions during failover mode.

- all – Causes do_advisory to investigate all the parameters.
- help – Displays information and syntax about the do_advisory parameter.
- group attribute – Limits do_advisory to investigate only the group attributes.
- base attribute – Limits do_advisory to investigate only the base attributes.

with_proxydb –

Creates proxy databases on the secondary companion for all database other than the system databases – and all subsequent databases that are added – when this parameter is included in the initial configuration of the companion servers. By default, with_proxydb is disabled.

srvlogin

– Is a user’s login to access the companion server. By default, the value of srvlogin is “sa”.

srvpassword

– Is the user’s password to access the companion server. By default, the value of srvpassword is null.

cluster_login

–Is the user’s login to log into the cluster. By default, the value of cluster_login is sa.

cluspassword

– Is the users password you must provide to log into the cluster. By default, the value of cluspassword is null.

Examples

Example 1

```
sp_companion "MONEY1", configure
```

Configures the Adaptive Server MONEY1 as the primary companion.

Example 2

```
sp_companion "MONEY1", configure, with_proxydb,  
"sa", "sapsswd"
```

Configures the Adaptive Server MONEY1 as the primary companion and creates proxy databases on the secondary companion.

Example 3

```
sp_companion "PERSONEL1", "drop"
```

Drops the Adaptive Server PERSONEL1 from the failover configuration. After the command has completed, both the primary companion and the secondary companion will be in single-server mode.

Example 4

```
sp_companion "MONEY1", "resume"
```

Resumes normal companion mode for the companion server (in this example, MONEY1).

Example 5

```
sp_companion "PERSONEL1", "prepare_failback"
```

Prepares the primary companion (in this example, PERSONEL1) to change to normal companion mode and resume control of the Adaptive Server that failed over.

Example 6

```
sp_companion "PERSONEL1", do_advisory, "all"
```

Checks to make sure a cluster operation with the PERSONEL1 companion will be successful. Because `do_advisory` in this example uses the `all` parameter, it checks all the `do_advisory` attributes of PERSONEL1 to make sure that none of them will prevent a successful cluster operation, and makes sure that the secondary companion can successfully perform the primary companion's operations after failover is complete.

Example 7

```
sp_companion "PERSONEL1", do_advisory, "CIS"
```

Checks to make sure that none of the attributes for the Component Integration Services (CIS) on the companion server is compatible with the local server.

Usage

- `sp_companion` performs cluster operations such as configuring Adaptive Server as a secondary companion in a high availability system. `sp_companion` also moves companion servers from one failover mode to another (for example, from failover mode back to normal companion mode). `sp_companion` is run from the secondary companion.

- sp_companion is installed with the *installhasvss* (*insthasv* on Windows NT), not the *installmaster* script. *installhasvss* is located in *\$SYBASE/ASE-12_0/scripts*
- sp_companion automatically disables Sybase’s mirroring. Sybase recommends that you use a third-party mirroring software to protect your data from disk failures.

For complete information, see *Using Sybase Failover in A High Availability System*. Before running the *do_advisory* command, make sure to read the configuration chapter of this book as well as the *do_advisory* chapter.

Permissions Only users with the *ha_role* can issue *sp_companion*.

sp_configure

Description Displays or changes configuration parameters.

Syntax `sp_configure [configname [, configvalue] | group_name | non_unique_parameter_fragment]`
`sp_configure "configuration file", 0, {"write" | "read" | "verify" | "restore"} "file_name"`

Parameters

- `sp_configure`
– displays configuration parameters by group, their current values, their default values, the value to which they have most recently been set, and the amount of memory used by this setting. Displays only the parameters whose display level is the same as or below that of the user.
- `sp_configure configname`
– displays the current value, default value, most recently changed value, and amount of memory used by the setting for all parameters matching *parameter*.
- `sp_configure configname, configvalue`
– resets *configname* to *configvalue* and displays the current value, default value, configured value, and amount of memory used by *configname*
- `sp_configure configname, 0, “default”`
– resets *configname* to its default value and displays current value, default value, configured value, and amount of memory used by *configname*.

`sp_configure group_name`

– displays all configuration parameters in `group_name`, their current values, their default values, the value (if applicable) to which they have most recently been set, and the amount of memory used by this setting.

`sp_configure non_unique_parameter_fragment`

– displays all parameter names that match `non_unique_parameter_fragment`, their current values, default values, configured values, and the amount of memory used.

`sp_configure "configuration file", 0, "write", "file_name"`

– creates `file_name` from the current configuration. If `file_name` already exists, a message is written to the error log and the existing file is renamed using the convention `file_name.001`, `file_name.002`, and so on. If you have changed a static parameter but have not restarted your server, "write" gives you the currently running value for that parameter.

`sp_configure "configuration file", 0, "read", "file_name"`

– performs validation checking on values contained in `file_name` and reads those values that pass validation into the server. If any parameters are missing from `file_name`, the current running values for those parameters are used.

`sp_configure "configuration file", 0, "verify", "file_name"`

– performs validation checking on the values in `file_name`.

`sp_configure "configuration file", 0, "restore", "file_name"`

– creates `file_name` with the values in sysconfigures. This is useful if all copies of the configuration file have been lost and you need to generate a new copy.

Examples

Example 1

```
sp_configure
```

Displays all configuration parameters by group, their current values, their default values, the value (if applicable) to which they have most recently been set, and the amount of memory used by this setting.

Example 2

```
sp_configure "identity"
```

Configuration option is not unique.

Parameter Name	Default	Memory Used	Config Value	Run Value
identity burning set factor	5000	0	5000	5000
identity grab size	1	0	1	1

```
size of auto identity column      10          0          10          10
```

Displays all configuration parameters that include the word “identity.”

Example 3

```
sp_configure "recovery interval in minutes", 3
```

Parameter Name	Default	Memory Used	Config Value	Run Value
recovery interval in minutes	5	0	3	3

Configuration option changed. The SQL Server need not be rebooted since the option is dynamic.

Sets the system recovery interval in minutes to 3 minutes.

Example 4

```
sp_configure "number of device", 0, "default"
```

Resets the value for number of devices to the Adaptive Server default.

Usage

- Any user can execute sp_configure to display information about parameters and their current values, but not to modify parameters. System Administrators can execute sp_configure to change the values of most configuration parameters. Only System Security Officers can execute certain parameters. These are listed under Permissions in this section.
- sp_configure allows you to specify the value for configuration parameters in unit specifiers. The unit specifiers are p or P for pages, m or M for megabytes, and g or G for gigabytes. If you do not specify a unit, and you are configuring a parameter that controls memory, Adaptive Server uses the logical page size for the basic unit.
- When you execute sp_configure to modify a dynamic parameter:
 - The configuration and run values are updated.
 - The configuration file is updated.
 - The change takes effect immediately.
- When you execute sp_configure to modify a static parameter:
 - The configuration value is updated.
 - The configuration file is updated.
 - The change takes effect only when you restart Adaptive Server.

- When issued with no parameters, *sp_configure* displays a report of all configuration parameters by group, their current values, their default values, the value (if applicable) to which they have most recently been set, and the amount of memory used by this setting:
 - The default column in the report displays the value Adaptive Server is shipped with. If you do not explicitly reconfigure a parameter, it retains its default value.
 - The memory used column displays the amount of memory used by the parameter at its current value. Some related parameters draw from the same memory pool. For instance, the memory used for stack size and stack guard size is already accounted for in the memory used for number of user connections. If you added the memory used by each of these parameters separately, it would total more than the amount actually used. In the memory used column, parameters that “share” memory with other parameters are marked with a hash mark (#).
 - The *config_value* column displays the most recent value to which the configuration parameter has been set with *sp_configure*.
 - The *run_value* column displays the value being used by Adaptive Server. It changes after you modify a parameter’s value with *sp_configure* and, for static parameters, after you restart Adaptive Server. This is the value stored in *syscurconfigs.value*.

Note If the server uses a case-insensitive sort order, *sp_configure* with no parameters returns a list of all configuration parameters and groups in alphabetical order with no grouping displayed.

- Each configuration parameter has an associated display level. There are three display levels:
 - The “basic” level displays only the most basic parameters. It is appropriate for very general server tuning.
 - The “intermediate” level displays parameters that are somewhat more complex, as well as showing you all the “basic” parameters. This level is appropriate for a moderately complex level of server tuning.
 - The “comprehensive” level displays all parameters, including the most complex ones. This level is appropriate for users who do highly detailed server tuning.

The default display level is “comprehensive”. Setting one of the other display levels lets you work with a subset of the configuration parameter, shortening the amount of information displayed by sp_configure.

The syntax for showing your current display level is:

```
sp_displaylevel
```

- For information on the individual configuration parameters, see the System Administration Guide.

Permissions

Any user can execute sp_configure to display information about parameters and their current values.

Only System Administrators and System Security Officers can execute sp_configure to modify configuration parameters.

Only System Security Officers can execute sp_configure to modify values for:

- allow procedure grouping
- allow select on syscomments.text
- allow updates
- audit queue size
- auditing
- current audit table
- remote access
- suspend auditing when full
- systemwide password expiration

System Administrators can modify all other parameters.

See also

Commands – set

System procedures – sp_dboption, sp_displaylevel, sp_helpconfig, sp_monitorconfig

sp_copy_all_qplans

Description	Copies all plans for one abstract plan group to another group.
Syntax	<code>sp_copy_all_qplans src_group, dest_group</code>
Parameters	<p><i>src_group</i></p> <ul style="list-style-type: none"> – is the name of the source abstract plan group. <p><i>dest_group</i></p> <ul style="list-style-type: none"> – is the name of the abstract plan group to which the plans are to be copied.
Examples	<pre>sp_copy_all_qplans dev_plans, ap_stdin</pre> <p>Copies all of the abstract plans in the <code>dev_plans</code> group to the <code>ap_stdin</code> group.</p>
Usage	<ul style="list-style-type: none"> • The destination group must exist before you can copy plans into it. It may contain plans. • <code>sp_copy_all_qplans</code> calls <code>sp_copy_qplan</code> for each plan in the source group. Each plan is copied as a separate transaction, so any problem that keeps <code>sp_copy_all_qplans</code> from completing does not affect the plans that have already been copied. • <code>sp_copy_qplan</code> prints messages when it cannot copy a particular abstract plan. You also see these messages when running <code>sp_copy_all_qplans</code>. • If the query text for a plan in the destination group exactly matches the query text in the source group and the user ID is the same, the plan is not copied, and a message giving the plan ID is sent to the user, but the copying process continues with the next plan in the source group. • Copying a very large number of abstract plans can take considerable time, and also requires space on the system segment in the database and space to log the changes to the database. Use <code>sp_spaceused</code> to check the size of <code>sysqueryplans</code>, and <code>sp_helpsegment</code> for the system and <code>logsegment</code> to check the space available.
Permissions	Any user can execute <code>sp_copy_all_qplans</code> to copy an abstract plan that he or she owns. Only the System Administrator or Database Owner can copy plans that are owned by other users.
See also	<i>System procedures</i> – <code>sp_copy_qplan</code> , <code>sp_help_qpgroup</code>

sp_copy_qplan

Description	Copies one abstract plan to an abstract plan group.
Syntax	<code>sp_copy_qplan src_id, dest_group</code>
Parameters	<i>src_id</i> – is the ID of the abstract plan to copy. <i>dest_group</i> – is the name of the destination abstract plan group.
Examples	<code>sp_copy_qplan 2140534659, ap_stdin</code>
Usage	<ul style="list-style-type: none">• The destination group must exist before you can copy an abstract plan into it. You do not need to specify a source group, since plans are uniquely identified by the plan ID.• A new plan ID is generated when the plan is copied. The plan retains the ID of the user who created it, even if the System Administrator or Database Owner copies the plan. To assign a different user ID, a System Administrator or Database Owner can use <code>sp_export_qpgroup</code> and <code>sp_import_qpgroup</code>.• If the query text for a plan in the destination group exactly matches the query text in the source group and the user ID, the plan is not copied, and a message giving the plan IDs is sent to the user.• To copy all of the plans in an abstract plan group, use <code>sp_copy_all_qplans</code>.
Permissions	Any user can execute <code>sp_copy_qplan</code> to copy a plan that he or she owns. Only the System Administrator or Database Owner can copy plans that are owned by other users.
See also	<i>System procedures</i> – <code>sp_copy_all_qplans</code> , <code>sp_help_qpgroup</code> , <code>sp_help_qplan</code> , <code>sp_import_qpgroup</code>

sp_countmetadata

Description	Displays the number of indexes, objects, or databases in Adaptive Server.
Syntax	<code>sp_countmetadata "configname" [, dbname]</code>
Parameters	<i>configname</i> – is either "open indexes", "open objects", or "open databases".

dbname

– is the name of the database on which to run *sp_countmetadata*. If no database name is given, *sp_countmetadata* provides a total count for all databases.

Examples

Example 1

```
sp_countmetadata "open objects"
```

```
There are 283 user objects in all database(s),
requiring 117.180 Kbytes of memory. The 'open
objects' configuration parameter is currently set to
a run value of 500.
```

Reports on the number of user objects in Adaptive Server. Use this value to set the number of objects allowed in the database, plus space for additional objects and temporary tables. For example:

```
sp_configure "number of open objects", 310
```

Example 2

```
sp_countmetadata "open indexes", pubs2
```

```
There are 21 user indexes in pubs2 database(s),
requiring 8.613 kbytes of memory. The 'open indexes'
configuration parameter is currently set to 600.
```

Reports on the number of indexes in Adaptive Server.

Usage

- *sp_countmetadata* displays the number of indexes, objects, or databases in Adaptive Server, including the number of system databases such as *model* and *tempdb*.
- Avoid running *sp_countmetadata* during Adaptive Server peak times. It can cause contention on the *sysindexes*, *sysobjects*, and *sysdatabases* system tables.
- You can run *sp_countmetadata* on a specified database if you want information on a particular database. However, when configuring caches for indexes, objects, or databases, run *sp_countmetadata* without the *database_name* option.
- The information on memory returned by *sp_countmetadata* can vary by platform. For example, a database on Adaptive Server for Windows NT could have a different *sp_countmetadata* result than the same database on Sun Solaris. Information on the number of user indexes, objects, or databases should be consistent, however.

- sp_countmetadata does not include temporary tables in its calculation. Add 5 percent to the open objects value and 10 percent to the open indexes value to accommodate temporary tables.
- If you specify a nonunique fragment of "open indexes", "open objects", or "open databases" for *configname*, sp_countmetadata returns a list of matching configuration parameter names with their configured values and current values. For example:

```
sp_countmetadata "open"
Configuration option is not unique.
option_name                config_value run_value
-----
current change w/ open cursors          1          1
number of open databases                 12          12
number of open indexes                   500         500
number of open objects                   500         500
open index hash spinlock ratio           100         100
open index spinlock ratio                 100         100
open object spinlock ratio                100         100
```

Permissions Only a System Administrator or the Database Owner can execute sp_countmetadata.

See also *System procedures* – sp_configure, sp_helpconfig, sp_monitorconfig

sp_cursorinfo

Description Reports information about a specific cursor or all cursors that are active for your session.

Syntax sp_cursorinfo [{*cursor_level* | null}] [, *cursor_name*]

Parameters

cursor_level | null

– is the level at which Adaptive Server returns information for the cursors. You can specify the following for *cursor_level*:

Level	Types of Cursors
<i>N</i>	Any cursors declared inside stored procedures at a specific procedure nesting level. You can specify any positive number for its level.
0	Any cursors declared outside stored procedures.
-1	Any cursors from either of the above. You can substitute any negative number for this level.

If you want information about cursors with a specific *cursor_name*, regardless of cursor level, specify null for this parameter.

cursor_name

– is the specific name for the cursor. Adaptive Server reports information about all active cursors that use this name at the *cursor_level* you specify. If you omit this parameter, Adaptive Server reports information about all the cursors at that level.

Examples

Example 1

```
sp_cursorinfo 0, authors_crshr
```

```
Cursor name 'authors_crshr' is declared at nesting level '0'.
The cursor id is 327681
The cursor has been successfully opened 1 times.
The cursor was compiled at isolation level 0.
The cursor is not open.
The cursor will remain open when a transaction is committed or rolled back.
The number of rows returned for each FETCH is 1.
The cursor is read only.
There are 3 columns returned by this cursor.
The result columns are:
Name = 'au_id', Table = 'authors', Type = ID,
    Length = 11 (read only)
Name = 'au_lname', Table = 'authors', Type = VARCHAR,
    Length = 40 (read only)
Name = 'au_fname', Table = 'authors', Type = VARCHAR,
    Length = 20 (read only)
```

Displays the information about the cursor named *authors_crshr* at level 0.

Example 2

```
sp_cursorinfo null, author_sales
```

```
Cursor name 'author_sales' is declared on procedure 'au_sales'.
```

```
Cursor name 'author_sales' is declared at nesting level '1'.
The cursor id is 327682
The cursor has been successfully opened 1 times.
The cursor was compiled at isolation level 1.
The cursor is currently scanning at a nonzero isolation level.
The cursor is positioned after the last row.
The cursor will be closed when a transaction is committed or rolled back.
The number of rows returned for each FETCH is 1.
The cursor is updatable.
There are 3 columns returned by this cursor.
The result columns are:
Name = 'title_id', Table = 'titleauthor', Type = ID,
      Length = 11 (updatable)
Name = 'title', Table = 'titles', Type = VARCHAR,
      Length = 80 (updatable)
Name = 'total_sales', Table = 'titles', Type = INT (updatable)
```

Displays the information about any cursors named `author_sales` declared by a user across all levels.

Usage

- If you do not specify either *cursor_level* or *cursor_name*, Adaptive Server displays information about all active cursors. Active cursors are those declared by you and allocated by Adaptive Server.
- Adaptive Server reports the following information about each cursor:
 - The cursor name, its nesting level, its cursor ID, and the procedure name (if it is declared in a stored procedure).
 - The number of times the cursor has been opened.
 - The isolation level (0, 1, or 3) in which it was compiled and in which it is currently scanning (if open).
 - Whether the cursor is open or closed. If the cursor is open, it indicates the current cursor position and the number of rows fetched.
 - Whether the open cursor will be closed if the cursor's current position is deleted.
 - Whether the cursor will remain open or be closed if the cursor's current transaction is committed or rolled back.
 - The number of rows returned for each fetch of that cursor.
 - Whether the cursor is updatable or read-only.

- The number of columns returned by the cursor. For each column, it displays the column name, the table name or expression result, and whether it is updatable.

The output from *sp_cursorinfo* varies, depending on the status of the cursor. In addition to the information listed, *sp_cursorinfo* displays the showplan output for the cursor. For more information about showplan, see the *Performance and Tuning Guide*.

Permissions

Any user can execute *sp_cursorinfo*.

See also

Commands – declare cursor, set

Procedures: *sp_dboption* – *sp_displayroles*

sp_dboption

Description	Displays or changes database options.
Syntax	<code>sp_dboption [dbname, optname, {true false}]</code>
Parameters	<p><i>dbname</i></p> <ul style="list-style-type: none"> – is the name of the database in which the option is to be set. You must be using master to execute <code>sp_dboption</code> with parameters (that is, to change a database option). You cannot, however, change option settings in the master database. <p><i>optname</i></p> <ul style="list-style-type: none"> – is the name of the option to be set. Adaptive Server understands any unique string that is part of the option name. Use quotes around the option name if it is a keyword or includes embedded blanks or punctuation. <p>{true false}</p> <ul style="list-style-type: none"> – true to turn the option on, false to turn it off.
Examples	<p>Example 1</p> <pre>sp_dboption Settable database options database_options ----- abort tran on log full allow nulls by default auto identity dbo use only ddl in tran identity in nonunique index no chkpt on recovery no free space acctg read only select into/bulkcopy/pllsort single user</pre>

```
trunc log on chkpt
trunc. log on chkpt.
unique auto_identity index
```

Displays a list of the database options.

Example 2

```
use pubs2
go
master..sp_dboption pubs2, "read", true
go
checkpoint
go
```

Makes the database pubs2 read only. The read string uniquely identifies the read only option from among all available database options. Note the use of quotes around the keyword read.

Example 3

```
pubs2..sp_dboption pubs2, "read", false
go
checkpoint
go
```

Makes the database pubs2 writable again.

Example 4

```
use pubs2
go
master..sp_dboption pubs2, "select into", true
go
checkpoint
go
```

Allows select into, bcp and parallel sort operations on tables in the pubs2 database. The select into string uniquely identifies the select into/ bulkcopy option from among all available database options. Note that quotes are required around the option because of the embedded space.

Example 5

```
use mydb
go
master..sp_dboption mydb, "auto identity", true
go
checkpoint
go
```


Automatically defines 10-digit IDENTITY columns in new tables created in mydb. The IDENTITY column, SYB_IDENTITY_COL, is defined in each new table that is created without specifying either a primary key, a unique constraint, or an IDENTITY column.

Example 6

```
use master
go
sp_dboption mydb, "identity in nonunique index",
true
go
use mydb
go
checkpoint
go
```

Automatically includes an IDENTITY column in the mydb tables' index keys, provided these tables already have an IDENTITY column. All indexes created on the tables will be internally unique.

Example 7

```
use master
go
sp_dboption pubs2, "unique auto_identity index",
true
go
use pubs2
go
checkpoint
go
```

Automatically includes an IDENTITY column with a unique, nonclustered index for new tables in the pubs2 database.

Usage

- The master database option settings cannot be changed.
- To display a list of database options, execute *sp_dboption* with no parameters from inside the master database.
- For a report on which database options are set in a particular database, execute *sp_helpdb*.
- The Database Owner or System Administrator can set or unset particular database options for all new databases by executing *sp_dboption* on model.

- After sp_dboption has been executed, the change does not take effect until the checkpoint command is issued in the database for which the option was changed.

Database options

- The abort tran on log full option determines the fate of a transaction that is running when the last-chance threshold is crossed in the log segment of the specified database. The default value is false, meaning that the transaction is suspended and is awakened only when space has been freed. If you change the setting to true, all user queries that need to write to the transaction log are killed until space in the log has been freed.
- Setting the allow nulls by default option to true changes the default value of a column from not null to null, in compliance with the SQL standards. The Transact-SQL default value for a column is not null, meaning that null values are not allowed in a column unless null is specified in the create table or alter table column definition. allow nulls by default true reverses this.
- While the auto identity option is set to true (on), a 10-digit IDENTITY column is defined in each new table that is created without specifying either a primary key, a unique constraint, or an IDENTITY column. The column is not visible when you select all columns with the select * statement. To retrieve it, you must explicitly mention the column name, SYB_IDENTITY_COL, in the select list.

To set the precision of the automatic IDENTITY column, use the size of auto identity column configuration parameter.

Though you can set auto identity to true in tempdb, it is not recognized or used, and temporary tables created there do not automatically include an IDENTITY column.

For a report on indexes in a particular table that includes the IDENTITY column, execute sp_helpindex.

- While the dbo use only option is set to true (on), only the database's owner can use the database.

- When the `ddl in tran` option is set to `true (on)`, you can use certain data definition language commands in transactions. If `ddl in tran` is `true` in a particular database, commands such as `create table`, `grant`, and `alter table` are allowed inside transactions in that database. If `ddl in tran` is `true` in the model database, the commands are allowed inside transactions in all databases created after `ddl in tran` was set in model.

Warning! Data definition language commands hold locks on system tables such as `sysobjects`. Avoid using them inside transactions; if you must use them, keep the transactions short.

Using any data definition language commands on `tempdb` within transactions may cause your system to grind to a halt. Always leave `ddl in tran` set to `false` in `tempdb`.

- Table 8-1 lists the commands that can be used inside a user-defined transaction when the `ddl in tran` option is set to `true`:

Table 8-1: DDL commands allowed in transactions

<code>alter table</code>	<code>create default</code>	<code>drop default</code>	<code>grant</code>
(clauses other than partition and unpartition are allowed)	<code>create index</code>	<code>drop index</code>	<code>revoke</code>
	<code>create procedure</code>	<code>drop procedure</code>	
	<code>create rule</code>	<code>drop rule</code>	
	<code>create schema</code>	<code>drop table</code>	
	<code>create table</code>	<code>drop trigger</code>	
	<code>create trigger</code>	<code>drop view</code>	
	<code>create view</code>		

- Table 8-2 lists the commands that cannot be used inside a user-defined transaction under any circumstances:

Table 8-2: DDL commands not allowed in transactions

<code>alter database</code>	<code>dump database</code>	<code>select into</code>
<code>alter table...partition</code>	<code>dump transaction</code>	<code>truncate table</code>
<code>alter table...unpartition</code>	<code>drop database</code>	<code>update statistics</code>
<code>create database</code>	<code>load transaction</code>	
<code>disk init</code>	<code>load database</code>	

In addition, system procedures that create temporary tables or change the master database cannot be used inside user-defined transactions.

- The identity in nonunique index option automatically includes an IDENTITY column in a table's index keys, so that all indexes created on the table are unique. This database option makes logically nonunique indexes internally unique, and allows these indexes to be used to process updatable cursors and isolation level 0 reads.

The table must already have an IDENTITY column for the identity in nonunique index option to work, either from a create table statement or by setting the auto identity database option to true before creating the table.

Use identity in nonunique index if you plan to use cursors and isolation level 0 reads on tables with nonunique indexes. A unique index ensures that the cursor will be positioned at the correct row the next time a fetch is performed on that cursor. If you plan to use cursors on tables with unique indexes and any isolation level, you may want to use the unique auto_identity index option.

For a report on indexes in a particular table that includes the IDENTITY column, execute sp_helpindex.

- The no free space acctg option suppresses free-space accounting and execution of threshold actions for the non-log segments. This speeds recovery time because the free-space counts are not recomputed for those segments.
- The no chkpt on recovery option is set to true (on) when an up-to-date copy of a database is kept. In these situations, there is a "primary" and a "secondary" database. Initially, the primary database is dumped and loaded into the secondary database. Then, at intervals, the transaction log of the primary database is dumped and loaded into the secondary database.

If this option is set to false (off), the default condition, a checkpoint record is added to a database after it is recovered when you restart Adaptive Server. This checkpoint, which ensures that the recovery mechanism will not be rerun unnecessarily, changes the sequence number and causes a subsequent load of the transaction log from the primary database to fail.

Setting this option to true (on) for the secondary database causes it not to get a checkpoint from the recovery process so that subsequent transaction log dumps from the primary database can be loaded into it.

- The read only option means that users can retrieve data from the database, but cannot modify any data.

- Setting the `select into/bulkcopy/pllsort` option to `true` (on) enables the use of `writetext`, `select into` a permanent table, “fast” bulk copy into a table that has no indexes or triggers, using `bcp` or the bulk copy library routines, and parallel sort. A transaction log dump cannot recover these minimally logged operations, so dump transaction to a dump device is prohibited. After non-logged operations are completed, set `select into/bulk copy/pllsort` to `false` (off) and issue `dump database`.

Issuing the `dump transaction` statement after unlogged changes have been made to the database with `select into`, bulk copy, or parallel sort produces an error message instructing you to use `dump database` instead. (The `writetext` command does not have this protection.)

You do not have to set the `select into/bulkcopy/pllsort` option to `true` in order to `select into` a temporary table, since `tempdb` is never recovered. The option need not be set to `true` in order to run `bcp` on a table that has indexes, because tables with indexes are always copied with the slower version of bulk copy and are logged.

- When `single user` is set to `true`, only one user at a time can access the database (single-user mode).

You cannot set `single user` to `true` in a user database from within a stored procedure or while users have the database open. You cannot set `single user` to `true` for `tempdb`.

- The `trunc log on chkpt` option means that if the transaction log has more than 50 rows of committed transactions, the transaction log is truncated (the committed transactions are removed) every time the checkpoint checking process occurs (usually more than once per minute). When the Database Owner runs `checkpoint` manually, however, the log is *not* truncated. It may be useful to turn this option on while doing development work, to prevent the log from growing.

While the `trunc log on chkpt` option is on, `dump transaction` to a dump device is prohibited, since dumps from the truncated transaction log cannot be used to recover from a media failure. Issuing the `dump transaction` statement produces an error message instructing you to use `dump database` instead.

- When the unique auto_identity index option is set to true, it adds an IDENTITY column with a unique, nonclustered index to new tables. By default, the IDENTITY column is a 10-digit numeric datatype, but you can change this default with the size of auto identity column configuration parameter. As with auto identity, the IDENTITY column is not visible when you select all columns with the select * statement. To retrieve it, you must explicitly mention the column name, SYB_IDENTITY_COL, in the select list.

If you need to use cursors or isolation level 0 reads with nonunique indexes, use the identity in nonunique index option.

Though you can set unique auto_identity index to true in tempdb, it is not recognized or used, and temporary tables created there do not automatically include an IDENTITY column with a unique index.

- For more information on database options, see the System Administration Guide.

Permissions

Only a System Administrator or the Database Owner can execute sp_dboption with parameters to change database options. A user aliased to the Database Owner cannot execute sp_dboption to change database options. Any user can execute sp_dboption with no parameters to view database options.

See also

Commands – checkpoint, select

System procedures – sp_configure, sp_helpdb, sp_helpindex, sp_helpjoins

Utilities – bcp

sp_dbrecovery_order

Description

Specifies the order in which user databases are recovered and lists the user-defined recovery order of a database or all databases.

Syntax

```
sp_dbrecovery_order  
    [database_name [, rec_order [, force]]]
```

Parameters

database_name

– The name of the database being assigned a recovery order or the database whose user-defined recovery order is to be listed.

rec_order

- The order in which the database is to be recovered. A *rec_order* of -1 deletes a specified database from the user-defined recovery sequence.

force

- allows the user to insert a database into an existing recovery sequence without putting it at the end.

Examples

Example 1

```
sp_dbrecovery_order pubs2, 1
```

Makes the pubs2 database the first user database to be recovered following a system failure.

Example 2

```
sp_dbrecovery_order pubs3, 3, force
```

Inserts the pubs3 database into third position in a user-defined recovery sequence. If another database was initially in third position, it is moved to fourth position, and all databases following it are moved accordingly.

Example 3

```
sp_dbrecovery_order pubs2, -1
```

Removes the pubs2 database from the user-defined recovery sequence. Subsequently, pubs2 will be recovered after all databases with a user-specified recovery order have recovered.

Example 4

```
sp_dbrecovery_order
```

Lists the current recovery order of all databases with a recovery order assigned through *sp_dbrecovery_order*.

Usage

- You must be in the master database to use *sp_dbrecovery_order* to enter or modify a user-specified recovery order. You can list the user-defined recovery order of databases from any database.
- To change the user-defined recovery position of a database, use *sp_dbrecovery_order* to delete the database from the recovery sequence, then use *sp_dbrecovery_order* to insert it into a new position.
- System databases are always recovered before user databases. The system databases and their recovery order are:
 - master

- model
 - tempdb
 - sybssystemdb
 - sybsecurity
 - sybssystemprocs
- If no database is assigned a recovery order through `sp_dbrecovery_order`, all user databases are recovered in order, by database ID, after system databases.
 - If `database_name` is specified, but no `rec_order` is given, `sp_dbrecovery_order` shows the user-defined recovery position of the specified database.
 - If `database_name` is not specified, `sp_dbrecovery_order` lists the recovery order of all databases with a user-assigned recovery order.
 - The order of recovery assigned through `sp_dbrecovery_order` must be consecutive, starting with 1 and containing no gaps between values. The first database assigned a recovery order must be assigned a `rec_order` of 1. If three databases have been assigned a recovery order of 1, 2, and 3, you cannot assign the next database a recovery order of 5.

Permissions

Only a System Administrator can execute `sp_dbrecovery_order`.

sp_dbremap

Description

Forces Adaptive Server to recognize changes made by `alter database`. Run this procedure only when instructed to do so by an Adaptive Server message.

Syntax

`sp_dbremap dbname`

Parameters

dbname

– is the name of the database in which the `alter database` command was interrupted.

Examples

```
sp_dbremap sample_db
```

An `alter database` command changed the database `sample_db`. This command makes the changes visible to Adaptive Server.

Usage	<ul style="list-style-type: none">• If an alter database statement issued on a database that is in the process of being dumped is interrupted, Adaptive Server prints a message instructing the user to execute <i>sp_dbremap</i>.• Any changes to sysusages during a database or transaction dump are not copied into active memory until the dump completes, to ensure that database mapping does not change during the dump. Running alter database makes changes to system tables on the disk immediately. In-memory allocations cannot be changed until a dump completes. This is why alter database pauses. When you execute <i>sp_dbremap</i>, it must wait until the dump process completes. <ul style="list-style-type: none">• If you are instructed to run <i>sp_dbremap</i>, but do not do it, the space you have allocated with alter database does not become available to Adaptive Server until the next restart.
Permissions	Only a System Administrator or Database Owner can execute <i>sp_dbremap</i> .
See also	<i>Commands</i> – alter database, dump database, dump transaction

sp_defaultloc

Description	<i>Component Integration Services only</i> – Defines a default storage location for objects in a local database.
Syntax	<i>sp_defaultloc</i> <i>dbname</i> , {" <i>defaultloc</i> " NULL} [, " <i>defaulttype</i> "]
Parameters	<i>dbname</i> – is the name of a database being mapped to a remote storage location. The database must already have been defined by a create database statement. You cannot map system databases to a remote location. <i>defaultloc</i> – is the remote storage location to which the database is being mapped. To direct the server to delete an existing default mapping for a database, supply NULL for this parameter. The value of <i>defaultloc</i> must end in a period (.), as follows: <i>server.dbname.owner.</i>

defaulttype

– is one of the values that specify the format of the object named by *object_loc*. Table 8-3 describes the valid values. Enclose the *defaulttype* value in quotes.

Table 8-3: Allowable values for defaulttype

Value	Description
table	Indicates that the object named by <i>object_loc</i> is a table accessible to a remote server. This value is the default for <i>defaulttype</i> .
view	Indicates that the object named by <i>object_loc</i> is a view managed by a remote server, processed as a table.
rpc	Indicates that the object named by <i>object_loc</i> is an RPC managed by a remote server; processes the result set from the RPC as a read-only table.

Examples

Example 1

```
sp_defaultloc pubs, "SYBASE.pubs.dbo.", "table"
create table pubs.dbo.book1 (bridges char(15))
```

sp_defaultloc defines the remote storage location pubs.dbo. in the remote server named SYBASE. It maps the database pubs to the remote location. A “create table book1” statement would create a table named book1 at the remote location. A create existing table statement for bookN would require that pubs.dbo.bookN already exist at the remote location, and information about table bookN would be stored in the local table bookN.

Example 2

```
sp_defaultloc pubs, NULL
```

Removes the mapping of the database pubs to a remote location.

Example 3

```
sp_defaultloc ticktape, "wallst.nasdaq.dbo.", "rpc"
create existing table sybase (bestbuy integer)
```

Identifies the remote storage location wallst.nasdaq.dbo where “wallst” is the value provided for *server_name*, “nasdaq” is provided for *database*, and “dbo” is provided for *owner*. The RPC sybase must already exist at the remote location. A “create existing table sybase” statement would store information about the result set from RPC sybase in local table ticktape. The result set from RPC sybase is regarded as a read-only table. Inserts, updates and deletes are not supported for RPCs.

Usage

- *sp_defaultloc* defines a default storage location for tables in a local database. It maps table names in a database to a remote location. It permits the user to establish a default for an entire database, rather than issue an *sp_addobjectdef* command before every create table and create existing table command.
- When *defaulttype* is table, view, or rpc, the *defaultloc* parameter takes the form:

server_name.dbname.owner.

- Note that the *defaultloc* specification ends in a period (.).
- *server_name* represents a server already added to sys.servers by *sp_addserver*. The *server_name* parameter is required.
- *dbname* might not be required. Some server classes do not support it.
- *owner* should always be provided to avoid ambiguity. If it is not provided, the remote object actually referenced could vary, depending on whether the external login corresponds to the remote object owner.
- Issue *sp_defaultloc* before any create table or create existing table statement. When either statement is used, the server uses the sysattributes table to determine whether any table mapping has been specified for the object about to be created or defined. If the mapping has been specified, a create table statement directs the table to be created at the location specified by *object_loc*. A create existing table statement stores information about the existing remote object in the local table.
- If you issue *sp_defaultloc* on defaulttype view and then issue create table, Component Integration Services creates a new table, not a view, on the remote server.
- Changing the default location for a database does not affect tables that have previously been mapped to a different default location.
- After tables in the database have been created, all future references to tables in *dbname* (by select, insert, delete, and update) are mapped to the correct location.

Permissions

Any user can execute *sp_defaultloc*.

See also

Commands – create existing table, create table

System procedures – *sp_addobjectdef*, *sp_addserver*, *sp_helpserver*

sp_depends

Description Displays information about database object dependencies—the view(s), trigger(s), and procedure(s) in the database that depend on a specified table or view, and the table(s) and view(s) in the database on which the specified view, trigger, or procedure depends.

Syntax sp_depends *objname*

Parameters *objname*
is the name of the table, view, Transact-SQL stored procedure, SQLJ stored procedure, SQLJ function, or trigger to be examined for dependencies. You cannot specify a database name. Use owner names if the object owner is not the user running the command and is not the Database Owner.

Examples **Example 1**

```
sp_depends sysobjects
```

Lists the database objects that depend on the table sysobjects.

Example 2

```
sp_depends titleview
```

Things that the object references in the current database.

object	type	updated	selected
dbo.authors	user table	no	no
dbo.titleauthor	user table	no	no
dbo.titles	user table	no	no

Things inside the current database that reference the object.

object	type
dbo.tvview2	view

Lists the database objects that depend on the titleview view, and the database objects on which the titleview view depends.

Example 3

```
sp_depends "mary.titles"
```

Lists the database objects that depend on the titles table owned by the user “mary”. The quotes are needed, since the period is a special character.

Example 4

The following example shows the column-level dependencies for all columns of the sysobjects table:

```
sp_depends sysobjects
```

Things inside the current database that reference the object.

object	type
dbo.sp_dbupgrade	stored procedure
dbo.sp_procxmode	stored procedure

Dependent objects that reference all columns in the table. Use *sp_depends* on each column to get more information.

Columns referenced in stored procedures, views or triggers are not included in this report.

Column	Type	Object Names or Column Names
cache	permission	column permission
ckfirst	permission	column permission
crdate	permission	column permission
deltrig	permission	column permission
expdate	permission	column permission
id	index	sysobjects (id)
id	logical RI	From syscolumns (id) To sysobjects (id)
id	logical RI	From syscomments (id) To sysobjects (id)
id	logical RI	From sysdepends (id) To sysobjects (id)
id	logical RI	From sysindexes (id) To sysobjects (id)
id	logical RI	From syskeys (depid) To sysobjects (id)
id	logical RI	From syskeys (id) To sysobjects (id)
id	logical RI	From sysobjects (id) To sysprocedures (id)
id	logical RI	From sysobjects (id) To sysprotects (id)
id	logical RI	sysobjects (id)
id	permission	column permission
indexdel	permission	column permission
instrig	permission	column permission
loginame	permission	column permission
name	index	ncsysobjects (name, uid)
name	permission	column permission
objspare	permission	column permission
schemacnt	permission	column permission
seltrig	permission	column permission
sysstat	permission	column permission
sysstat2	permission	column permission
type	permission	column permission

uid	index	ncsysobjects (name, uid)
uid	logical RI	From sysobjects (uid) To sysusers (uid)
uid	permission	column permission
updtrig	permission	column permission
userstat	permission	column permission
versionts	permission	column permission

Example 5

The following example shows more details about the column-level dependencies for the id column of the sysobjects table:

```
sp_depends sysobjects, id
```

Things inside the current database that reference the object.

object	type
-----	-----
dbo.sp_dbupgrade	stored procedure
dbo.sp_procxmode	stored procedure

Dependent objects that reference column id.

Columns referenced in stored procedures, views or triggers are not included in this report.

Type	Property	Object Names or Column Names Also see/Use command
-----	-----	-----
index	index	sysobjects (id) sp_helpindex, drop index, sp_helpconstraint, alter table drop constraint
logical RI	primary	sysobjects (id) sp_helpkey, sp_dropkey
logical RI	foreign	From syskeys (id) To sysobjects (id) sp_helpkey, sp_dropkey
logical RI	common	From syscolumns (id) To sysobjects (id) sp_helpkey, sp_dropkey
logical RI	common	From sysdepends (id) To sysobjects (id) sp_helpkey, sp_dropkey
logical RI	common	From sysindexes (id) To sysobjects (id) sp_helpkey, sp_dropkey
logical RI	common	From syskeys (depid) To sysobjects (id) sp_helpkey, sp_dropkey
logical RI	common	From syscomments (id) To sysobjects (id) sp_helpkey, sp_dropkey
logical RI	common	From sysobjects (id) To sysprotects (id) sp_helpkey, sp_dropkey
logical RI	common	From sysobjects (id) To sysprocedures (id)

```

permission      permission      sp_helpkey, sp_dropkey
                 column permission
                 sp_helprotect, grant/revoke
    
```

Example 6

The following example shows the column-level dependencies for all columns of the user-created table, titles:

```
1> sp_depends titles
```

Things inside the current database that reference the object.

object	type
dbo.delttitle	trigger
dbo.history_proc	stored procedure
dbo.title_proc	stored procedure
dbo.titleid_proc	stored procedure
dbo.titleview	view
dbo.totalsales_trig	trigger

Dependent objects that reference all columns in the table. Use *sp_depends* on each column to get more information. Columns referenced in stored procedures, views or triggers are not included in this report.

Column	Type	Object Names or Column Names
pub_id	logical RI	From titles (pub_id) To publishers (pub_id)
pubdate	default	datedflt
title	index	titleind (title)
title	statistics	(title)
title_id	index	titleidind (title_id)
title_id	logical RI	From roysched (title_id) To titles (title_id)
title_id	logical RI	From salesdetail (title_id) To titles (title_id)
title_id	logical RI	From titleauthor (title_id) To titles (title_id)
title_id	logical RI	titles (title_id)
title_id	rule	title_idrule
title_id	statistics	(title_id)
type	default	typedflt

Example 7

The following example shows more details about the column-level dependencies for the *pub_id* column of the user-created titles table:

```
sp_depends titles, pub_id
```

Things inside the current database that reference the object.

object	type
dbo.delttitle	trigger
dbo.history_proc	stored procedure
dbo.title_proc	stored procedure
dbo.titleid_proc	stored procedure
dbo.titleview	view
dbo.totalsales_trig	trigger

Dependent objects that reference column pub_id.

Columns referenced in stored procedures, views or triggers are not included in this report.

Type	Property	Object Names or Column Names Also see/Use command
logical RI	foreign	From titles (pub_id) To publishers (pub_id) sp_helpkey, sp_dropkey

Usage

- Executing sp_depends lists all objects in the current database that depend on *objname*, and on which *objname* depends. For example, views depend on one or more tables and can have procedures or other views that depend on them. An object that references another object is dependent on that object. References to objects outside the current database are not reported.
- Before you modify or drop a column, use sp_depends to determine if the table contains any dependent objects that could be affected by the modification. For example, if you modify a column to use a new datatype, objects tied to the table may need to be redefined to be consistent with the column's new datatype.
- The sp_depends procedure determines the dependencies by looking at the sysdepends table.

If the objects were created out of order (for example, if a procedure that uses a view was created before the view was created), no rows exist in sysdepends for the dependencies, and sp_depends does not report the dependencies.

- The updated and selected columns in the report from sp_depends are meaningful if the object being reported on is a stored procedure or trigger. The values for the updated column indicate whether the stored procedure updates the object. The selected column indicates whether the object is being used for a read cursor or a data modification statement.

- *sp_depends* follows these Adaptive Server rules for finding objects:
 - If the user does not specify an owner name, and the user executing the command owns an object with the specified name, that object is used.
 - If the user does not specify an owner name, and the user does not own an object of that name, but the Database Owner does, the Database Owner's object is used.
 - If neither the user nor the Database Owner owns an object of that name, the command reports an error condition, even if an object exists in the database with that object name, but with a different owner.
 - If both the user and the Database Owner own objects with the specified name, and the user wants to access the Database Owner's object, the name must be specified, as in *dbo.objectname*.
- Objects owned by database users other than the user executing a command and the Database Owner must always be qualified with the owner's name, as in example 3.
- SQLJ functions and SQLJ stored procedures are Java methods wrapped in SQL wrappers. See *Java in Adaptive Server Enterprise* for more information.
 - SQLJ functions and SQLJ stored procedures are database objects for which you can list dependencies. The only dependencies of SQLJ stored procedures and SQLJ functions are Java classes.
 - If *objname* is a SQLJ stored procedure or SQLJ function, *sp_depends* lists the Java class in the routine's external name declared in the create statement, not classes specified as the return type or datatypes in the parameter list.
 - SQLJ stored procedures and SQLJ functions can be listed as dependencies of other database objects.

Permissions

Any user can execute *sp_depends*.

See also

Commands – create procedure, create table, create view, execute

System procedures – *sp_help*

sp_deviceattr

Description	Changes the dsync setting of an existing database device file.
Syntax	<code>sp_deviceattr logicalname, optname, optvalue</code>
Parameters	<p><i>logicalname</i></p> <ul style="list-style-type: none">– is the name of an existing database device. The device can be stored on either an operating system file or a raw partition, but the dsync setting is ignored for raw partitions. <p><i>optname</i></p> <ul style="list-style-type: none">– is the name of the setting to change. Currently, the only acceptable value for <i>optname</i> is dsync. <p><i>optvalue</i></p> <ul style="list-style-type: none">– can be either “true” or “false.”
Examples	<pre>sp_deviceattr file_device1, dsync, true</pre> <p>Sets dsync on for the device named “file_device1.”</p>
Usage	<ul style="list-style-type: none">• For database devices stored on UNIX files, dsync determines whether updates to the device take place directly on the storage media, or are buffered by the UNIX file system. <p>When dsync is on, writes to the database device occur directly to the physical storage media, and Adaptive Server can recover data on the device in the event of a system failure.</p> <p>When dsync is off, writes to the database device may be buffered by the UNIX file system. The UNIX file system may mark an update as being completed, even though the physical media has not yet been modified. In the event of a system failure, there is no guarantee that requests to update data have ever taken place on the physical media, and Adaptive Server may be unable to recover the database.</p> <ul style="list-style-type: none">• After using sp_deviceattr to change the dsync setting, you must reboot Adaptive Server before the change takes affect.• dsync is always on for the master device file. You cannot change the dsync setting for a master device file with sp_deviceattr.• The dsync value should be turned off only when the databases on the device need not be recovered after a system failure. For example, you may consider turning dsync off for a device that stores only the tempdb database.

- Adaptive Server ignores the *dsync* setting for devices stored on raw partitions—updates to those devices are never buffered, regardless of the *dsync* setting.
- *dsync* is not used on the Windows NT platform.

Permissions The user executing *sp_deviceattr* must have permission to update the *sysdevices* table.

See also *System procedures* – *sp_helpdevice*

sp_diskdefault

Description Specifies whether or not a database device can be used for database storage if the user does not specify a database device or specifies *default* with the *create database* or *alter database* commands.

Syntax *sp_diskdefault logicalname*, {*defaulton* | *defaultoff*}

Parameters *logicalname*
– is the logical name of the device as given in *master.dbo.sysdevices.name*. The device must be a database device rather than a dump device.

defaulton | *defaultoff*
– *defaulton* designates the database device as a default database device; *defaultoff* designates that the specified database device is not a default database device.

Use *defaulton* after adding a database device to the system with *disk init*. Use *defaultoff* to change the default status of the master device (which is designated as a default device when Adaptive Server is first installed).

Examples

```
sp_diskdefault master, defaultoff
```

The master device is no longer used by *create database* or *alter database* for default storage of a database.

Usage

- A default database device is one that is used for database storage by *create database* or *alter database* if the user does not specify a database device name or specifies the keyword *default*.

- You can have multiple default devices. They are used in the order they appear in the master.dbo.sysdevices table (that is, alphabetical order). When the first default device is filled, the second default device is used, and so on.
- When you first install Adaptive Server, the master device is the only default database device.

Note Once you initialize devices to store user databases, use `sp_diskdefault` to turn off the master device's default status. This prevents users from accidentally creating databases on the master device and simplifies recovery of the master database.

- To find out which database devices are default database devices, execute `sp_helpdevice`.

Permissions

Only a System Administrator can execute `sp_diskdefault`.

See also

Commands – alter database, create database, disk init

System procedures – sp_helpdevice

sp_displayaudit

Description

Displays the status of audit options.

Syntax

```
sp_displayaudit ["procedure" | "object" | "login" | "database" | "global" |  
"default_object" | "default_procedure" [, "name"]]
```

Parameters

procedure

– displays the status of audit options for the specified stored procedure or trigger. If you do not specify a value for *name*, `sp_displayaudit` displays the active audit options for all procedures and triggers in the current database.

object

– displays the status of audit options for the specified table or view. If you do not specify a value for *name*, `sp_displayaudit` displays the active audit options for all tables and views in the current database.

login

– displays the status of audit options for the specified user login. If you do not specify a value for *name*, `sp_displayaudit` displays the active audit options for all logins in the master database.

database

– displays the status of audit options for the specified database. If you do not specify a value for *name*, *sp_displayaudit* displays the active audit options for all databases on the server.

global

– displays the status of the specified global audit option. If you do not specify a value for *name*, *sp_displayaudit* displays the active audit options for all procedures and triggers in the current database.

default_object

– displays the default audit options that will be used for any new table or view created on the specified database. If you do not specify a value for *name*, *sp_displayaudit* displays the default audit options for all databases with active default audit settings.

default_procedure

– displays the default audit options that will be used for any new procedure or trigger created on the specified database. If you do not specify a value for *name*, *sp_displayaudit* displays the default audit options for all databases with active default audit settings.

name

– is the information for the specified parameter, as described in the following table:

Parameter	Value for <i>name</i>
procedure	Procedure or trigger name
object	Table or view name
login	User login
database	Database name
global	Global audit option
default_object	Database name
default_procedure	Database name

Examples

Example 1

```

sp_displayaudit

Procedure/Trigger      Audit Option  Value Database
-----
dbo.sp_altermessage    exec_procedure on   sybssystemprocs
dbo.sp_help            exec_procedure on   sybssystemprocs
dbo.sp_who             exec_procedure on   sybssystemprocs
No databases currently have default sproc/trigger auditing enabled.
No objects currently have auditing enabled.
    
```

No databases currently have default table/view auditing enabled.
 No logins currently have auditing enabled.
 No databases currently have auditing enabled.

Option Name	Value
adhoc	off
dbcc	off
disk	off
errors	off
login	off
logout	off
navigator_role	off
oper_role	off
replication_role	off
rpc	off
sa_role	off
security	off
sso_role	off

When no parameter is specified, the status of each category and all auditing options is displayed.

Example 2

```
sp_displayaudit "procedure"
```

Procedure/Trigger	Audit Option	Value	Database
dbo.sp_altermessage	exec_procedure	on	sybsystemprocs
dbo.sp_help	exec_procedure	on	sybsystemprocs
dbo.sp_who	exec_procedure	on	sybsystemprocs

When no procedure name is specified, the status of all procedure audit options is displayed.

Example 3

```
sp_displayaudit "procedure", "sp_who"
```

Procedure/Trigger	Audit Option	Value	Database
dbo.sp_who	exec_procedure	on	sybsystemprocs

When you specify a name for the procedure, only the status of that procedure is displayed.

Example 4

```
sp_displayaudit "global"
```

Option Name	Value
adhoc	off
dbcc	off
disk	off
errors	off
login	off
logout	off
navigator_role	off
oper_role	off
replication_role	off
rpc	off
sa_role	off
security	off
sso_role	off

When no global audit option is specified, the status of all global audit options is displayed.

Usage

- *sp_displayaudit* displays the status of audit options.
- The following table shows the valid auditing options for each parameter:

Object Type Parameter	Valid Auditing Options
procedure	exec_procedure, exec_trigger
object	delete, func_obj_access, insert, reference, select, update
login	all, cmdtext, table_access, view_access
database	alter, bcp, bind, create, dbaccess, drop, dump, func_dbaccess, grant, load, revoke, setuser, truncate, unbind
global	adhoc, dbcc, disk, errors, login, logout, navigator_role, oper_role, replication_role, rpc, sa_role, security, sso_role
default_object	delete, func_obj_access, insert, reference, select, update
default_procedure	exec_procedure, exec_trigger

- You cannot specify a value for name unless you first specify an object type parameter.
- For information on setting up auditing, see the *System Administration Guide*.

Permissions

Only a System Security Officer can execute *sp_displayaudit*.

See also *System procedures – sp_audit*
Utilities – bcp

sp_displaylevel

Description Sets or shows which Adaptive Server configuration parameters appear in `sp_configure` output.

Syntax `sp_displaylevel [loginname [, level]]`

Parameters *loginname*
– is the Adaptive Server login of the user for whom you want to set or show the display level.

level

– sets the display level to one of the following:

- “basic” display level shows just the most basic configuration parameters. This level is appropriate for very general server tuning.
- “intermediate” display level shows configuration parameters that are somewhat more complex, as well as all the “basic” level parameters. This level is appropriate for moderately complex server tuning.
- “comprehensive” display level shows all configuration parameters, including the most complex ones. This level is appropriate for highly detailed server tuning.

Examples

Example 1

```
sp_displaylevel
```

```
The current display level for login 'sa' is  
'comprehensive'.
```

Shows the current display level for the user who invoked `sp_displaylevel`.

Example 2

```
sp_displaylevel jerry
```

```
The current display level for login 'jerry' is  
'intermediate'.
```

Shows the current display level for the user “jerry”.

Example 3

```
sp_displaylevel jerry, comprehensive
```

The display level for login 'jerry' has been changed to 'comprehensive'.

Sets the display level to “comprehensive” for the user “jerry”.

Usage

For details about display levels and configuration parameters, see the *System Administration Guide*.

Permissions

Only a System Administrator can execute *sp_displaylevel* to set the display level for another user. Any user can execute *sp_displaylevel* to set and show his or her own display level.

See also

System procedures – *sp_configure*

sp_displaylogin

Description

Displays information about a login account. Also displays information about the hierarchy tree above or below the login account when you so specify.

Syntax

```
sp_displaylogin [loginame [, expand_up | expand_down]]
```

Parameters

loginame

– is the user login account about which you want information if it is other than your own. You must be a System Security Officer or System Administrator to get information about someone else’s login account.

expand_up

– specifies that Adaptive Server display all roles in the role hierarchy that contain the *loginame* role.

expand_down

– specifies that Adaptive Server display all roles in the role hierarchy that are contained by the *loginame* role.

Examples

Example 1

```
sp_displaylogin

Suid: 1
Loginame: sa
Fullname:
Default Database: master
```

```
Default Language:
Configured Authorization:
    sa_role (default ON)
    sso_role (default ON)
    oper_role (default ON)
Locked: NO
Date of Last Password Change: Nov 16 1994 10:08AM
```

Displays information about your server login account.

Example 2

```
sp_displaylogin susanne

Suid: 12
Loginame: susanne
Fullname:
Default Database: pubs2
Default Language:
Configured Authorization:
    supervisor (default OFF)
Locked: NO
Date of Last Password Change: May 12 1997 11:09AM
```

Displays information about the login account “susanne”. The information displayed varies, depending on the role of the user executing sp_displaylogin.

Example 3

```
sp_displaylogin pillai, expand_up
```

Displays information about all roles containing the role of the login account “pillai”. The information displayed varies, depending on the role of the user executing sp_displaylogin.

Example 4

Displays the login security-related parameters configured for a login.

```
sp_displaylogin joe

Suid: 294
Loginame: joe
Fullname: Joseph Resu
Default Database: master
Default Language:
Configured Authorization: intern_role (default OFF)
Locked: NO
Date of Last Password Change: Nov 24 1998 3:46PM
Password expiration interval : 5
```

```
Password expired : NO
Minimum password length:4
Maximum failed logins : 10
Current failed logins : 3
```

Usage	<ul style="list-style-type: none">• <i>sp_displaylogin</i> displays configured roles, so even if you have made a role inactive with the <i>set</i> command, it is displayed.• When you use <i>sp_displaylogin</i> to get information about your own account, you do not need to use the <i>loginame</i> parameter. <i>sp_displaylogin</i> displays your server user ID, login name, full name, any roles that have been granted to you, date of last password change, default database, default language, and whether your account is locked.• If you are a System Security Officer or System Administrator, you can use the <i>loginame</i> parameter to access information about any account.
Permissions	Only a System Administrator or a System Security Officer can execute <i>sp_displaylogin</i> with the <i>loginame</i> and <i>expand</i> parameters to get information about other users' login accounts. Any user can execute <i>sp_displaylogin</i> to get information about his or her own login account.
See also	<i>Stored procedures</i> – <i>sp_activeroles</i> , <i>sp_displayroles</i> , <i>sp_helprotect</i> , <i>sp_modifylogin</i>

sp_displayroles

Description	Displays all roles granted to another role, or displays the entire hierarchy tree of roles in table format.
Syntax	<i>sp_displayroles</i> [<i>grantee_name</i> [, <i>mode</i>]]
Parameters	<i>grantee_name</i> <ul style="list-style-type: none">– is the login name of a user whose roles you want information about, or the name of a role you want information about.

mode

– is one of the following:

- expand_up, which shows the role hierarchy tree for the parent levels
- expand_down, which shows the role hierarchy tree for the child levels
- display_info, which shows the login security-related parameters configured for the specified role

Examples

Example 1

```
sp_displayroles  
  
Role Name  
-----  
supervisor_role
```

Displays all roles granted to the user issuing the command.

Example 2

```
sp_displayroles "supervisor_role"  
  
Role Name  
-----  
clerk
```

Displays all roles granted to supervisor_role.

Example 3

```
sp_displayroles susanne, expand_down  
  
Role Name          Parent Role Name      Level  
-----  
supervisor_role    NULL                   1  
clerk_role          supervisor_role        2
```

Displays the active roles granted to login “susanne” and the roles below it in the hierarchy.

Example 4

```
sp_displayroles "intern_role", expand_up
```

Displays the active roles granted to intern_role and the roles above it in the hierarchy.

Example 5

```
sp_displayroles physician_role, "display_info"
Role name = physician_role
Locked : NO
Date of Last Password Change : Oct 31 1999 3:33PM
Password expiration interval = 5
Password expired : NO
Minimum password length = 4
Maximum failed logins = 10
Current failed logins = 3
```

Shows the login security-related parameters configured for the specified role.

Usage

- When you specify the optional parameter `expand_up` or `expand_down` all directly granted roles contained by or containing the specified role name are displayed.
- For more information, see “User-Defined Login Security” in the *System Administration Guide*.

Permissions

Only a System Administrator or a System Security Officer can execute `sp_displayroles` to display information on roles activated by any other user. Any user can execute `sp_displayroles` to see his or her own active roles.

See also

Commands – alter role, create role, drop role, grant, revoke, set

System procedures – `sp_activeroles`, `sp_displaylogin`, `sp_helprotect`, `sp_modifylogin`

Procedures: *sp_dropalias* – *sp_dropmessage*

sp_dropalias

Description	Removes the alias user name identity established with <i>sp_addalias</i> .
Syntax	<i>sp_dropalias loginame</i>
Parameters	<i>loginame</i> – is the name (in <i>master.dbo.syslogins</i>) of the user who was aliased to another user.
Examples	<pre>sp_dropalias victoria</pre> Assuming that “victoria” was aliased (for example, to the Database Owner) in the current database, this statement drops “victoria” as an aliased user from the database.
Usage	<ul style="list-style-type: none"> • Executing the <i>sp_dropalias</i> procedure deletes an alternate <i>suid</i> mapping for a user from the <i>sysalternates</i> table. • When a user’s alias is dropped, he or she no longer has access to the database for which the alias was created. • You cannot drop the alias of a user who owns objects in the database that were created in version 12.0 or later. You must drop the objects before dropping the login.
Permissions	Only the Database Owner or a System Administrator can execute <i>sp_dropalias</i> .
See also	<i>System procedures</i> – <i>sp_addalias</i> , <i>sp_adduser</i> , <i>sp_droplogin</i> , <i>sp_dropuser</i> , <i>sp_helpuser</i>

sp_drop_all_qplans

Description	Deletes all abstract plans in an abstract plan group.
Syntax	<i>sp_drop_all_qplans name</i>

Parameters	<i>name</i> – is the name of the abstract plan group from which to drop all plans.
Examples	<code>sp_drop_all_qplans dev_test</code>
Usage	<ul style="list-style-type: none">• To drop individual plans, use <code>sp_drp_qplan</code>.• To see the names of abstract plan groups in the current database, use <code>sp_help_qpgroup</code>.• <code>sp_drop_all_qplans</code> silently drops all plans in the group that belong to the specified user, or all plans in the group, if it is executed by a System Administrator or Database Owner.
Permissions	Any user can execute <code>sp_drop_all_qplans</code> to drop plans that he or she owns. Only a System Administrator or Database Owner can drop plans owned by other users.
See also	<i>System procedures</i> – <code>sp_drop_qplan</code> , <code>sp_help_qpgroup</code>

sp_dropdevice

Description	Drops an Adaptive Server database device or dump device.
Syntax	<code>sp_dropdevice logicalname</code>
Parameters	<i>logicalname</i> – is the name of the device as listed in <code>master.dbo.sysdevices.name</code> .
Examples	<code>sp_dropdevice tape5</code> Drops the device named <code>tape5</code> from Adaptive Server. <code>sp_dropdevice fredsdta</code> Drops the database device named <code>fredsdta</code> from Adaptive Server. The device must not be in use by any database.
Usage	<ul style="list-style-type: none">• The <code>sp_dropdevice</code> procedure drops a device from Adaptive Server, deleting the device entry from <code>master.dbo.sysdevices</code>.• <code>sp_dropdevice</code> does not remove a file that is being dropped as a database device; it makes the file inaccessible to Adaptive Server. Use operating system commands to delete a file after using <code>sp_dropdevice</code>.
Permissions	Only a System Administrator can execute <code>sp_dropdevice</code> .

See also *Commands* – drop database
System procedures – *sp_addumpdevice*, *sp_helpdb*, *sp_helpdevice*

sp_dropengine

Description Drops an engine from a specified engine group or, if the engine is the last one in the group, drops the engine group.

Syntax `sp_dropengine engine_number, engine_group`

Parameters *engine_number*
– is the number of the engine you are dropping from the group. Values are between 0 and a maximum equal to the number of configured online engines, minus one.

engine_group
– is the name of the engine group from which to drop the engine.

Examples
`sp_dropengine 2, DS_GROUP`
This statement drops engine number 2 from the group called DS_GROUP. If it is the last engine in the group, the group is also dropped.

Usage

- *sp_dropengine* can be invoked only from the master database.
- If *engine_number* is the last engine in *engine_group*, Adaptive Server also drops *engine_group*.
- The *engine_number* you specify must exist in *engine_group*.

Permissions Only a System Administrator can execute *sp_dropengine*.

See also *System procedures* – *sp_addengine*

sp_dropexeclass

Description Drops a user-defined execution class.

Syntax `sp_dropexeclass classname`

Parameters *classname*
– is the name of the user-defined execution class to be dropped.

Examples
`sp_dropexeclass 'DECISION'`

Usage	<p>This statement drops the user-defined execution class DECISION.</p> <ul style="list-style-type: none">• An execution class helps define the execution precedence used by Adaptive Server to process tasks. For more information on execution classes and execution attributes, see the Performance and Tuning Guide.• <code>classname</code> must not be bound to any client application, login, or stored procedure. Unbind the execution class first, using <code>sp_unbindexclass</code>, then drop the execution class, using <code>sp_dropexclass</code>.• You cannot drop system-defined execution classes.
Permissions	Only a System Administrator can execute <code>sp_dropexclass</code> .
See also	<i>System procedures</i> – <code>sp_addexclass</code> , <code>sp_bindexclass</code> , <code>sp_showexclass</code> , <code>sp_unbindexclass</code>

sp_dropextendedproc

Description	Removes an extended stored procedure (ESP).
Syntax	<code>sp_dropextendedproc esp_name</code>
Parameters	<i>esp_name</i> – is the name of the extended stored procedure to be dropped.
Examples	<pre>sp_dropextendedproc xp_echo</pre> <p>Removes <code>xp_echo</code>.</p>
Usage	<ul style="list-style-type: none">• <code>sp_dropextendedproc</code> must be executed from the master database.• The <i>esp_name</i> is case sensitive. It must precisely match the name with which the ESP was created.
Permissions	Only a System Administrator can execute <code>sp_dropextendedproc</code> .
See also	<i>Commands</i> – drop procedure <i>System procedures</i> – <code>sp_addextendedproc</code> , <code>sp_freelld</code> , <code>sp_helpextendedproc</code>

sp_dropexternlogin

Description	<i>Component Integration Services only</i> – Drops the definition of a remote login previously defined by <i>sp_addexternlogin</i> .
Syntax	<code>sp_dropexternlogin remote_server [, login_name]</code>
Parameters	<i>remote_server</i> – is the name of the remote server from which the local server is dropping account access. The <i>remote_server</i> is known to the local server by an entry in the <code>master.dbo.sys.servers</code> table. <i>login_name</i> – is a login account known to the local server. If <i>login_name</i> is not specified, the current account is used. <i>login_name</i> must exist in the <code>master.dbo.syslogins</code> table.
Examples	Example 1 <pre>sp_dropexternlogin JOBSERV, sa</pre> Drops the definition of an external login to the remote server JOBSERV from <i>login_name</i> “sa”. Example 2 <pre>sp_dropexternlogin CIS1012, bobj</pre> Drops the definition of an external login to the remote server CIS1012 from “bobj”. Only the “bobj” account and the “sa” account can add or modify a remote login for “bobj”.
Usage	<ul style="list-style-type: none">• <i>sp_dropexternlogin</i> drops the definition of a remote login previously defined to the local server by <i>sp_addexternlogin</i>.• You cannot execute <i>sp_dropexternlogin</i> from within a transaction.• The <i>remote_server</i> must be defined to the local server by <i>sp_addserver</i>.• To add and drop local server users, use the system procedures <i>sp_addlogin</i> and <i>sp_droplogin</i>.
Permissions	Only <i>login_name</i> or a System Administrator can execute <i>sp_dropexternlogin</i> .
See also	<i>System procedures</i> – <i>sp_addexternlogin</i>

sp_dropglockpromote

Description	Removes lock promotion values from a table or database.
Syntax	sp_dropglockpromote {"database" "table"}, <i>objname</i>
Parameters	<p>database table</p> <ul style="list-style-type: none">– specifies whether to remove the lock promotion thresholds from a database or table. The quotes are required because these are Transact-SQL keywords. <p><i>objname</i></p> <ul style="list-style-type: none">– is the name of the table or database from which to remove the lock promotion thresholds.
Examples	<pre>sp_dropglockpromote "table", titles</pre> <p>Removes the lock promotion values from titles. Lock promotion for titles now uses the database or server-wide values.</p>
Usage	<ul style="list-style-type: none">• Use sp_dropglockpromote to drop lock promotion values set with sp_setpglockpromote.• When you drop a database's lock promotion thresholds, tables that do not have lock promotion thresholds configured will use the server-wide values.• When a table's values are dropped, Adaptive Server uses the database's lock promotion thresholds if they are configured or the server-wide values if they are not.• Server-wide values can be changed with sp_setpglockpromote, but cannot be dropped.
Permissions	Only a System Administrator can execute sp_dropglockpromote.
See also	<i>System procedures</i> – sp_setpglockpromote

sp_dropgroup

Description	Drops a group from a database.
Syntax	sp_dropgroup <i>grpname</i>
Parameters	<p><i>grpname</i></p> <ul style="list-style-type: none">– is the name of a group in the current database.
Examples	<pre>sp_changegroup accounting, martha</pre>

```
sp_changegroup "public", george
sp_dropgroup purchasing
```

The “purchasing” group has merged with the “accounting” group. These commands move “martha” and “george”, members of the “purchasing” group, to other groups before dropping the group. The group name “public” is quoted because “public” is a reserved word.

Usage

- Executing *sp_dropgroup* drops a group name from a database’s *sysusers* table.
- You cannot drop a group if it has members. You must execute *sp_changegroup* for each member before you can drop the group.

Permissions

Only the Database Owner, a System Administrator, or a System Security Officer can execute *sp_dropgroup*.

See also

System procedures – *sp_addgroup*, *sp_changegroup*, *sp_helpgroup*

sp_dropkey

Description

Removes from the *syskeys* table a key that had been defined using *sp_primarykey*, *sp_foreignkey*, or *sp_commonkey*.

Syntax

```
sp_dropkey keytype, tablename [, deptabname]
```

Parameters

keytype

– is the type of key to be dropped. The *keytype* must be primary, foreign, or common.

tablename

– is the name of the key table or view that contains the key to be dropped.

deptabname

– specifies the name of the second table in the relationship, if the *keytype* is foreign or common. If the *keytype* is primary, this parameter is not needed, since primary keys have no dependent tables. If the *keytype* is foreign, this is the name of the primary key table. If the *keytype* is common, give the two table names in the order in which they appear with *sp_helpkey*.

Examples**Example 1**

```
sp_dropkey primary, employees
```

Drops the primary key for the employees table. Any foreign keys that were dependent on the primary key for employees are also dropped.

Example 2

```
sp_dropkey common, employees, projects
```

Drops the common keys between the employees and projects tables.

Example 3

```
sp_dropkey foreign, titleauthor, titles
```

Drops the foreign key between the titleauthor and titles tables.

Usage

- Executing sp_dropkey deletes the specified key from syskeys. Only the owner of a table can drop a key from that table.
- Keys are created to make explicit a logical relationship that is implicit in your database design. This information can be used by an application.
- Dropping a primary key automatically drops any foreign keys associated with it. Dropping a foreign key has no effect on a primary key specified on that table.
- Executing sp_commonkey, sp_primarykey, or sp_foreignkey adds the key to the syskeys system table. To display a report on the keys that have been defined, execute sp_helpkey.

Permissions

Only the owner of *tablename* can execute sp_dropkey.

See also

System procedures – sp_commonkey, sp_foreignkey, sp_helpkey, sp_primarykey

sp_droplanguage

Description

Drops an alternate language from the server and removes its row from master.dbo.syslanguages.

Syntax

```
sp_droplanguage language [, dropmessages]
```

Parameters

language
– is the official name of the language to be dropped.

dropmessages

- drops all Adaptive Server system messages in *language*. You cannot drop a language with associated system messages without also dropping its messages.

Examples

Example 1

```
sp_droplanguage french
```

This command drops French from the available alternate languages, if there are no associated messages.

Example 2

```
sp_droplanguage french, dropmessages
```

This command drops French from the available alternate languages, if there are associated messages.

Usage

- Executing *sp_droplanguage* drops a language from a list of alternate languages by deleting its entry from the *master.dbo.syslanguages* table.
- If you try to drop a language that has system messages, the request fails unless you supply the *dropmessages* parameter.

Permissions

Only a System Administrator can execute *sp_droplanguage*.

See also

System procedures – *sp_addlanguage*, *sp_helplanguage*

sp_droplogin

Description

Drops an Adaptive Server user login by deleting the user's entry from *master.dbo.syslogins*.

Syntax

```
sp_droplogin loginame
```

Parameters

loginame

- is the name of the user, as listed in *master.dbo.syslogins*.

Examples

```
sp_droplogin victoria
```

Drops the “victoria” login from Adaptive Server.

Usage

- Executing *sp_droplogin* drops a user login from Adaptive Server, deleting the user's entry from *master.dbo.syslogins*.

- Adaptive Server reuses a dropped login's server user ID, which compromises accountability. You can avoid dropping accounts entirely and, instead, use sp_locklogin to lock any accounts that will no longer be used.

If you need to drop logins, be sure to audit these events (using sp_audit) so that you have a record of them.

- sp_droplogin deletes all resource limits associated with the dropped login.
- sp_droplogin fails if the login to be dropped is a user in any database on the server. Use sp_dropuser to drop the user from a database. You cannot drop a user from a database if that user owns any objects in the database.
- If the login to be dropped is a System Security Officer, sp_droplogin verifies that at least one other unlocked System Security Officer's account exists. If not, sp_droplogin fails. Similarly, sp_droplogin ensures that there is always at least one unlocked System Administrator account.

Permissions Only a System Administrator or a System Security Officer can execute sp_droplogin.

See also *System procedures* – sp_addlogin, sp_locklogin

sp_dropmessage

Description Drops user-defined messages from sysusermessages.

Syntax sp_dropmessage *message_num* [, *language*]

Parameters *message_num*
– is the message number of the message to be dropped. Message numbers must have a value of 20000 or higher.

language
– is the language of the message to be dropped.

Examples sp_dropmessage 20002, french

Removes the French version of the message with the number 20002 from sysusermessages.

Usage	<ul style="list-style-type: none">• The <i>language</i> parameter is optional. If included, only the message with the indicated <i>message_num</i> in the indicated language is dropped. If you do not specify a <i>language</i>, all messages with the indicated <i>message_num</i> are dropped.
Permissions	Only the Database Owner, a System Administrator, or the user who created the message being dropped can execute <i>sp_dropmessage</i> .
See also	<i>System procedures</i> – <i>sp_addmessage</i> , <i>sp_getmessage</i>

Procedures: *sp_dropobjectdef* – *sp_drop_qplan*

sp_dropobjectdef

Description	<i>Component Integration Services only</i> – Deletes the external storage mapping provided for a local object.
Syntax	<code>sp_dropobjectdef "object_name"</code>
Parameters	<p><i>object_name</i> has the form <code>dbname.owner.object</code>, where:</p> <ul style="list-style-type: none"> • <i>dbname</i> is the name of the database containing the object whose storage location you are dropping. <i>dbname</i> is optional; if present, it must be the current database, and the <i>owner</i> or a placeholder is required. • <i>owner</i> is the name of the owner of the object whose storage location you are dropping. <i>owner</i> is optional; it is required if <i>dbname</i> is specified. • <i>object</i> is the name of the local table for which external storage mapping is to be dropped.
Examples	<p>Example 1</p> <pre>sp_dropobjectdef "personnel.dbo.colleges"</pre> <p>Deletes the entry from <code>sysattributes</code> that provided the external storage mapping for a table known to the server as the <code>colleges</code> table in database <code>personnel</code>.</p> <p>Example 2</p> <pre>sp_dropobjectdef "andrea.fishbone"</pre> <p>Deletes the entry from <code>sysattributes</code> that provided the external storage mapping for the <code>andrea.fishbone</code> object, where <code>andrea</code> is the owner and the local table name is <code>fishbone</code>.</p>
Usage	<ul style="list-style-type: none"> • <code>sp_dropobjectdef</code> deletes the external storage mapping provided for a local object. It replaces <code>sp_droptabledef</code>.

- Use sp_dropobjectdef after dropping a remote table with drop table.
- Dropping a table does not remove the mapping information from the sysattributes table if it was added using sp_addobjectdef. It must be explicitly removed using sp_dropobjectdef.
- The *object_name* can be in any of these forms:
 - *object*
 - *owner.object*
 - *dbname..object*
 - *dbname.owner.object*

Permissions Only the Database Owner or a System Administrator can execute sp_dropobjectdef. Only a System Administrator can execute sp_dropobjectdef to remove mapping information for another user's object.

See also *Commands* – create existing table, create table, drop table
System procedures – sp_addobjectdef

sp_drop_qpgroup

Description Drops an abstract plan group.

Syntax sp_drop_qpgroup *group*

Parameters *group*
 – is the name of the abstract plan group to drop.

Examples sp_drop_qpgroup dev_test

Usage

- You cannot drop the default groups, ap_stdin and ap_stdout.
- You cannot drop a group that contains plans. To drop all of the plans in a group, use sp_drop_all_qpplans. To see a list of groups and the number of plans they contain, use sp_help_qpgroup.
- sp_drop_qpgroup cannot be run in a transaction.

Permissions Only a System Administrator or Database Owner can execute sp_drop_qpgroup.

See also *System procedures* – sp_drop_all_qpplans, sp_help_qpgroup

sp_drop_qplan

Description	Drops an abstract plan.
Syntax	<code>sp_drop_qplan id</code>
Parameters	<i>id</i> – is the ID of the abstract plan to drop.
Examples	<code>sp_drop_qplan 1760009301</code> The abstract plan with the specified ID is dropped.
Usage	<ul style="list-style-type: none">• To find the ID of a plan, use <code>sp_help_qpgroup</code>, <code>sp_help_qplan</code>, or <code>sp_find_qplan</code>. Plan IDs are also returned by <code>create plan</code> and are included in <code>showplan</code> output.• To drop all abstract plans in a group, use <code>sp_drop_all_qplans</code>.
Permissions	Any user can execute <code>sp_drop_qplan</code> to drop a plan he or she owns. Only the System Administrator or the Database Owner can drop plans owned by other others.
See also	<i>Commands</i> – <code>create plan</code> <i>System procedures</i> – <code>sp_drop_all_qplans</code> , <code>sp_find_qplan</code> , <code>sp_help_qpgroup</code> , <code>sp_help_qplan</code>

Procedures: *sp_dropremotelogin* – *sp_dropuser*

sp_dropremotelogin

Description	Drops a remote user login.
Syntax	<code>sp_dropremotelogin remoteserver [, loginname [, remotename]]</code>
Parameters	<p><i>remoteserver</i></p> <ul style="list-style-type: none"> – is the name of the server that has the remote login to be dropped. <p><i>loginname</i></p> <ul style="list-style-type: none"> – is the local server’s user name that is associated with the remote server in the sysremotelogins table. <p><i>remotename</i></p> <ul style="list-style-type: none"> – is the remote user name that gets mapped to <i>loginname</i> when logging in from the remote server.
Examples	<p>Example 1</p> <pre>sp_dropremotelogin GATEWAY</pre> <p>Drops the entry for the remote server named GATEWAY.</p> <p>Example 2</p> <pre>sp_dropremotelogin GATEWAY, churchy</pre> <p>Drops the entry for mapping remote logins from the remote server GATEWAY to the local user named “churchy”.</p> <p>Example 3</p> <pre>sp_dropremotelogin GATEWAY, churchy, pogo</pre> <p>Drops the login for the remote user “pogo” on the remote server GATEWAY that was mapped to the local user named “churchy”.</p>
Usage	<ul style="list-style-type: none"> • Executing <code>sp_dropremotelogin</code> drops a user login from a remote server, deleting the user’s entry from <code>master.dbo.sysremotelogins</code>. • For a more complete discussion on remote logins, see <code>sp_addremotelogin</code>.

- To add and drop local server users, use the system procedures `sp_addlogin` and `sp_droplogin`.

Permissions Only a System Administrator can execute `sp_dropremotelogin`.

See also *System procedures* – `sp_addlogin`, `sp_addremotelogin`, `sp_addserver`, `sp_droplogin`, `sp_helpremotelogin`, `sp_helpserver`

sp_drop_resource_limit

Description Removes one or more resource limits from Adaptive Server.

Syntax `sp_drop_resource_limit {name, appname }
[, rangename, limittype, enforced, action, scope]`

Parameters

name
– is the Adaptive Server login to which the limit applies. To drop resource limits that apply to all users of a particular application, specify the *appname* and a *name* of NULL.

appname
– is the application to which the limit applies. To drop resource limits that apply to all applications used by the specified login, specify the login name and an *appname* of NULL. To drop a limit that applies to a particular application, specify the application name that the client program passes to the Adaptive Server in the login packet.

rangename
– is the time range during which the limit is enforced. This must be an existing time range stored in the `systimeranges` system table or NULL to delete all resource limits for the specified *name*, *appname*, *limittype*, *action*, and *scope*, without regard to *rangename*.

limittype
– is the type of resource being limited. This must be one of the following:

Limit Type	Description
row_count	Drops only limits that restrict the number of rows a query can return.
elapsed_time	Drops only limits that restrict the number of seconds that a query batch or transaction can run.
io_cost	Drops only limits that restrict actual or estimated query processing cost.

Limit Type	Description
NULL	Drops all resource limits with the specified <i>name</i> , <i>appname</i> , <i>rangename</i> , enforcement time, <i>action</i> , and <i>scope</i> , without regard to <i>limittype</i> .

enforced

– determines whether the limit is enforced prior to or during query execution. The following table lists the valid values for each limit type:

Enforced Code	Description	Limit Type
1	Drops only limits for which action is taken when the estimated cost of execution exceeds the specified limit.	io_cost
2	Drops only limits for which action is taken when the actual row count, elapsed time, or cost of execution exceeds the specified limit.	row_count elapsed_time io_cost
3	Drops only limits for which action is taken when either the estimated cost (1) or the actual cost (2) exceeds the specified limit.	io_cost
NULL	Drops all resource limits with the specified <i>name</i> , <i>appname</i> , <i>rangename</i> , <i>limittype</i> , and <i>scope</i> , without regard to when the <i>action</i> is enforced.	

action

– is the action taken when the limit is exceeded. This must be one of the following:

Action Code	Description
1	Drops only limits that issue a warning.
2	Drops only limits that abort the query batch.
3	Drops only limits that abort the transaction.
4	Drops only limits that kill the session.
NULL	Drops all resource limits with the specified <i>name</i> , <i>appname</i> , <i>rangename</i> , <i>limittype</i> , enforcement time, and <i>scope</i> , without regard to the <i>action</i> they take.

scope

– is the scope of the limit. This must be one of the following:

Scope Code	Description
1	Drops only limits that apply to queries.
2	Drops only limits that apply to query batches.
4	Drops only limits that apply to transactions.
6	Drops only limits that apply to both query batches and transactions.
NULL	Drops all resource limits with the specified <i>name</i> , <i>appname</i> , <i>rangename</i> , <i>limittype</i> , enforcement time, and <i>action</i> , without regard to their <i>scope</i> .

Examples

Example 1

```
sp_drop_resource_limit joe, payroll,  
friday_afternoon, io_cost, 2, 4, 1
```

Drops the single resource limit that kills the session whenever joe’s use of the *payroll* application runs a query during the *friday_afternoon* time range that results in excessive execution-time I/O cost.

If no resource limit matches these selection criteria, *sp_drop_resource_limit* returns without error.

Example 2

```
sp_drop_resource_limit joe, payroll
```

Drops all limits that apply to joe’s use of the *payroll* application.

Example 3

```
sp_drop_resource_limit joe
```

Drops all limits that apply to the user “joe”.

Example 4

```
sp_drop_resource_limit NULL, payroll
```

Drops all resource limits that apply to the *payroll* application.

Example 5

```
sp_drop_resource_limit NULL, payroll, NULL, NULL,  
NULL, 4, NULL
```

Drops all resource limits on the *payroll* application whose action is to kill the session.

Usage	<ul style="list-style-type: none">• Use the <code>sp_help_resource_limit</code> system procedure to determine which resource limits apply to a given user, application, or time of day.• When you use <code>sp_droplogin</code> to drop an Adaptive Server login, all resource limits associated with that login are also dropped.• The deletion of a resource limit causes the limits for each session for that login and/or application to be rebound at the beginning of the next query batch for that session.• For more information on resource limits, see the System Administration Guide.
Permissions	Only a System Administrator can execute <code>sp_drop_resource_limit</code> .
See also	<i>System procedures</i> – <code>sp_add_resource_limit</code> , <code>sp_help_resource_limit</code> , <code>sp_modify_resource_limit</code>

sp_droprowlockpromote

Description	Removes row lock promotion threshold values from a database or table.
Syntax	<code>sp_droprowlockpromote</code> {"database" "table"}, <i>objname</i>
Parameters	<code>database</code> <code>table</code> – specifies whether to remove the row lock promotion thresholds from a database or table. <i>objname</i> – is the name of the database or table from which to remove the row lock promotion thresholds.
Examples	<pre>sp_droprowlockpromote "table", "sales"</pre> <p>Removes the row lock promotion values from the sales table. Lock promotion for sales now uses the database or server-wide values.</p>
Usage	<ul style="list-style-type: none">• Use <code>sp_droprowlockpromote</code> to drop row lock promotion values set with <code>sp_setrowlockpromote</code>.• When you drop a database's row lock promotion thresholds, datarows-locked tables that do not have row lock promotion thresholds configured use the server-wide values. Use <code>sp_configure</code> to check the value of the row lock promotion configuration parameters.

- When a table's row lock promotion values are dropped, Adaptive Server uses the database's row lock promotion thresholds, if they are configured, or the server-wide values, if no thresholds are set for the database.
- To change the lock promotion thresholds for a database, you must be using the master database. To change the lock promotion thresholds for a table in a database, you must be using the database where the table resides.
- Server-wide values can be changed with `sp_setrowlockpromote`. This changes the values in the row lock promotion configuration parameters, so there is no corresponding server option for `sp_droprowlockpromote`.

Permissions Only a System Administrator can execute `sp_droprowlockpromote`.

See also *System procedures* – `sp_setrowlockpromote`

sp_dropsegment

Description Drops a segment from a database or unmaps a segment from a particular database device.

Syntax `sp_dropsegment segname, dbname [, device]`

Parameters

- segname*
 - is the name of the segment to be dropped.
- dbname*
 - is the name of the database from which the segment is to be dropped.
- device*
 - is the name of the database device from which the segment *segname* is to be dropped. This parameter is optional, except when the system segment system, default, or logsegment is being dropped from a database device.

Examples **Example 1**

```
sp_dropsegment indexes, pubs2
```

This command drops the segment indexes from the pubs2 database.

Example 2

```
sp_dropsegment indexes, pubs2, dev1
```

	<p>This command unmaps the segment indexes from the database device <i>dev1</i>.</p>
Usage	<ul style="list-style-type: none">• You can drop a segment if it is not referenced by any table or index in the specified database.• If you do not supply the optional argument <i>device</i>, the segment is dropped from the specified database. If you do supply a <i>device</i> name, the segment is no longer mapped to the named database device, but the segment is not dropped.• Dropping a segment drops all thresholds associated with that segment.• When you unmap a segment from one or more devices, Adaptive Server drops any thresholds that exceed the total space on the segment. When you unmap the logsegment from one or more devices, Adaptive Server recalculates the last-chance threshold.• <i>sp_placeobject</i> changes future space allocations for a table or index from one segment to another, and removes the references from the original segment. After using <i>sp_placeobject</i>, you can drop the original segment name with <i>sp_dropsegment</i>.• For the system segments <i>system</i>, <i>default</i>, and <i>logsegment</i>, you must specify the device name from which you want the segments dropped.
Permissions	Only the Database Owner or a System Administrator can execute <i>sp_dropsegment</i> .
See also	<i>System procedures</i> – <i>sp_addsegment</i> , <i>sp_addthreshold</i> , <i>sp_helpsegment</i> , <i>sp_helpthreshold</i> , <i>sp_placeobject</i>

sp_dropserver

Description	Drops a server from the list of known servers or drops remote logins and external logins in the same operation.
Syntax	<code>sp_dropserver server [, droplogins]</code>
Parameters	<i>server</i> <ul style="list-style-type: none">– is the name of the server to be dropped. <i>droplogins</i> <ul style="list-style-type: none">– indicates that any remote logins for <i>server</i> should also be dropped.

Examples

Example 1

```
sp_dropserver GATEWAY
```

This command drops the remote server GATEWAY.

Example 2

```
sp_dropserver RDBAM_ALPHA, droplogins
```

Drops the entry for the remote server RDBAM_ALPHA and drops all remote logins and external logins for that server.

Usage

- Executing sp_dropserver drops a server from the list of known servers by deleting the entry from the master.dbo.sysservers table.
- Running sp_dropserver on a server that has associated entries in the master.dbo.sysremotelogins table results in an error message stating that you must drop the remote users before you can drop the server. To drop all remote logins for a server when dropping the server, use droplogins.
- Running sp_dropserver without droplogins against a server that has associated entries in the sysattributes table results in an error. You must drop the remote logins and external logins before you can drop the server.
- The checks against sysattributes for external logins and for default mapping to a server apply when Component Integration Services is configured.

Permissions

Only a System Security Officer can execute sp_dropserver.

See also

System procedures – sp_addserver, sp_dropremotelogs, sp_helpremotelogs, sp_helpserver

sp_dropthreshold

Description

Removes a free-space threshold from a segment.

Syntax

```
sp_dropthreshold dbname, segname, free_space
```

Parameters

dbname

– is the database from which you are dropping the threshold. This must be the name of the current database.

	<i>segname</i> – is the segment whose free space is monitored by the threshold. Use quotes when specifying the “default” segment.
	<i>free_space</i> – is the number of free pages at which the threshold is crossed.
Examples	<code>sp_dropthreshold mydb, segment1, 200</code> Removes a threshold from <code>segment1</code> of <code>mydb</code> . You must specify the database, segment, and amount of free space to identify the threshold.
Usage	<ul style="list-style-type: none">• You cannot drop the last-chance threshold from the log segment.• You can use the <code>no free space acctg</code> option of <code>sp_dboption</code> as an alternative to <code>sp_dropthreshold</code>. This option disables free-space accounting on non-log segments. You cannot disable free-space accounting on log segments.
Permissions	Only the Database Owner or a System Administrator can execute <code>sp_dropthreshold</code> .
See also	<i>System procedures</i> – <code>sp_addthreshold</code> , <code>sp_dboption</code> , <code>sp_helpthreshold</code> , <code>sp_thresholdaction</code>

sp_drop_time_range

Description	Removes a user-defined time range from Adaptive Server.
Syntax	<code>sp_drop_time_range name</code>
Parameters	<i>name</i> – is the name of the time range to be dropped.
Examples	<code>sp_drop_time_range evenings</code> Removes the “evenings” time range.
Usage	<ul style="list-style-type: none">• You cannot remove the “at all times” time range.• You cannot drop a time range if a resource limit exists for that time range.• Dropping a time range does not affect the active time ranges for sessions currently in progress.• For more information on time ranges, see the <i>System Administration Guide</i>.

Permissions Only a System Administrator can execute sp_drop_time_range.
See also *System procedures* – sp_add_resource_limit, sp_add_time_range, sp_modify_time_range

sp_droptype

Description Drops a user-defined datatype.
Syntax `sp_droptype typename`
Parameters *typename*
– is the name of a user-defined datatype that you own.
Examples `sp_droptype birthday`
Drops the user-defined datatype named birthday.
Usage

- Executing sp_droptype deletes a user-defined datatype from systypes.
- A user-defined datatype cannot be dropped if it is referenced by tables or another database object.

Permissions Only the Database Owner or datatype owner can execute sp_droptype.
See also *Datatypes* – User-defined datatypes
System procedures – sp_addtype, sp_rename

sp_dropuser

Description Drops a user from the current database.
Syntax `sp_dropuser name_in_db`
Parameters *name_in_db*
– is the user’s name in the current database’s sysusers table.
Examples `sp_dropuser albert`
Drops the user “albert” from the current database. The user “albert” can no longer use the database.
Usage

- sp_dropuser drops a user from the current database by deleting the user’s row from sysusers.

- You cannot drop a user who owns objects in the database.
- You cannot drop a user who has granted permissions to other users.
- You cannot drop the Database Owner from a database.
- If other users are aliased to the user being dropped, their aliases are also dropped. They no longer have access to the database.
- You cannot drop a user from a database if the user owns a stored procedure that is bound to an execution class in that database. See *sp_bindexeclss*.

Permissions

Only the Database Owner, a System Administrator, or a System Security Officer can execute *sp_dropuser*.

See also

Commands – grant, revoke, use

System procedures – *sp_addalias*, *sp_adduser*, *sp_bindexeclss*, *sp_droplogin*

sp_dumpoptimize

Description	Specifies the amount of data dumped by Backup Server during the dump database operation.
Syntax	<pre>sp_dumpoptimize ['archive_space = {maximum minimum default}'] sp_dumpoptimize ['reserved_threshold = {<i>nnn</i> default}'] sp_dumpoptimize ['allocation_threshold = {<i>nnn</i> default}']</pre>
Parameters	<p>archive_space – specifies the amount of the database you want dumped.</p> <p>maximum – dumps the whole database without determining which pages are allocated or not. The total space used by the archive image or images is equal to the size of the database. Using this option has the same effect as using the options <code>reserved_threshold=0</code> and <code>allocation_threshold=0</code>.</p> <p>minimum – dumps only the allocated pages, which results in the smallest possible archive image. This option is useful when dumping to archive devices for which the throughput is much smaller than that of the database devices such as QIC tape drives. Using this option has the same effect as using the options <code>reserved_threshold=100</code> and <code>allocation_threshold=100</code>.</p> <p>default – specifies that default values should be used.</p> <p>When used with <code>archive_space</code>, this option dumps the database with the <code>reserved_threshold</code> and <code>allocation_threshold</code> options set to their default values. Use this to reset Backup Server to the default configuration.</p> <p>When used with <code>reserved_threshold</code>, default specifies 85 percent.</p> <p>When used with <code>allocation_threshold</code>, default specifies 40 percent.</p>

reserved_threshold

– dumps all the pages belonging to the database in a database disk if the percentage of reserved pages in the disk is equal to or greater than *nnn*. For example, if you specify *nnn* as 60 and if a database disk has a percentage of reserved pages equal to or greater than 60 percent, then the entire disk is dumped without determining which pages within that disk are allocated. The default for this option is 85 percent.

nnn

– an integer value between 0 and 100 that represents the value of the threshold. It is used to determine how much data to dump.

When used with `reserved_threshold`, if the percentage of reserved pages in the disk is greater than the value specified, all the pages of the database in a database disk are dumped.

When used with `allocation_threshold`, if the percentage of allocated pages in an allocation unit is greater than the percentage specified for `allocation_threshold`, all the pages within an allocation unit are dumped.

allocation_threshold

– dumps all the pages in the allocation unit if the percentage of allocated pages in the unit is equal to or greater than *nnn*. For example, if *nnn* is specified as 70 and if the percentage of allocated pages in an allocation unit is equal to or greater than 70 percent, then the entire allocation unit is dumped without determining whether pages within that allocation unit are allocated or not. If the `reserved_threshold` setting causes the whole disk to be dumped, the `allocation_threshold` setting is ignored for the disk. The default for this option is 40 percent.

Examples

Example 1

```
sp_dumpoptimize 'archive_space=maximum'
```

```
Backup Server: 4.172.1.1: The value of 'reserved pages threshold' has been set to 0%.
```

```
Backup Server: 4.172.1.2: The value of 'allocated pages threshold' has been set to 0%.
```

This causes the whole database to be dumped.

Example 2

```
sp_dumpoptimize 'archive_space=minimum'
```

```
Backup Server: 4.172.1.1: The value of 'reserved pages threshold' has been set to 100%.
```

```
Backup Server: 4.172.1.2: The value of 'allocated pages threshold' has been
```

set to 100%.

This causes only the allocated pages to be dumped, thereby resulting in the smallest archive image.

Example 3

```
sp_dumpoptimize 'archive_space=default'
```

Backup Server: 4.172.1.1: The value of 'reserved pages threshold' has been set to 85%.

Backup Server: 4.172.1.2: The value of 'allocated pages threshold' has been set to 40%.

This causes the reserved threshold to be set to 85 percent and the allocation threshold to be set to 40 percent.

Example 4

```
sp_dumpoptimize 'reserved_threshold=60'
```

Backup Server: 4.172.1.3: The value of 'reserved pages threshold' has been set to 60%.

Those disks in the database whose percentage of reserved pages is greater than or equal to 60 percent are dumped without reading allocation pages on this disk. For the remaining disks, the allocation pages are read, and the last set value for the `allocation_threshold` is used. If the `allocation_threshold` was not set after Backup Server was started, default `allocation_threshold` of 40 percent is used.

Example 5

```
sp_dumpoptimize 'reserved_threshold=default'
```

Backup Server: 4.172.1.3: The value of 'reserved pages threshold' has been set to 85%.

This causes the reserved threshold to be set to 85 percent. It does not affect the allocation page threshold.

Example 6

```
sp_dumpoptimize 'allocation_threshold=80'
```

Backup Server: 4.172.1.4: The value of 'allocated pages threshold' has been set to 80%.

Allocation pages are read for those disks whose reserved page percentage is less than the last set value for the `reserved_threshold` and if an allocation unit has 80 percent or more pages allocated, then the whole allocation unit is dumped.

Example 7

```
sp_dumpoptimize 'allocation_threshold=default'
```

Backup Server: 4.172.1.4: The value of 'allocated pages threshold' has been set to 40%.

This causes the allocation page threshold to be set to the default of 40 percent. It does not affect the reserved pages threshold.

Example 8

```
sp_dumpoptimize 'reserved_threshold=60',  
'allocation_threshold=30'
```

Backup Server: 4.172.1.3: The value of 'reserved pages threshold' has been set to 60%.

Backup Server: 4.172.1.4: The value of 'allocated pages threshold' has been set to 30%.

Those disks in the database whose percentage of reserved pages is greater than or equal to 60 percent are dumped without reading allocation pages on this disk. For the remaining disks, the allocation pages are read and if an allocation unit has 30 percent or more pages allocated, then the whole allocation unit is dumped.

Example 9

```
sp_dumpoptimize
```

Backup Server: 4.171.1.1: The current value of 'reserved pages threshold' is 60%

Backup Server: 4.171.1.2: The current value of 'allocated pages threshold' is 30%.

This displays the current value of the thresholds.

Usage

- When you set values with sp_dumpoptimize, those values are immediately in affect without the need to restart Backup Server. However, the changes are effective only until the Backup Server is restarted. When Backup Server is restarted, the default values are used.
- If you issue sp_dumpoptimize multiple times, the thresholds specified by the last instance are used by later dumps. For example, if you first set the reserved_threshold value, and later issue archive_space=maximum, then that value overwrites the previous value you set for reserved_threshold.

- Dumps of different databases can use different thresholds by changing the *sp_dumpoptimize* values before each database dump.
- The optimal threshold values can vary from one database to another. Therefore, the performance of a dump depends on both the I/O configuration and the amount of used space in the database. The DBA can determine the appropriate configuration for a database by experimenting with dumps using different values and choosing the one that results in the shortest dump time.
- You can use *sp_dumpoptimize* for both local and remote dumps.
- *sp_dumpoptimize* has no effect on the performance of a transaction log dump or a load. Therefore, it need not be issued before dump transaction, load database or load transaction operations.
- If *sp_dumpoptimize* is issued without any parameters, the current value of the thresholds is displayed on the client.
- On configurations in which the archive device throughput is equal to or higher than the cumulative throughput of all the database disks, using *archive_space=maximum* may result in a faster dump. However, on configurations in which the archive device throughput is less than the cumulative throughput of all the database disks, using this option may result in a slower dump.
- The option names and the values for this procedure can be abbreviated to the unique substring that identifies them. For example, *ar = ma* is sufficient to uniquely identify the option *archive_space=maximum*.
- There can be zero or more blank space characters around the equal sign (=) in the option string.
- The option names and their values are case insensitive.
- For information on allocation pages, see the *System Administration Guide*.

Permissions

Only the System Administrator, the Database Owner, or users with the Operator role can execute *sp_dumpoptimize*.

See also

Commands – dump database, dump transaction, load database, load transaction

Procedures: *sp_estspace* – *sp_grantlogin*

sp_estspace

Description	Estimates the amount of space required for a table and its indexes, and the time needed to create the index.
Syntax	<code>sp_estspace table_name, no_of_rows [, fill_factor [, cols_to_max [, textbin_len [, iosec]]]]</code>
Parameters	<p><i>table_name</i> – is the name of the table. It must already exist in the current database.</p> <p><i>no_of_rows</i> – is the estimated number of rows that the table will contain.</p> <p><i>fill_factor</i> – is the index fillfactor. The default is null, which means that Adaptive Server uses its default fillfactor.</p> <p><i>cols_to_max</i> – is a comma-separated list of the variable-length columns for which you want to use the maximum length instead of the average. The default is the average declared length of the variable-length columns.</p> <p><i>textbin_len</i> – is the length, per row, of all text and image columns. The default value is 0. You need to provide a value only if the table stores text or image data. text and image columns are stored in a separate set of data pages from the rest of the table's data. The actual table row stores a pointer to the text or image value. <i>sp_estspace</i> provides a separate line of information about the size of the text or image pages for a row.</p> <p><i>iosec</i> – is the number of disk I/Os per second on this machine. The default is 30 I/Os per second.</p>
Examples	<p>Example 1</p> <pre>sp_estspace titles, 10000, 50, "title,notes", 0, 25</pre>

name	type	idx_level	Pages	Kbytes
titles	data	0	3364	6728
titles	text/image	0	0	0
titleidind	clustered	0	21	43
titleidind	clustered	1	1	2
titleind	nonclustered	0	1001	2002
titleind	nonclustered	1	54	107
titleind	nonclustered	2	4	8
titleind	nonclustered	3	1	2

Total_Mbytes

8.68

name	type	total_pages	time_mins
titleidind	clustered	3386	13
titleind	nonclustered	1060	5
titles	data	0	2

Calculates the space requirements for the titles table and its indexes, and the time required to create the indexes. The number of rows is 10,000, the fillfactor is 50 percent, two variable-length columns are computed using the maximum size for the column, and the disk I/O speed is 25 I/Os per second.

Example 2

```
declare @i int
select @i = avg(datalength(pic)) from au_pix
exec sp_estspace au_pix, 1000, null, null, 16, @i
```

au_pix has no indexes

name	type	idx_level	Pages	Kbytes
au_pix	data	0	31	63
au_pix	text/image	0	21000	42000

Total_Mbytes

41.08

Uses the average length of existing image data in the au_pix table to calculate the size of the table with 1000 rows. You can also provide this size as a constant.

Example 3

```

sp_estspace titles, 50000
name          type          idx_level  Pages      Kbytes
-----
titles        data              0          4912      9824
titleidind   clustered         0           31         61
titleidind   clustered         1            1            2
titleind     nonclustered     0          1390      2780
titleind     nonclustered     1            42           84
titleind     nonclustered     2             2            4
titleind     nonclustered     3             1            2

```

```

Total_Mbytes
-----
12.46

```

```

name          type          total_pages  time_mins
-----
titleidind   clustered         4943         19
titleind     nonclustered     1435          8

```

Calculates the size of the titles table with 50,000 rows, using defaults for all other values.

Usage

- To estimate the amount of space required by a table and its indexes:
 - a Create the table.
 - b Create all indexes on the table.
 - c Run *sp_estspace*, giving the table name, the estimated number of rows for the table, and the optional arguments, as needed.

You do not need to insert data into the tables. *sp_estspace* uses information in the system tables—not the size of the data in the tables—to calculate the size of tables and indexes.

- If the auto identity option is set in a database, Adaptive Server automatically defines a 10-digit **IDENTITY** column in each new table that is created without specifying a primary key, a unique constraint, or an **IDENTITY** column. To estimate how much extra space is required by this column:
 - a In the master database, use *sp_dboption* to turn on the auto identity option for the database.
 - b Create the table.

- c Run sp_estspace on the table and record the results.
 - d Drop the table.
 - e Turn the auto identity option off for the database.
 - f Re-create the table.
 - g Rerun sp_estspace on the table, and record the results.
- For information about tables or columns, use sp_help *tablename*.

Permissions

Any user can execute sp_estspace.

See also

Commands – create index, create table

System procedures – sp_help

sp_export_qpgroup

Description

Exports all plans for a specified user and abstract plan group to a user table.

Syntax

sp_export_qpgroup *usr, group, tab*

Parameters

usr

– is the name of the user who owns the abstract plans to be exported.

group

– is the name of the abstract plan group that contains the plans to be exported.

tab

– is the name of a table into which to copy the plans. It must be a table in the current database. You can specify a database name, but not an owner name, in the form *dbname..tablename*. The total length must be 30 characters or less.

Examples

```
sp_export_qpgroup freidak, ap_stdout,  
"tempdb..moveplans"
```

Creates a table called moveplans containing all the plans for the user “freidak” that are in the ap_stdout group.

Usage

- sp_export_qpgroup copies plans from an abstract plan group to a user table. With sp_import_qpgroup, it can be used to copy abstract plans groups between servers and databases or to assign user IDs to copied plans.

- The user table name that you specify cannot exist before you run *sp_export_qpgroup*. The table is created with a structure identical to that of *sysqueryplans*.
- *sp_export_qpgroup* uses *select...into* to create the table to store the copied plans. You must use *sp_dboption* to enable *select into/bulkcopy/pllsort* in order to use *sp_export_qpgroup*, or create the table in *tempdb*.

Permissions Only a System Administrator or the Database Owner can execute *sp_export_qpgroup*.

See also *System procedures* – *sp_copy_all_qplans*, *sp_copy_qplan*, *sp_import_qpgroup*

sp_extendsegment

Description Extends the range of a segment to another database device.

Syntax *sp_extendsegment segname, dbname, devname*

Parameters

segname

– is the name of the existing segment previously defined with *sp_addsegment*.

dbname

– is the name of the database on which to extend the segment. *dbname* must be the name of the current database.

devname

– is the name of the database device to be added to the current database device range already included in *segname*.

Examples `sp_extendsegment indexes, pubs2, dev2`

This command extends the range of the segment *indexes* for the database *pubs2* on the database device *dev2*.

Usage

- A segment can be extended over several database devices.
- If the *logsegment* segment is extended, any other segments on the device are dropped and the device is used for the log segment exclusively.
- When you extend the *logsegment* segment, Adaptive Server recalculates its last-chance threshold.

- To associate a segment with a database device, create or alter the database with a reference to that device. A database device can have more than one segment associated with it.
- After defining a segment, you can use it in the create table and create index commands to place the table or index on the segment. If you create a table or index on a particular segment, subsequent data for the table or index is located on that segment.

Permissions Only the Database Owner or a System Administrator can execute sp_extendsegment.

See also *Commands* – alter database, create index, create table

System procedures – sp_addsegment, sp_dropsegment, sp_helpdb, sp_helpdevice, sp_helpsegment, sp_placeobject

sp_extengine

Description Starts and stops EJB Server. Displays status information about EJB Server.

Syntax sp_extengine 'ejb_server', '{ start | stop | status }'

Parameters *ejb_server*
The logical name of the EJB Server.

start
Starts the EJB Server.

stop
Shuts down the EJB Server.

status
Displays status information about the EJB Server.

Examples **Example 1**

```
sp_extengine 'SYB_EJB', 'status'  
Enterprise java bean server is up and running.
```

Informs user that the EJB Server SYB_EJB is running.

Example 2

```
sp_extengine 'SYB_EJB', 'stop'
```

Shuts down the EJB Server SYB_EJB.

Usage	<ul style="list-style-type: none"> You must have a valid Adaptive Server EJB Server site license to use <i>sp_extengine</i>. See the <i>User's Guide to EJB Server</i> for more information.
Permissions	Only a System Administrator can execute <i>sp_extengine</i> .

sp_familylock

Description	Reports information about all the locks held by a family (coordinating process and its worker processes) executing a statement in parallel.
Syntax	<i>sp_familylock</i> [<i>fpid1</i> [, <i>fpid2</i>]]
Parameters	<p><i>fpid1</i></p> <ul style="list-style-type: none"> – is the family identifier for a family of worker processes from the <i>master.dbo.sysprocesses</i> table. Run <i>sp_who</i> or <i>sp_lock</i> to get the <i>spid</i> of the parent process. <p><i>pfid</i></p> <ul style="list-style-type: none"> 2 – is the Adaptive Server process ID number for another lock.

Examples

```

sp_familylock 5
  fid  spid  locktype      table_id  page  dbname  class      context
  ---  ---  -
  5    5    Sh_intent      176003658  0    userdb  Non cursor lock Sync-pt
duration request
  5    5    Sh_intent-blk  208003772  0    userdb  Non cursor lock Sync-pt
duration request
  5    6    Sh_page        208003772  3972 userdb  Non cursor lock Sync-pt
duration request
  5    7    Sh_page        208003772  3973 userdb  Non cursor lock Sync-pt
duration request
  5    8    Sh_page        208003772  3973 userdb  Non cursor lock Sync-pt
duration request

```

Displays information about the locks held by all members of the family with an *fid* of 5.

Usage	<ul style="list-style-type: none"> <i>sp_familylock</i> with no parameter reports information on all processes belonging to families that currently hold locks. The report is identical to the output from <i>sp_lock</i>; however, <i>sp_familylock</i> allows you to generate reports based on the family ID, rather than the process ID. It is useful for detecting family deadlocks.
-------	---

- Use the `object_name` system function to derive a table's name from its ID number.
- The "locktype" column indicates whether the lock is a shared lock ("Sh" prefix), an exclusive lock ("Ex" prefix) or an update lock, and whether the lock is held on a table ("table" or "intent") or on a page ("page").

The "blk" suffix in the "locktype" column indicates that this process is blocking another process that needs to acquire a lock. As soon as this process completes, the other process(es) moves forward. The "demand" suffix indicates that the process is attempting to acquire an exclusive lock.

- The "class" column indicates whether a lock is associated with a cursor. It displays one of the following:
 - "Non cursor lock" indicates that the lock is not associated with a cursor.
 - "Cursor Id *number*" indicates that the lock is associated with the cursor ID number for that Adaptive Server process ID.
 - A cursor name indicates that the lock is associated with the cursor *cursor_name* that is owned by the current user executing `sp_lock`.
- The "fid" column identifies the family (including the coordinating process and its worker processes) to which a lock belongs. Values for "fid" are as follows:
 - A zero value indicates that the task represented by the `spid` is executed in serial. It is not participating in parallel execution.
 - A nonzero value indicates that the task (`spid`) holding the lock is a member of a family of processes (identified by "fid") executing a statement in parallel. If the value is equal to the `spid`, it indicates that the task is the coordinating process in a family executing a query in parallel.
- The "context" column identifies the context of the lock. Worker processes in the same family have the same context value. Values for "context" are as follows:
 - "NULL" means that the task holding this lock is either executing a query in serial or is a query being executed in parallel in transaction isolation level 1.

- “FAM_DUR” means that the task holding the lock will hold the lock until the query is complete.

A lock’s context may be “FAM_DUR” if the lock is a table lock held as part of a parallel query, if the lock is held by a worker process at transaction isolation level 3, or if the lock is held by a worker process in a parallel query and must be held for the duration of the transaction.

Permissions	Any user can execute <i>sp_familylock</i> .
See also	<i>Commands</i> – kill, select <i>System procedures</i> – <i>sp_lock</i> , <i>sp_who</i>

sp_find_qplan

Description	Finds an abstract plan, given a pattern from the query text or plan text.
Syntax	<i>sp_find_qplan pattern</i> [, group]
Parameters	<i>pattern</i> – is a string to find in the text of the query or abstract plan. <i>group</i> – is the name of the abstract plan group.
Examples	Example 1

```
sp_find_qplan "%from titles%"
```

```
gid          id
           text
-----
2 921054317
select count(*) from titles
2 921054317
( plan
( i_scan t_pub_id_ix titles )
( )
)
( prop titles
( parallel 1 )
( prefetch 16 )
( lru )
```

```
)
    5 937054374
    select type, avg(price) from titles group by type
    5 937054374
    ( plan
    ( store Worktab1
      ( i_scan type_price titles )
    )
    ( t_scan ( work_t Worktab1 ) )
  )
( prop titles
  ( parallel 1 )
  ( prefetch 16 )
  ( lru )
```

Reports on all abstract plans that have the string “from titles” in the query.

Example 2

```
sp_find_qplan "%t_scan%"
```

Finds all plans that include a table scan operator.

Example 3

```
sp_find_qplan "%table[0-9]%", dev_plans
```

Uses the range pattern matching to look for strings such as “table1”, “table2”, and so forth, in plans in the dev_plans group.

Usage

- Use sp_find_qplan to find an abstract plan that contains a particular string. You can match strings from either the query text or from the abstract plan text.
- For each matching plan, sp_find_qplan prints the group ID, plan ID, query text and abstract plan text.
- If you include a group name, sp_find_qplan searches for the string in the specified group. If you do not provide a group name, sp_find_qplan searches all queries and plans for all groups.
- You must supply the “%” wildcard characters, as shown in the examples, unless you are searching for a string at the start or end of a query or plan. You can use any Transact-SQL pattern matching syntax, such as that shown in Example 3.
- The text of queries in sysqueryplans is broken into 255-byte column values. sp_find_qplan may miss matches that span one of these boundaries, but finds all matches that are less than 127 bytes, even if they span two rows.

Permissions	Any user can execute <code>sp_find_qplan</code> . It reports only on abstract plans owned by the user who executes it, except when executed by a System Administrator or the Database Owner.
See also	<i>System procedures</i> – <code>sp_help_qpgroup</code> , <code>sp_help_qplan</code>

sp_flushstats

Description	Flushes statistics from in-memory storage to the <code>systabstats</code> system table.
Syntax	<code>sp_flushstats objname</code>
Parameters	<i>objname</i> – is the name of a table.
Examples	<code>sp_flushstats titles</code> Flushes statistics for the <code>titles</code> table.
Usage	<ul style="list-style-type: none">• Some statistics in the <code>systabstats</code> table are updated in in-memory storage locations and flushed to <code>systabstats</code> periodically, to reduce overhead and contention on <code>systabstats</code>.• If you query <code>systabstats</code> using SQL, executing <code>sp_flushstats</code> guarantees that in-memory statistics are flushed to <code>systabstats</code>.• The <code>optdiag</code> command always flushes in-memory statistics before displaying output.• The statistics in <code>sysstatistics</code> are changed only by data definition language commands and do not require the use of <code>sp_flushstats</code>.
Permissions	Only a System Administrator can execute <code>sp_flushstats</code> .

sp_forceonline_db

Description	Provides access to all the pages in a database that were previously marked suspect by recovery.
Syntax	<code>sp_forceonline_db dbname,</code> <code>{ "sa_on" "sa_off" "all_users" }</code>
Parameters	<i>dbname</i> – is the name of the database to be brought online.

sa_on
– allows only users with the sa_role access to the specified page.

sa_off
– revokes access privileges created by a previous invocation of sp_forceonline_page with sa_on.

all users
– allows all users access to the specified page.

Examples

Example 1

```
sp_forceonline_db pubs2, "sa_on"
```

Allows the System Administrator access to all suspect pages in the pubs2 database.

Example 2

```
sp_forceonline_db pubs2, "sa_off"
```

Revokes access to all suspect pages in the pubs2 database from the System Administrator. Now, no one can access the suspect pages in pubs2.

Example 3

```
sp_forceonline_db pubs2, "all_users"
```

Allows all users access to all pages in the pubs2 database.

Usage

- A page that is forced online is not necessarily repaired. Corrupt pages can also be forced online. Adaptive Server does not perform any consistency checks on pages that are forced online.
- sp_forceonline_page with all users cannot be reversed. When pages have been brought online for all users, you cannot take them offline again.
- sp_forceonline_db cannot be used in a transaction.
- To bring only specific offline pages online, use sp_forceonline_page.

Permissions

Only a System Administrator can execute sp_forceonline_db.

See also

System procedures – sp_forceonline_page, sp_listsuspect_db, sp_listsuspect_page, sp_setsuspect_granularity, sp_setsuspect_threshold

sp_forceonline_object

Description	Provides access to an index previously marked suspect by recovery.
Syntax	<code>sp_forceonline_object dbname, objname, indid, {sa_on sa_off all_users} [, no_print]</code>
Parameters	<p><i>dbname</i> – is the name of the database containing the index to be brought online.</p> <p><i>objname</i> – is the name of the table.</p> <p><i>indid</i> – is the index ID of the suspect index being brought online.</p> <p><i>sa_on</i> – allows only users with the <i>sa_role</i> to access the specified index.</p> <p><i>sa_off</i> – revokes access privileges created by a previous invocation of <i>sp_forceonline_object</i> with <i>sa_on</i>.</p> <p><i>all_users</i> – allows all users to access the specified index.</p> <p><i>no_print</i> – skips printing a list of other suspect objects after the specified object is brought online.</p>
Examples	<p>Example 1</p> <pre>sp_forceonline_object pubs2, titles, 3 , sa_on</pre> <p>Allows a System Administrator to access the index with <i>indid</i> 3 on the <i>titles</i> table in the <i>pubs2</i> database.</p> <p>Example 2</p> <pre>sp_forceonline_object pubs2, titles, 3, sa_off</pre> <p>Revokes access to the index from the System Administrator. Now, no one has access to this index.</p> <p>Example 3</p> <pre>sp_forceonline_object pubs2, titles, 3, all_users</pre> <p>Allows all users to access the index on the <i>titles</i> table in the <i>pubs2</i> database.</p>

Usage

- If an index on a data-only-locked table has suspect pages, the entire index is taken offline during recovery. Offline indexes are not considered by the query optimizer. Indexes on allpages-locked tables are not taken completely offline during recovery; only individual pages of these indexes are taken offline. These pages can be brought online with sp_forceonline_page.
- Use sp_listsuspect_object to see a list of databases that are offline.
- To repair a suspect index, use sp_forceonline_object with sa_on access. Then, drop and re-create the index.

Note If the index is on systabstats or sysstatistics (the only data-only-locked system tables) call Sybase Technical Support for assistance.

- sp_forceonline_object with all_users cannot be reversed. When an index has been brought online for all users, you cannot take it offline again.
- An index that is forced online is not necessarily repaired. Corrupt indexes can be forced online. Adaptive Server does not perform any consistency checks on indexes that are forced online.
- sp_forceonline_object cannot be used in a transaction.
- sp_forceonline_object works only for databases in which the recovery fault isolation mode is "page." Use sp_setsuspect_granularity to display the recovery fault isolation mode for a database.
- To bring all of a database's offline pages and indexes online in a single command, use sp_forceonline_db.
- For more information on recovery fault isolation, see the System Administration Guide.

Permissions

Only a System Administrator can execute sp_forceonline_object.

See also

System procedures – sp_listsuspect_object

sp_forceonline_page

Description

Provides access to pages previously marked suspect by recovery.

Syntax

```
sp_forceonline_page dbname, pgid,  
{"sa_on" | "sa_off" | "all_users"}
```

Parameters	<p><i>dbname</i></p> <p>– is the name of the database containing the pages to be brought online.</p> <p><i>pgid</i></p> <p>– is the page identifier of the page being brought online.</p> <p><i>sa_on</i></p> <p>– allows only users with the <i>sa_role</i> access to the specified page.</p> <p><i>sa_off</i></p> <p>– revokes access privileges created by a previous invocation of <i>sp_forceonline_page</i> with <i>sa_on</i>.</p> <p><i>all_users</i></p> <p>– allows all users access to the specified page.</p>
Examples	<p>Example 1</p> <pre>sp_forceonline_page pubs2, 312, "sa_on"</pre> <p>Allows a System Administrator access to page 312 in the pubs2 database.</p> <p>Example 2</p> <pre>sp_forceonline_page pubs2, 312, "sa_off"</pre> <p>Revokes access to page 312 in the pubs2 database from the System Administrator. Now, no one has access to this page.</p> <p>Example 3</p> <pre>sp_forceonline_page pubs2, 312, "all_users"</pre> <p>Allows all users access to page 312 in the pubs2 database.</p>
Usage	<ul style="list-style-type: none">• <i>sp_forceonline_page</i> with <i>all_users</i> cannot be reversed. When pages have been brought online for all users, you cannot take them offline again.• A page that is forced online is not necessarily repaired. Corrupt pages can also be forced online. Adaptive Server does not perform any consistency checks on pages that are forced online.• <i>sp_forceonline_page</i> cannot be used in a transaction.• <i>sp_forceonline_page</i> works only for databases in which the recovery fault isolation mode is "page." Use <i>sp_setsis[ect_granularity]</i> to display the recovery fault isolation mode for a database.• To bring all of a database's offline pages online in a single command, use <i>sp_forceonline_db</i>.

Permissions Only a System Administrator can use sp_forceonline_page.
See also *System procedures* – sp_forceonline_db, sp_listsuspect_db, sp_listsuspect_page, sp_setsuspect_granularity, sp_setsuspect_threshold

sp_foreignkey

Description Defines a foreign key on a table or view in the current database.

Syntax `sp_foreignkey tablename, pktabname, col1 [, col2] ...
[, col8]`

Parameters

tablename
– is the name of the table or view that contains the foreign key to be defined.

pktabname
– is the name of the table or view that has the primary key to which the foreign key applies. The primary key must already be defined.

col1
– is the name of the first column that makes up the foreign key. The foreign key must have at least one column and can have a maximum of eight columns.

Examples

Example 1

```
sp_foreignkey titles, publishers, pub_id
```

The primary key of the publishers table is the pub_id column. The titles table also contains a pub_id column, which is a foreign key of publishers.

Example 2

```
sp_foreignkey orders, parts, part, subpart
```

The primary key of the parts table has been defined with sp_primarykey as the partnumber and subpartnumber columns. The orders table contains the columns part and subpart, which make up a foreign key of parts.

Usage

- sp_foreignkey adds the key to the syskeys table. Keys make explicit a logical relationship that is implicit in your database design.
- sp_foreignkey does not enforce referential integrity constraints; use the foreign key clause of the create table or alter table command to enforce a foreign key relationship.

- The number and order of columns that make up the foreign key must be the same as the number and order of columns that make up the primary key. The datatypes (and lengths) of the primary and foreign keys must agree, but the null types need not agree.
- The installation process runs *sp_foreignkey* on the appropriate columns of the system tables.
- To display a report on the keys that have been defined, execute *sp_helpkey*.

Permissions

Only the owner of the table or view can execute *sp_foreignkey*.

See also

Commands – alter table, create table, create trigger

System procedures – *sp_commonkey*, *sp_dropkey*, *sp_helpjoins*, *sp_helpkey*, *sp_primarykey*

sp_freedll

Description

Unloads a dynamic link library (DLL) that was previously loaded into XP Server memory to support the execution of an extended stored procedure (ESP).

Syntax

sp_freedll *dll_name*

Parameters

dll_name

– is the file name of the DLL being unloaded from XP Server memory.

Examples

```
sp_freedll "sqlsrvdll.dll"
```

Unloads the *sqlsrvdll.dll* DLL.

Usage

- *sp_freedll* cannot be executed from within a transaction.
- *sp_freedll* cannot free the DLL of a system ESP.
- An alternative to unloading a DLL explicitly, using *sp_freedll*, is to specify that DLLs always be unloaded after the ESP request that invoked them terminates. To do this, set the *esp unload dll* configuration parameter to 1 or start *xpserver* with the *-u* option.
- *sp_freedll* can be used to update an ESP function in a DLL without shutting down XP Server or Adaptive Server.
- If you use *sp_freedll* to unload a DLL that is in use, *sp_freedll* will succeed, causing the ESP currently using the DLL to fail.

Permissions Only a System Administrator can execute sp_freelI.

See also *System procedures* – sp_addextendedproc, sp_dropextendedproc, sp_helpextendedproc

sp_getmessage

Description Retrieves stored message strings from sysmessages and sysusermessages for print and raiserror statements.

Syntax sp_getmessage *message_num*, *result* output [, *language*]

Parameters *message_num*
– is the number of the message to be retrieved.

result output
– is the variable that receives the returned message text, followed by a space and the keyword output. The variable must have a datatype of char, unichar, nchar, varchar, univarchar, or nvarchar.

language
– is the language of the message to be retrieved. *language* must be a valid language name in syslanguages table. If you include *language*, the message with the indicated *message_num* and *language* is retrieved. If you do not include *language*, then the message for the default session language, as indicated by the variable @@langid, is retrieved.

Examples **Example 1**

```
declare @myvar varchar(200)
exec sp_getmessage 20001, @myvar output
```

Retrieves message number 20001 from sysusermessages.

Example 2

```
declare @myvar varchar(200)
exec sp_getmessage 20010, @myvar output, french
```

Retrieves the French language version of message number 20010 from sysusermessages.

Usage

- Any application can use sp_getmessage, and any user can read the messages stored in sysmessages and sysusermessages.

Permissions Any user can execute sp_getmessage.

See also *Commands* – print, raiserror
System procedures – *sp_addressmessage*, *sp_dropmessage*

sp_grantlogin

Description *Windows NT only* – Assigns Adaptive Server roles or default permissions to Windows NT users and groups when Integrated Security mode or Mixed mode (with Named Pipes) is active.

Syntax `sp_grantlogin {login_name | group_name}
[role_list | default]`

Parameters *login_name*
– is the network login name of the Windows NT user.

group_name
– is the Windows NT group name.

role_list
– is a list of the Adaptive Server roles granted. The role list can include one or more of the following role names: *sa_role*, *sso_role*, *oper_role*. If you specify more than one role, separate the role names with spaces, not commas.

default
– specifies that the *login_name* or *group_name* receive default permissions assigned with the grant statement or *sp_role* procedure.

Examples

Example 1

```
sp_grantlogin jeanluc, oper_role
```

Assigns the Adaptive Server *oper_role* to the Windows NT user “jeanluc”.

Example 2

```
sp_grantlogin valle
```

Assigns the default value to the Windows NT user “valle”. User “valle” receives any permissions that were assigned to her via the grant command or *sp_role* procedure.

Example 3

```
sp_grantlogin Administrators, "sa_role sso_role"
```

Assigns the Adaptive Server sa_role and sso_role to all members of the Windows NT administrators group.

Usage

- You must create the Windows NT login name or group before assigning roles with sp_grantlogin. See your Windows NT documentation for details.
- sp_grantlogin is active only when Adaptive Server is running in Integrated Security mode or Mixed mode when the connection is Named Pipes. If Adaptive Server is running under Standard mode or Mixed mode with a connection other than Named Pipes, use grant and sp_role instead.
- If you do not specify a *role_list* or default, the procedure automatically assigns the default value.
- The default value does not indicate an Adaptive Server role. It specifies that the user or group should receive any permissions that were assigned to it via the grant command or sp_role procedure.
- Using sp_grantlogin with an existing *login_name* or *group_name* overwrites the user's or group's existing roles.

Permissions

Only a System Administrator can execute sp_grantlogin.

See also

Commands – grant, setuser

System procedures – sp_addlogin, sp_displaylogin, sp_droplogin, sp_locklogin, sp_logininfo, sp_modifylogin, sp_revokelgin

sp_ha_admin

Description	Performs administrative tasks on Adaptive Servers configured with Sybase Failover in a high availability system. <i>sp_ha_admin</i> is installed with the <i>installhavss</i> script on UNIX platforms or the <i>insthasv</i> script on Windows NT.
Syntax	<i>sp_ha_admin</i> [cleansessions help]
Parameters	<p><i>cleansessions</i></p> <ul style="list-style-type: none"> – Removes old entries from <i>sysessions</i>. Old <i>sysessions</i> entries are typically left behind because either Adaptive Server failed to clean up <i>sysessions</i> during a reboot, or because a client failed to connect to Adaptive Server. <p><i>help</i></p> <ul style="list-style-type: none"> – displays the syntax for <i>sp_ha_admin</i>.
Examples	<p>Example 1</p> <pre>sp_ha_admin cleansessions (return status = 0)</pre> <p>Removes old entries from <i>sysessions</i> left by a client connection that did not exit correctly.</p> <p>Example 2</p> <pre>sp_ha_admin "help" sp_ha_admin Usage: sp_ha_admin command [, option1 [, option2]] sp_ha_admin commands: sp_ha_admin 'cleansessions' sp_ha_admin 'help' (return status = 0)</pre> <p>Displays the syntax for <i>sp_ha_admin</i></p>

Usage

- `sp_ha_admin` performs administrative tasks on Adaptive Server that are configured for Sybase's Failover in a high availability system. `sp_ha_admin` is not installed using the *installmaster* script; instead, use the *installhavss* script that installs and configures for Sybase's Failover (*insthasv* on Windows NT).
- `sp_ha_admin` returns a 0 if it successfully cleaned up `sysessions`, and returns a 1 if it encounters an error.
- `sp_ha_admin` enters a message in the errorlog if it could not remove any entries from `sysessions` (for example, if it could not get a lock on `sysessions`).
- To view all the current entries in `sysessions`, enter:

```
select * from sysessions
```

Permissions

Only the a System Administrator with the `ha_role` can execute `sp_ha_admin`.

Procedures: *sp_help* – *sp_helpindex*

sp_help

Description	Reports information about a database object (any object listed in sysobjects) and about system or user-defined datatypes.
Syntax	<code>sp_help [objname]</code>
Parameters	<p><i>objname</i></p> <ul style="list-style-type: none"> – is the name of any object in sysobjects or any user-defined datatype or system datatype in systypes. You cannot specify database names. <i>objname</i> can include tables, views, stored procedures, logs, rules, defaults, triggers, referential constraints, and check constraints. Use owner names if the object owner is not the user running the command and is not the Database Owner.
Examples	<p>Example 1</p> <pre>sp_help</pre> <p>Displays a list of objects in sysobjects and displays each object's name, owner, and object type. Also displays a list of each user-defined datatype in systypes, indicating the datatype name, storage type, length, null type, default name, and rule name. Null type is 0 (null values not allowed) or 1 (null values allowed).</p> <p>Example 2</p> <pre>sp_help publishers</pre>

```

Name                               Owner                               Type
-----
publishers                          dbo                                  user table

Data_located_on_segment            When_created
-----
default                             Apr 12 1999  3:31PM

Column_name  Type  Length  Prec  Scale  Nulls  Default_name  Rule_name  Identity
-----
pub_id       char  4  NULL  NULL   0  NULL         pub_idrule  0

```

sp_help

```
pub_name varchar      40 NULL NULL      1 NULL      NULL      0
city      varchar      20 NULL NULL      1 NULL      NULL      0
state     char         2  NULL NULL      1 NULL      NULL      0
```

```
attribute_class attribute      int_value char_value      comments
-----
buffer manager  cache binding      1 publishers_cache      NULL
```

```
index_name      index_description      index_keys
index_max_rows_per_page
-----
```

```
pubind          clustered, unique located on default  pub_id
0
```

```
name      attribute_class attribute  int_value char_value
comments
```

```
pubind  buffer manager  cache name      NULL cache for index pubind
NULL
```

```
keytype      object      related_object
object_keys
related_keys
```

```
primary      publishers      - none --
```

```
pub_id, *, *, *, *, *, *, *, *
*, *, *, *, *, *, *, *, *
```

```
foreign      titles      publishers
```

```
pub_id, *, *, *, *, *, *, *, *
pub_id, *, *, *, *, *, *, *, *
```

Object is not partitioned.

Lock scheme Allpages

The attribute 'exp_row_size' is not applicable to tables with allpages lock scheme.

The attribute 'concurrency_opt_threshold' is not applicable to tables with allpages lock scheme.

```
exp_row_size reservepagegap fillfactor max_rows_per_page identity_gap
-----
0 0 0 0 0
```

```
concurrency_opt_threshold
-----
```

0

Displays information about the publishers table. sp_help also lists any attributes assigned to the specified table and its indexes, giving the attribute's class, name, integer value, character value, and comments. The above example shows cache binding attributes for the publishers table.

Example 3

```

                                sp_help partitioned_table
Name                               Owner           Type
-----
partitioned_table                 dbo             user table

Data_located_on_segment           When_created
-----
data1                             Mar 24 1995 10:48AM

Column_name  Type  Length  Prec  Scale  Nulls  Default_name
-----
coll        char    5     NULL  NULL    0     NULL

Rule_name      Identity
-----
NULL          0

Object does not have any indexes.
No defined keys for this object.
partitionid    firstpage    controlpage
-----
              1          145          146
              2          312          313

Lock_scheme Datarows

exp_row_size reservepagegap fillfactor max_rows_per_page identity_gap
-----
              1              0              0              0              0

concurrency_opt_threshold
-----
              15

```

Displays information about a partitioned table.

Example 4

```

                                sp_help "mary.marytrig"
Name                               Owner           Type
-----
marytrig                          mary            trigger

```

```
Data_located_on_segment  When_created
-----
not applicable           Mar 20 1992  2:03PM
```

Displays information about the trigger marytrig owned by user “mary”. The quotes are needed, because the period is a special character.

Example 5

```
sp_help money

Type_name  Storage_type Length Prec  Scale
-----
money      money          8  NULL  NULL
Nulls      Default_name  Rule_name  Identity
-----
1          NULL         NULL      0
```

Displays information about the system datatype money.

Example 6

```
sp_help identype

Type_name  Storage_type Length Nulls Default_name
-----
identype   numeric      4      0  NULL
Rule_name  Identity
-----
NULL      1
```

Displays information about the user-defined datatype identype. The report indicates the base type from which the datatype was created, whether it allows nulls, the names of any rules and defaults bound to the datatype, and whether it has the IDENTITY property.

Example 7

```
sp_help titles

Name          Owner          Type
-----
titles        dbo            user table

Data_located_on_segment  When_created
-----
default                 Dec  6 1997 12:07PM
```

Column_name	Type	Length	Prec	Scale	Nulls	Default_name	Rule_name
Identity							
title_id	tid	6	NULL	NULL	0	NULL	0
title	varchar	80	NULL	NULL	0	NULL	0
type	char	12	NULL	NULL	0	typedflt	0
pub_id	char	4	NULL	NULL	1	NULL	0
price	money	8	NULL	NULL	1	NULL	0
advance	money	8	NULL	NULL	1	NULL	0
total_sales	int	4	NULL	NULL	1	NULL	0
notes	varchar	200	NULL	NULL	1	NULL	0
pubdate	datetime	8	NULL	NULL	0	datedflt	0
contract	bit	1	NULL	NULL	0	NULL	0

```

attribute_class      attribute      int_value
char_value
comments

```

```

lock strategy      row lock promotion      NULL
PCT = 95, LWM = 300, HWM = 300
NULL

```

```

index_name      index_description
index_keys
index_max_rows_per_page index_fillfactor index_reservepagegap

```

```

-----
titleidind      clustered, unique located on default
title_id
0 0 0
titleind      nonclustered located on default
title
0 0 0
type_price      nonclustered located on default
type, price DESC
0 0 0

```

No defined keys for this object.

Object is not partitioned.

Lock scheme Datarows

```

exp_row_size reservedpagegap fillfactor max_rows_per_page identity_gap
-----
224 16 0 0 0
concurrency_opt_threshold

```

0

Reports on the titles table, including information about the locking scheme, expected row size, reserve page gap, and row lock promotion settings.

Usage

- sp_help looks for an object in the current database only.
- sp_help follows the Adaptive Server rules for finding objects:
 - If you do not specify an owner name, and you own an object with the specified name, sp_help reports on that object.
 - If you do not specify an owner name, and do not own an object of that name, but the Database Owner does, sp_help reports on the Database Owner's object.
 - If neither you nor the Database Owner owns an object with the specified name, sp_help reports an error condition, even if an object with that name exists in the database for a different owner. Qualify objects that are owned by database users other than yourself and the Database Owner with the owner's name, as shown in example 4.
 - If both you and the Database Owner own objects with the specified name, and you want to access the Database Owner's object, specify the name in the format *dbo.objectname*.
- sp_help works on temporary tables if you issue it from tempdb.
- Columns with the IDENTITY property have an "Identity" value of 1; others have an "Identity" value of 0. In example 2, there are no IDENTITY columns.
- sp_help lists any indexes on a table, including indexes created by defining unique or primary key constraints in the create table or alter table statements. It also lists any attributes associated with those indexes. However, sp_help does not describe any information about the integrity constraints defined for a table. Use sp_helpconstraint for information about any integrity constraints.
- sp_help displays the following new settings:
 - The locking scheme, which can be set with create table and changed with alter table
 - The expected row size, which can be set with create table and changed with sp_chgattribute

- The reserve page gap, which can be set with `create table` and changed with `sp_chgattribute`
- The row lock promotion settings, which can be set or changed with `sp_setrowlockpromote` and dropped with `sp_droprowlockpromote`
- `sp_help` includes the report from `sp_helpindex`, which shows the order of the keys used to create the index and the space management properties.
- When Component Integration Services is enabled, `sp_help` displays information on the storage location of remote objects.
- `sp_help` reports information about SQLJ stored procedures and SQLJ functions. See *Java in Adaptive Server Enterprise* for more information about SQLJ routines.

Permissions Any user can execute `sp_help`.

See also *Commands* – `create table`, `alter table`

System procedures – `sp_chgattribute`

sp_helppartition

Description Lists the partition number, first page, control page, and number of data pages and summary size information for each partition in a partitioned table.

Syntax `sp_helppartition [table_name]`

Parameters *table_name*
– is the name of a partitioned table in the current database. If the table name is not supplied, the owner, table name, and number of partitions is printed for all user tables in the database.

Examples `sp_helppartition sales`

partitionid	firstpage	controlpage	ptn_data_pages
1	313	314	4227
2	12802	12801	4285
3	25602	25601	4404
4	38402	38401	4523
5	51202	51201	4347

6 64002 64001 4285

(6 rows affected)

Partitions	Average Pages	Maximum Pages	Minimum Pages	Pages Ratio (Max/Avg)
6	4345	4523	4227	1.040967

Returns information about the partitions in *sales*.

Usage

- sp_helppartition lists the partition number, first page, control page, and number of data pages for each partition in a partitioned table. The number of pages per partition shows how evenly the data is distributed between partitions.

The summary information display the number of partitions, the average number of pages per partition, the minimum and maximum number of pages, and the ratio between the average number of pages and the maximum number. This ratio is used during query optimization. If the ratio is 2 or greater (meaning that the maximum size is twice as large as the average size), the optimizer chooses a serial query plan rather than a parallel plan.

- Partitioning a table creates additional page chains. Use the partition clause of the alter table command to partition a table. Each chain has its own last page, which is available for concurrent insert operations. This improves insert performance by reducing page contention. If the table is spread over multiple physical devices, partitioning improves insert performance by reducing I/O contention while Adaptive Server is flushing data from cache to disk.
- Partitioning a table does not affect its performance for update or delete commands.
- Use the unpartition clause of the alter table command to concatenate all existing page chains.
- Neither partitioning nor unpartitioning a table moves existing data.
- To change the number of partitions in a table, first use the unpartition clause of alter table to concatenate its page chains. Then use the partition clause of alter table to repartition the table.
- sp_helppartition looks only in the current database for the table.
- Use sp_helpsegment to display the number of used and free pages on the segment on where the partitioned table is stored.

Accuracy of results

- The values reported in the “data_pages” column may be greater than the actual values. To determine whether the count is inaccurate, run *sp_statistics* and *sp_helppartition* to compare the data page count. The count provided by *sp_statistics* is always accurate.

If the page count reported by *sp_statistics* differs from the sum of the partition pages reported by *sp_helppartition* by more than 5 percent, run one of the following commands to update the partition statistics:

- `dbcc checkalloc`
- `dbcc checkdb`
- `dbcc checktable`
- `update all statistics`
- `update partition statistics`

Then, rerun *sp_helppartition* for an accurate report.

Permissions

Any user can execute *sp_helppartition*.

See also

Commands – `alter table`, `insert`

System procedures – *sp_helpsegment*

sp_helpcache

Description

Displays information about the objects that are bound to a data cache or the amount of overhead required for a specified cache size.

Syntax

```
sp_helpcache {cache_name | "cache_size[P|K|M|G]"}
```

Parameters

cache_name

– is the name of an existing data cache.

cache_size

– specifies the size of the cache, specified by P for pages, K for kilobytes, M for megabytes, or G for gigabytes. The default is K.

Examples

Example 1

```
sp_helpcache pub_cache
```

Displays information about items bound to `pub_cache`.

Example 2

```
sp_helpcache "80M"
```

Shows the amount of overhead required to create an 80MB data cache.

Example 3

```
sp_helpcache
```

Displays information about all caches and all items bound to them.

Usage

- To see the size, status, and I/O size of all data caches on the server, use `sp_cacheconfig`.
- When you configure data caches with `sp_cacheconfig`, all the memory that you specify is made available to the data cache. Overhead for managing the cache is taken from the default data cache. The `sp_helpcache` displays the amount of memory required for a cache of the specified size.
- To bind objects to a cache, use `sp_bindcache`. To unbind a specific object from a cache, use `sp_unbindcache`. To unbind all objects that are bound to a specific cache, use `sp_unbindcache_all`.
- The procedure `sp_cacheconfig` configures data caches. The procedure `sp_poolconfig` configures memory pools within data caches.
- `sp_helpcache` computes overhead accurately up to 74GB.

Permissions

Any user can execute `sp_helpcache`.

See also

System procedures – `sp_bindcache`, `sp_cacheconfig`, `sp_poolconfig`, `sp_unbindcache`, `sp_unbindcache_all`

sp_helpconfig

Description

Reports help information on configuration parameters.

Syntax

```
sp_helpconfig "configname", [size]
```

Parameters

configname

– is the configuration parameter being queried, or a non-unique parameter fragment.

size

– is the size of memory, specified by B (bytes), K (kilobytes), M (megabytes), G (gigabytes), or P (pages). Used without the type of size specified, *size* specifies the number of the entity being configured using this parameter, for examples, locks, open indexes, and so on. *size* is ignored if *configname* is not a unique parameter name.

Examples

Example 1

```
sp_helpconfig "allow"
```

Configuration option is not unique.

option_name	config_value	run_value
allow backward scans	1	1
allow nested triggers	1	1
allow procedure grouping	1	1
allow remote access	1	1
allow resource limits	0	0
allow sendmsg	0	0
allow sql server async i/o	1	1
allow updates to system tables	0	0

Example 2

```
sp_helpconfig "open objects", "421"
```

number of open objects sets the maximum number of database objects that are open at one time on SQL Server. The default run value is 500.

Minimum Value	Maximum Value	Default Value	Current Value	Memory Used
100	2147483647	500	500	243

Configuration parameter, 'number of open objects', will consume 207K of memory if configured at 421.

Returns a report on how much memory is needed to create a metadata cache for 421 object descriptors.

Example 3

```
sp_helpconfig "open databases", "1M"
```

number of open databases sets the maximum number of databases that can be open at one time on SQL Server. The default run value is 12.

Minimum Value	Maximum Value	Default Value	Current Value	Memory Used
5	2147483647	12	12	433

Configuration parameter, 'number of open databases', can be configured to 28 to fit in 1M of memory.

Returns a report on how many database descriptors would fill a 1MB database cache.

Example 4

```
sp_helpconfig "number of locks", "512K"
```

number of locks sets the number of available locks. The default run value is 5000.

Minimum Value	Maximum Value	Default Value	Current Value	Memory Used
1000	2147483647	5000	5000	528

Configuration parameter 'number of locks', can be configured to 4848 to fit in 512K of memory.

Returns a report on how many locks will use 512K of memory.

Example 5

```
sp_helpconfig "allow updates to system tables"
```

allow updates to system tables allows system tables to be updated directly. The default is 0 (off).

Minimum Value	Maximum Value	Default Value	Current Value	Memory Used
0	1	0	0	0

Returns a report on the status of the allow updates to system tables configuration parameter.

Usage

- sp_helpconfig reports help information on configuration parameters, such as how much memory would be needed if the parameter were set to a certain value. sp_helpconfig also displays the current setting, the amount of memory used for that setting, the default value, and the minimum and maximum settings.
- If you use a nonunique parameter fragment for *configname*, sp_helpconfig returns a list of matching parameters with their configured values and current values. See example 1.

Planning metadata cache configuration

- Use sp_helpconfig when you are planning a metadata cache configuration for a server.

For example, suppose you were planning to move a database that contained 2000 user indexes to a different server. To find how much memory you would need to configure for that server so that it would accommodate the database's user indexes, enter the following command:

```
sp_helpconfig "open indexes", "2000"
number of open indexes sets the maximum number of indexes that can be
open at one time on SQL Server. The default run value is 500.
```

Minimum Value	Maximum Value	Default Value	Current Value	Memory Used
-----	-----	-----	-----	-----
100	2147483647	500	500	208

Configuration parameter, 'number of open indexes', will consume 829k of memory if configured at 2000.

Alternatively, suppose you had 1MB of memory available for the index cache, and you needed to know how many index descriptors it would support. Run the following command:

```
sp_helpconfig "open indexes", "1M"
number of open indexes sets the maximum number of indexes that can be
open at one time on SQL Server. The default run value is 500.
```

Minimum Value	Maximum Value	Default Value	Current Value	Memory Used
-----	-----	-----	-----	-----
100	2147483647	500	500	208

Configuration parameter 'number of open indexes', can be configured to 2461 to fit in 1M of memory.

Based on this output, if you have 1MB of memory, you can create an index descriptor cache that can contain a maximum of 2461 index descriptors. To create this cache, set the number of open indexes configuration parameter as follows:

```
sp_configure "number of open indexes", 2461
```

Using *sp_helpconfig* with *sybdiagdb* (Sybase Technical Support only)

Note Sybase Technical Support may create the *sybdiagdb* database on your system for debugging purposes. This database holds diagnostic configuration data, and is for use by Sybase Technical Support only.

The following *configname* options have been added to sp_helpconfig for Sybase Technical Support to use with the sybdiagdb database:

- *number of ccbs* – the number of configurable action point control blocks available to aid debugging.
- *caps per ccb* – the maximum number of configurable action points that can be configured at any one time within one configurable action point.
- *average cap size* – the estimated number of bytes of memory required to store the information associated with a typical configurable action point.

For example:

```

                sp_helpconfig "number of ccbs"
Minimum Value Maximum Value Default Value Current Value Memory Used
-----
0             100           0           0           0
                sp_helpconfig "caps per ccb"
Minimum Value Maximum Value Default Value Current Value Memory Used
-----
5             500           50          50          0
                sp_helpconfig "average cap size"
Minimum Value Maximum Value Default Value Current Value Memory Used
-----
100           10000          200         200         0

```

Permissions The options specified in “Using sp_helpconfig with sybdiagdb (Sybase Technical Support only)” on page 243 can be used only by Sybase Technical Support. Any user can execute sp_helpconfig with other *configname* options.

See also *System procedures* – sp_configure, sp_countmetadata, sp_monitorconfig

sp_helpconstraint

Description Reports information about integrity constraints used in the specified tables.

Syntax sp_helpconstraint [*objname*] [, detail]

Parameters *objname*
 – is the name of a table that has one or more integrity constraints defined by a create table or alter table statement.

detail
 – returns information about the constraint’s user or error messages.

Examples **Example 1**

```

        sp_helpconstraint store_employees

name                                     defn
-----
store_empl_stor_i_272004000             store_employees FOREIGN KEY
                                         (stor_id) REFERENCES stores(stor_id)
store_empl_mgr_id_288004057             store_employees FOREIGN KEY
                                         (mgr_id) SELF REFERENCES
                                         store_employees(emp_id)
store_empl_2560039432                   UNIQUE INDEX( emp_id) :
                                         NONCLUSTERED, FOREIGN REFERENCE
    
```

(3 rows affected)

Total Number of Referential Constraints: 2

Details:

- Number of references made by this table: 2
- Number of references to this table: 1
- Number of self references to this table: 1

Formula for Calculation:

Total Number of Referential Constraints
 = Number of references made by this table
 + Number of references made to this table
 - Number of self references within this table

Displays the constraint information for the store_employees table in the pubs3 database. The store_employees table has a foreign key to the stores table (stor_id) and a self-reference (mgr_id references emp_id).

Example 2

```

        sp_helpconstraint titles, detail

name                                     type
defn
msg
-----
-----
    
```

```
-----
datedflt                default value
      create default datedflt as getdate()

typedflt                default value
      create default typedflt as "UNDECIDED"

titles_pub_id_96003373  referential constraint
      titles FOREIGN KEY (pub_id) REFERENCES publishers(pub_id)
      standard system error message number : 547

roysched_title__144003544 referential constraint
      roysched FOREIGN KEY (title_id) REFERENCES titles(title_id)
      standard system error message number : 547

salesdetai_title__368004342 referential constraint
      salesdetail FOREIGN KEY (title_id) REFERENCES titles(title_id)
      standard system error message number : 547

titleautho_title__432004570 referential constraint
      titleauthor FOREIGN KEY (title_id) REFERENCES titles(title_id)
      standard system error message number : 547

titles_800033162        unique constraint
      UNIQUE INDEX ( title_id) : NONCLUSTERED, FOREIGN REFERENCE
      standard system error message number : 2601
```

(7 rows affected)

Total Number of Referential Constraints: 4

Details:

```
-- Number of references made by this table: 1
-- Number of references to this table: 3
-- Number of self references to this table: 0
```

Formula for Calculation:

```
Total Number of Referential Constraints
= Number of references made by this table
+ Number of references made to this table
- Number of self references within this table.
```

Displays more detailed information about the pubs3..salesdetail constraints, including the constraint type and any constraint error messages.

Example 3

```

                                sp_helpconstraint
id          name                  Num_referential_constraints
-----
80003316 titles                      4
16003088 authors                    3
176003658 stores                    3
256003943 salesdetail               3
208003772 sales                     2
336004228 titleauthor              2
896006223 store_employees           2
48003202 publishers                 1
128003487 roysched                 1
400004456 discounts                 1
448004627 au_pix                   1
496004798 blurbs                   1

```

(11 rows affected)

Displays a listing of all tables in the pubs3 database.

Usage

- `sp_helpconstraint` prints the name and definition of the integrity constraint, and the number of references used by the table. The `detail` option returns information about the constraint's user or error messages.
- Running `sp_helpconstraint` with no parameters lists all the tables containing references in the current database, and displays the total number of references in each table. `sp_helpconstraint` lists the tables in descending order, based on the number of references in each table.
- `sp_helpconstraint` reports only the integrity constraint information about a table (defined by a `create table` or `alter table` statement). It does not report information about rules, triggers, or indexes created using the `create index` statement. Use `sp_help` to see information about rules, triggers, and indexes for a table.
- For constraints that do not have user-defined messages, Adaptive Server reports the system error message associated with the constraint. Query `sysmessages` to obtain the actual text of that error message.
- You can use `sp_helpconstraint` only for tables in the current database.

- If a query exceeds the configured number of auxiliary scan descriptors, Adaptive Server returns an error message. You can use sp_helpconstraint to determine the necessary number of scan descriptors. For more information, see the description of the number of aux scan descriptors configuration parameter in the *System Administration Guide*.
- A System Security Officer can prevent the source text of constraint definitions from being displayed to most users who execute sp_helpconstraint. To restrict select permission on the text column of the syscomments table to the object owner or a System Administrator, use sp_configure to set the select on syscomments.text column parameter to 0. This restriction is required to run Adaptive Server in the evaluated configuration. For more information about the evaluated configuration, see the *System Administration Guide*.

Permissions Any user can execute sp_helpconstraint.

See also *Commands* – alter table, create table

System procedures – sp_help, sp_helpdb, sp_monitorconfig

sp_helpdb

Description Reports information about a particular database or about all databases.

Syntax sp_helpdb [dbname]

Parameters *dbname*
 – is the name of the database on which to report information. Without this optional parameter, sp_helpdb reports on all databases. *dbname* can include wildcard characters to return all databases that match the specified pattern.

Examples **Example 1**

sp_helpdb

name	db_size	owner	dbid	created	status
master	5.0 MB	sa	1	Jan 01, 1900	no options set
model	2.0 MB	sa	3	Jan 01, 1900	no options set
pubs2	2.0 MB	sa	6	Sep 20, 1995	no options set
sybssystemprocs	16.0 MB	sa	4	Sep 20, 1995	trunc log on chkp
tempdb	2.0 MB	sa	2	Sep 20, 1995	select into/bulkcopy

Displays information about all the databases in Adaptive Server.

Example 2

```
sp_helpdb pubs2
```

(Issued from within pubs2.)

```

name      db_size  owner  dbid  created      status
-----
pubs2    2.0 MB  sa     4     Mar 05, 1993  abort tran when log full
device_fragments  size  usage      free kbytes
-----
master                2.0 MB  data and log          576
device                segment
-----
master                default
master                logsegment
master                system
name      attribute_class  attribute      int_value  char_value  comments
-----
pubs2    buffer manager  cache binding          1  pubs2_cache  NULL

```

Displays information about the pubs2 database, and includes segment information.

Example 3

```
sp_helpdb pubs2
```

(Not issued from within pubs2.)

```

name      db_size  owner  dbid  created      status
-----
pubs2    2.0 MB  sa     4     Mar 05, 1993  abort tran when log full
device_fragments  size  usage      free kbytes
-----
master                2.0 MB  data and log          576
name      attribute_class  attribute      int_value  char_value  comments
-----
pubs2    buffer manager  cache binding          1  pubs2_cache  NULL

```

Displays information about the pubs2 database.

Example 4

```
sp_helpdb pubtune
```

```

name                                attribute_class
attribute                            int_value
      char_value
      comments
-----
pubtune                              lock strategy
      row lock promotion              NULL
      PCT = 95, LWM = 300, HWM = 300

```

Displays the row lock promotion attributes set for the pubtune database.

Usage

- sp_helpdb reports on the specified database when *dbname* is given. If no value is supplied for *dbname*, sp_helpdb reports on all the databases listed in master.dbo.sysdatabases.
- *dbname* can include wildcard characters to return all databases that match the specified pattern. For details about using wildcard characters, see Chapter 7, “Expressions, Identifiers, and Wildcard Characters,” in *Reference Manual Volume 1, Building Blocks*.
- Executing sp_helpdb *dbname* from *dbname* includes free space and segment information in the report.
- sp_helpdb displays information about a database’s attributes, giving the attribute’s class, name, integer value, character value, and comments, if any attributes are defined. Example 3 shows cache binding attributes for the pubs2 database.
- sp_helpdb reports if a database is offline.
- sp_helpdb reports row lock promotion thresholds, if any are defined for the database.
- A database created with the for load option has a status of “don’t recover” in the output from sp_helpdb.
- When Component Integration Services is enabled, sp_helpdb lists the default storage location for the specified database or all databases. If there is no default storage location, the display indicates “NULL”.

Permissions

Any user can execute sp_helpdb.

See also

Commands – alter database, create database
System procedures – sp_configure, sp_dboption, sp_rename

sp_helpdevice

Description Reports information about a particular device or about all Adaptive Server database devices and dump devices.

Syntax `sp_helpdevice [devname]`

Parameters *devname*
– is the name of the device about which to report information. If you omit this parameter, *sp_helpdevice* reports on all devices.

Examples

Example 1

```
sp_helpdevice
```

device_name	physical_name	description		
diskdump	null	disk, dump device		
master	d_master	special, default disk, dsync on,	physical	disk, 10 MB

status	cntrltype	device_number	low	high
16	2	0	0	20000
3	0	0	0	5120

Displays information about all the devices on Adaptive Server.

Example 2

```
sp_helpdevice diskdump
```

Reports information about the dump device named *diskdump*.

Usage

- *sp_helpdevice* displays information on the specified device, when *devname* is given, or on all devices in *master.dbo.sysdevices*, when no argument is given.
- The *sysdevices* table contains dump devices and database devices.
Database devices can be designated as default devices, which means that they can be used for database storage. This can occur when a user issues *create database* or *alter database* and does not specify a database device name or gives the keyword *default*. To make a database device a default database device, execute the system procedure *sp_diskdefault*.
- Add database devices to the system with *disk init*. Add dump devices with *sp_addumpdevice*.

- The number in the “status” column corresponds to the status description in the “description” column.

The “cntrltype” column specifies the controller number of the device. The “cntrltype” is 2 for disk or file dump devices and 3–8 for tape dump devices. For database devices, the “cntrltype” is usually 0 (unless your installation has a special type of disk controller).

The “device_number” column is 0 for dump devices, 0 for the master database device, and between 1 and 255 for other database devices. sp_helpdevice may report erroneous negative numbers for device numbers greater than 126.

The “low” and “high” columns represent virtual page numbers, each of which is unique among all the devices in Adaptive Server.

Permissions

Any user can execute sp_helpdevice.

See also

Commands – disk init, dump database, dump transaction, load database, load transaction

System procedures – sp_addumpdevice, sp_deviceattr, sp_dropdevice, sp_logdevice

sp_helpextendedproc

Description

Displays extended stored procedures (ESPs) in the current database, along with their associated DLL files.

Syntax

sp_helpextendedproc [*esp_name*]

Parameters

esp_name

– is the name of the extended stored procedure. It must be a procedure in the current database.

Examples

Example 1

```
use sybssystemprocs
go
sp_helpextendedproc xp_cmdshell

ESP Name      DLL Name
-----
xp_cmdshell  sybsyesp
```

Lists the *xp_cmdshell* ESP and the name of the DLL file in which its function is stored.

Example 2

```
sp_helpextendedproc
ESP Name      DLL Name
-----
xp_freedl     sybsyesp
xp_cmdshell   sybsyesp
```

Lists all the ESPs in the current database, along with the names of the DLL files in which their functions are stored.

Usage

- If the *esp_name* is omitted, *sp_helpextendedproc* lists all the extended stored procedures in the database.
- The *esp_name* is case sensitive. It must match the *esp_name* used to create the ESP.

Permissions

Only a System Administrator can execute *sp_helpextendedproc* to see all the ESPs in the database. All users can execute *sp_helpextendedproc* to see ESPs owned by themselves or by the Database Owner.

See also

Commands – create procedure, drop procedure

System procedures – *sp_addextendedproc*, *sp_dropextendedproc*

sp_helpexternlogin

Description

Component Integration Services only – Reports information about external login names.

Syntax

```
sp_helpexternlogin [remote_server [, login_name]]
```

Parameters

remote_server

– is the name of the remote server that has been added to the local server with *sp_addserver*.

login_name

– is a login account on the local server.

Examples

Example 1

```
sp_helpexternlogin
```

Displays all remote servers, local login names, and external logins.

Example 2

```
sp_helpexternlogin SSB
```

Displays local login names and external logins for the server named SSB.

Example 3

```
sp_helpexternlogin NULL, milo
```

Displays remote servers, local login names and external logins for the user named “milo”.

Example 4

```
sp_helpexternlogin SSB, trixi
```

Displays external logins for remote server SSB where the local user name is “trixi”.

Usage

- sp_helpexternlogin displays all remote servers, the user’s local login name, and the user’s external login name.
- Add remote servers with sp_addserver. Add local logins with sp_addlogin.

Permissions

Any user can execute sp_helpexternlogin.

See also

System procedures – sp_addexternlogin, sp_addlogin, sp_addserver, sp_helpserver

sp_helpgroup

Description

Reports information about a particular group or about all groups in the current database.

Syntax

```
sp_helpgroup [grpname]
```

Parameters

grpname

– is the name of a group in the database created with sp_addgroup.

Examples

Example 1

```
sp_helpgroup

Group_name          Group_id
-----
hackers             16384
public              0
```

Displays information about all groups in the current database.

Example 2

```

                                sp_helpgroup hackers
Group_name      Group_id      Users_in_group      Userid
-----
hackers        16384          ann                 4
hackers        16384          judy                 3

```

Displays information about the group “hackers”.

Usage

- To get a report on the default group, “public,” enclose the name “public” in single or double quotes (“public” is a reserved word).
- If there are no members in the specified group, *sp_helpgroup* displays the header, but lists no users, as follows:

```

Group_name      Group_id      Users_in_group      Userid
-----

```

Permissions

Any user can execute *sp_helpgroup*.

See also

Commands – grant, revoke

System procedures – *sp_addgroup*, *sp_changegroup*, *sp_dropgroup*, *sp_helprotect*, *sp_helpuser*

sp_helpindex

Description

Reports information about the indexes created on a table.

Syntax

sp_helpindex objname

Parameters

objname

– is the name of a table in the current database.

Examples

Example 1

```

                                sp_helpindex sysobjects
index_name      index_description
index_keys
index_max_rows_per_page index_fillfactor index_reservepagegap
-----
sysobjects      clustered, unique located on system
id

```

```

                                0          0          0
ncsysobjects                    nonclustered, unique located on system
  name,uid
                                0          0          0

```

Displays the types of indexes on the sysobjects table.

Example 2

```
sp_helpindex titles
```

```

index_name      index_description
  index_keys
  index_max_rows_per_page  index_fillfactor  index_reservepagegap
-----
title_id_ix     nonclustered, unique located on default
  title_id
                                0          0          0
publ_ix         nonclustered located on default
  pub_id, pubdate DESC
                                0          0          8
title_ix        clustered, allow duplicate rows located on default
  title
                                0          90          0

```

The index on publ_ix was created with pub_id in ascending order and pubdate in descending order.

Usage

- sp_helpindex lists any indexes on a table, including indexes created by defining unique or primary key constraints defined by a create table or alter table statement.
- sp_helpindex displays any attributes (for example, cache bindings) assigned to the indexes on a table.
- sp_helpindex displays:
 - The max_rows_per_page setting of the indexes.
 - Information about clustered indexes on data-only locked tables
 - The index ID (indid) of a clustered index in data-only locked tables is not equal to 1.
 - The column order of the keys, to indicate whether they are in ascending or descending order.
 - Space manage property values

- The key column name followed by the order. Only descending order is displayed. For example, if there is an index on column a ASC, b DESC, c ASC, “*index_keys*” shows “a, b DESC, c”.

Permissions

Any user can execute *sp_helpindex*.

See also

Commands – create index, drop index, update statistics

System procedures – *sp_help*, *sp_helpkey*

Procedures: *sp_helpjava* – *sp_helpprotect*

sp_helpjava

Description	<p>Displays information about Java classes and associated JARs that are installed in the database.</p> <p>For more information about Java in the database, see <i>Java in Adaptive Server Enterprise</i>.</p>
Syntax	<pre>sp_helpjava ["class" [, <i>java_class_name</i> [, "detail" "depends"]] "jar" [, <i>jar_name</i> [, "depends"]]]</pre>
Parameters	<p>"class" "jar"</p> <ul style="list-style-type: none"> – specifies whether to display information about a class or a JAR. Both “class” and “jar” are keywords, so the quotes are required. <p><i>java_class_name</i></p> <ul style="list-style-type: none"> – the name of the class about which you want information. The class must be a system class or a user-defined class that is installed in the database. <p>detail</p> <ul style="list-style-type: none"> – specifies that you want to see detailed information about the class. <p>depends</p> <ul style="list-style-type: none"> lists all the database objects that depend on the specified class or classes in the JAR, including SQLJ functions, SQLJ stored procedures, views, Transact-SQL stored procedures, and tables. <p><i>jar_name</i></p> <ul style="list-style-type: none"> – the name of the JAR for which you want to see information. The JAR must be installed in the database using <i>installjava</i>.
Examples	<p>Example 1</p> <pre>sp_helpjava</pre> <p>Displays the names of all classes and associated JAR files installed in the database.</p>

Example 2

```
sp_helpjava "class"
```

Displays the name of all classes.

Example 3

```
sp_helpjava "class", Address, detail
```

Displays detailed information about the Address class. For example:

```
Class
-----
Address

(1 row affected)
Class Modifiers
-----
public synchronized

Implemented Interfaces
-----
java.io.Serializable

Extended Superclass
-----
java.lang.Object

Constructors
-----
public Address()
public Address(java.lang.String, java.lang.String)

Methods
-----
public final native java.lang.Class
java.lang.Object.getClass()
public native int java.lang.Object.hashCode()
public boolean
java.lang.Object.equals(java.lang.Object)
public java.lang.String
java.lang.Object.toString()
public final native void java.lang.Object.notify()
public final native void
java.lang.Object.notifyAll()
public final native void
java.lang.Object.wait(long) throws
java.lang.InterruptedException
```

```
public final void java.lang.Object.wait(long,int)
throws java.lang.InterruptedException
public final void java.lang.Object.wait() throws
java.lang.InterruptedException
public java.lang.StringAddress.display()
public void Address.removeLeadingBlanks()
```

Fields

```
-----
public java.lang.String Address.street
public java.lang.String Address.zip
```

- Usage
- The depends parameter lists dependencies of a class or classes if the class is listed in the external name clause of a create statement for a SQLJ routine or is used as a datatype of a column in the database.
 - For more information about Java in the database, see *Java in Adaptive Server Enterprise*.
- Permissions
- Any user can execute *sp_helpjava*.
- See also
- Commands* – remove java
- Utilities* – extractjava, installjava

sp_helpjoins

- Description
- Lists the columns in two tables or views that are likely join candidates.
- Syntax
- sp_helpjoins lefttab, righttab*
- Parameters
- lefttab*
- is the first table or view.
- righttab*
- is the second table or view. The order of the parameters does not matter.

Examples

Example 1

```
sp_helpjoins sales, salesdetail
a1      a2      b1      b2      c1      c2
  d1      d2      e1      e2      f1      f2
    g1      g2      h1      h2
-----
-----
```

```

-----
stor_id stor_id ord_num ord_num NULL NULL
NULL NULL NULL NULL NULL NULL
NULL NULL NULL NULL

```

Displays a list of columns that are likely join candidates in the sales and salesdetail tables.

Example 2

```

                sp_helpjoins sysobjects, syscolumns
a1  a2  b1  b2  c1  c2  d1  d2  e1  e2
      f1  f2  g1  g2  h1  h2
-----
id  id  NULL NULL NULL NULL NULL NULL NULL NULL
    NULL NULL NULL NULL NULL NULL

```

Displays a list of columns that are likely join candidates in the sysobjects and syscolumns system tables.

Usage

- The column pairs that sp_helpjoins displays come from either of two sources. sp_helpjoins checks the syskeys table in the current database to see if any foreign keys have been defined with sp_foreignkey on the two tables, then checks to see if any common keys have been defined with sp_commonkey on the two tables. If sp_helpjoins does not find any foreign keys or common keys there, it checks for keys with the same user-defined datatypes. If that fails, it checks for columns with the same name and datatype.
- sp_helpjoins does not create any joins.

Permissions

Any user can execute sp_helpjoins.

See also

System procedures – sp_commonkey, sp_foreignkey, sp_helpkey, sp_primarykey

sp_helpkey

Description

Reports information about a primary, foreign, or common key of a particular table or view, or about all keys in the current database.

Syntax

sp_helpkey [tablename]

Parameters *tablename*
 – is the name of a table or view in the current database. If you do not specify a name, the procedure reports on all keys defined in the current database.

Examples `sp_helpkey`

keytype	object	related_object	object_keys	related_keys
primary	authors	-- none --	au_id,*,*,*,*,*,*,*	*,*,*,*,*,*,*,*
foreign	titleauthor	authors	au_id,*,*,*,*,*,*,*	au_id,*,*,*,*,*,*,*,* *,*

Displays information about the keys defined in the current database. The “object_keys” and “related_keys” columns refer to the names of the columns that make up the key.

- Usage**
- *sp_helpkey* lists information about all primary, foreign, and common key definitions that reference the table *tablename* or, if *tablename* is omitted, about all the keys in the database. Define these keys with the *sp_primarykey*, *sp_foreignkey*, and *sp_commonkey* system procedures.
 - *sp_helpkey* does not provide information about the unique or primary key integrity constraints defined by a create table statement. Use *sp_helpconstraint* to determine what constraints are defined for a table.
 - Create keys to make explicit a logical relationship that is implicit in your database design so that applications can use the information.
 - If you specify an object name, *sp_helpkey* follows the Adaptive Server rules for finding objects:
 - If you do not specify an owner name, and you own an object with the specified name, *sp_helpkey* reports on that object.
 - If you do not specify an owner name, and you do not own an object of that name, but the Database Owner does, *sp_helpkey* reports on the Database Owner’s object.
 - If neither you nor the Database Owner owns an object with the specified name, *sp_helpkey* reports an error condition, even if an object with that name exists in the database for a different owner.
 - If both you and the Database Owner own objects with the specified name, and you want to access the Database Owner’s object, specify the name in the form *dbo.objectname*.

- Qualify objects that are owned by database users other than yourself and the Database Owner with the owner’s name, as in “mary.myproc”.

Permissions Any user can execute sp_helpkey.

See also *Commands* – create trigger

System procedures – sp_commonkey, sp_foreignkey, sp_primarykey

sp_helplanguage

Description Reports information about a particular alternate language or about all languages.

Syntax sp_helplanguage [*language*]

Parameters *language*
– is the name of the alternate language you want information about.

Examples **Example 1**

```
sp_helplanguage french

langid dateformat datefirst upgrade    name
      alias
      months
      shortmonths
      days
-----
1      dmy          1          0          french
french
janvier , février , mars , avril , mai , juin , juillet , août , septembre ,
octobre , novembre , décembre
jan , fév , mar , avr , mai , jui , juil , août , sep , oct , nov , déc
lundi , mardi , mercredi , jeudi , vendredi , samedi , dimanche
```

Displays information about the alternate language, “french”.

Example 2

```
sp_helplanguage
```


	Displays information about all installed alternate languages.
Usage	<ul style="list-style-type: none">• <i>sp_helplanguage</i> reports on a specified language, when the language is given, or on all languages in <code>master.dbo.syslanguages</code>, when no language is supplied.
Permissions	Any user can execute <i>sp_helplanguage</i> .
See also	<i>System procedures</i> – <i>sp_addlanguage</i> , <i>sp_droplanguage</i> , <i>sp_setlangalias</i>

sp_helplog

Description	Reports the name of the device that contains the first page of the transaction log.
Syntax	<code>sp_helplog</code>
Parameters	None.
Examples	<pre>sp_helplog In database 'master', the log starts on device 'master'.</pre>
Usage	<ul style="list-style-type: none">• <i>sp_helplog</i> displays the name of the device that contains the first page of the transaction log in the current database.
Permissions	Any user can execute <i>sp_helplog</i> .
See also	<i>Commands</i> – <i>alter database</i> , <i>create database</i> <i>System procedures</i> – <i>sp_helpdevice</i> , <i>sp_logdevice</i>

sp_helpobjectdef

Description	<i>Component Integration Services only</i> – Reports owners, objects, and type information for remote object definitions.
Syntax	<code>sp_helpobjectdef [object_name]</code>

Parameters	<p><i>object_name</i></p> <p>– is the name of the object as it is defined in the sysattributes table. The <i>object_name</i> can be in any of the following forms:</p> <ul style="list-style-type: none">• <i>dbname.owner.object</i>• <i>dbname..object</i>• <i>owner.object</i>• <i>object</i> <p><i>dbname</i> and <i>owner</i> are optional. <i>object</i> is required. If <i>owner</i> is not supplied, the <i>owner</i> defaults to the current user name. If <i>dbname</i> is supplied, it must be the current database, and <i>owner</i> must be supplied or marked with the placeholder <i>dbname..object</i>. Enclose a multipart <i>object_name</i> in quotes.</p>
Examples	<p>Example 1</p> <pre>sp_helpobjectdef</pre> <p>Displays all remote object definitions in the current database.</p> <p>Example 2</p> <pre>sp_helpobjectdef "dbo.tb1"</pre> <p>Displays remote object definitions for the tb1 table owned by the Database Owner.</p>
Usage	<ul style="list-style-type: none">• If no <i>object_name</i> is supplied, sp_helpobjectdef displays all remote object definitions.• A server name is not permitted in the <i>object_name</i> parameter.
Permissions	Any user can execute sp_helpobjectdef.
See also	<p><i>Commands</i> – create table, create existing table, drop table</p> <p><i>System procedures</i> – sp_addobjectdef, sp_dropobjectdef, sp_helpserver</p>

sp_help_qpgroup

Description	Reports information on an abstract plan group.
Syntax	sp_help_qpgroup [<i>group</i> [, <i>mode</i>]]

Parameters

group

– is the name of an abstract plan group.

mode

– is the type of report to print, one of the following:

Mode	Information Returned
full	The number of rows and number of plans in the group, the number of plans that use two or more rows, the number of rows and plan IDs for the longest plans, and number of hash keys and hash key collision information. This is the default report mode.
stats	All of the information from the “full” report, except hash key information.
hash	The number of rows and number of abstract plans in the group, the number of hash keys, and hash-key collision information.
list	The number of rows and number of abstract plans in the group, and the following information for each query/plan pair: hash key, plan ID, first few characters of the query, and the first few characters of the plan.
queries	The number of rows and number of abstract plans in the group, and the following information for each query: hash key, plan ID, first few characters of the query.
plans	The number of rows and number of abstract plans in the group, and the following information for each plan: hash key, plan ID, first few characters of the plan.
counts	The number of rows and number of abstract plans in the group, and the following information for each plan: number of rows, number of characters, hash key, plan ID, first few characters of the query.

Examples

Example 1

```
sp_help_qpgroup
```

```

Group                GID        Plans
-----
ap_stdin              1           0
ap_stdout             2           0
dev_test              3          209

```

Reports summary information about all abstract plan groups in the database.

Example 2

```
sp_help_qpgroup test_plans
```

```
Query plans group 'test_plans', GID 8
```

```
Total Rows  Total QueryPlans
-----
                6                3
sysqueryplans rows consumption, number of query
plans per row count
```

```
Rows          Plans
-----
                2                3
```

```
Hashkeys
-----
                3
```

There is no hash key collision in this group.

Reports on the test_plans group.

Usage

- When used with an abstract plan group name, and no mode parameter, the default mode for sp_help_qpgroup is full.
- Hash-key collisions indicate that more than one plan for a particular user has the same hash-key value. When there are hash key collisions, the query text of each query with the matching hash key must be compared to the user’s query text in order to identify the matching query, so performance is slightly degraded.

Permissions

Any user can execute sp_help_qpgroup.

See also

System procedures – sp_help_qplan

sp_help_qplan

Description

Reports information about an abstract plan.

Syntax

sp_help_qplan *id* [, *mode*]

Parameters

id

– is the ID of the abstract plan.

mode

– is the type of report to print, one of the following:

mode	Information returned
full	The plan ID, group ID, and hash key, and the full query and plan text.

mode	Information returned
brief	The same as full, but only prints about 80 characters of the query and plan, rather than the full query and plan. This is the default mode.
list	The hash key, ID, and first 20 characters of the query and plan.

Examples

Example 1

```

sp_help_qplan 800005881

gid          hashkey      id
-----
          5  2054169974   937054374

query
-----
select type, avg(price) from titles group by type

plan
-----
( plan
  ( store Worktabl
    ( i_scan type_price titles )
  )
  ( t_scan ( ...

```

Prints the brief abstract plan report.

Example 2

```
sp_help_qplan 784005824, full
```

Prints the full abstract plan report.

Usage

- If you do not supply a value for the mode parameter, the default is brief.

Permissions

Any user can execute *sp_help_qplan* to see the abstract plan of a query that he or she owns. Only the System Administrator and the Database Owner can display an abstract plan owned by another user.

See also

System procedures – *sp_find_qplan*, *sp_help_qpgroup*

sp_helpremotelogin

Description	Reports information about a particular remote server's logins or about all remote server logins.
Syntax	sp_helpremotelogin [<i>remoteserver</i> [, <i>remotename</i>]]
Parameters	<i>remoteserver</i> – is the name of the server about which to report remote login information. <i>remotename</i> – is the name of a particular remote user on the remote server.
Examples	Example 1 <pre>sp_helpremotelogin GATEWAY</pre> <p>Displays information about all the remote users of the remote server GATEWAY.</p> Example 2 <pre>sp_helpremotelogin</pre> <p>Displays information about all the remote users of all the remote servers known to the local server.</p>
Usage	<ul style="list-style-type: none">• sp_helpremotelogin reports on the remote logins for the specified server, when <i>remoteserver</i> is given, or on all servers, when no parameter is supplied.
Permissions	Any user can execute sp_helpremotelogin.
See also	<i>System procedures</i> – sp_addremotelogin, sp_droptremotelogin, sp_helpserver

sp_help_resource_limit

Description	Reports on resource limits.
Syntax	sp_help_resource_limit [<i>name</i> [, <i>appname</i> [, <i>limittime</i> [, <i>limitday</i> [, <i>scope</i> [, <i>action</i>]]]]]]]]

Parameters

name

– is the Adaptive Server login to which the limits apply. For information about limits that govern a particular login, specify the login *name*. For information about limits without regard to login, specify null.

Note If you are not a System Administrator, specify your own login, or a login of NULL, to display information about the resource limits that apply to you.

appname

– is the name of the application to which the limit applies. For information about limits that govern a particular application, specify the application name that the client program passes to the Adaptive Server in the login packet. For information about limits without regard to application, specify null.

limittime

– is the time during which the limit is enforced. For information about limits in effect at a given time, specify the time, with a value between “00:00” and “23:59”, using the following form:

"*HH* : *MM*"

For information about limits without regard to time, specify null.

limitday

– is any day on which the limit is enforced. For information about resource limits in effect on a given day of the week, specify the full weekday name for the default server language, as stored in the *syslanguages* system table of the master database. For information about limits without regard to the days on which they are enforced, specify null.

scope

– is the scope of the limit. Specify one of the following:

Scope Code	For Help on All Limits That Govern
1	Queries
2	Query batches (one or more SQL statements sent by the client to the server)
4	Transactions
6	Both query batches and transactions
NULL	The specified <i>name</i> , <i>appname</i> , <i>limittime</i> , <i>limitday</i> , and <i>action</i> , without regard to their <i>scope</i>

action

– is the action to take when the limit is exceeded. Specify one of the following:

Action Code	For Help on All Limits That
1	Issue a warning
2	Abort the query batch
3	Abort the transaction
4	Kill the session
NULL	Govern the specified <i>name</i> , <i>appname</i> , <i>limittime</i> , <i>limitday</i> , and <i>scope</i> , without regard to the <i>action</i> they take

Examples

Example 1

```
sp_help_resource_limit
```

Lists all resource limits stored in the sysresourcelimits system table.

Example 2

```
sp_help_resource_limit joe_user
```

Lists all limits for the user “joe_user”.

Example 3

```
sp_help_resource_limit NULL, my_app
```

Lists all limits for the application *my_app*.

Example 4

```
sp_help_resource_limit NULL, NULL, "09:00"
```

Lists all limits enforced at 9:00 a.m.

Example 5

```
sp_help_resource_limit @limittype = "09:00"
```

An alternative way of listing the limits enforced at 9:00 a.m.

Example 6

```
sp_help_resource_limit NULL, NULL, NULL, Monday
```

Lists all limits enforced on Mondays.

Example 7

```
sp_help_resource_limit joe_user, NULL, "09:00",  
Monday
```


Usage	<p>Lists any limit in effect for “joe_user” on Mondays at 9:00 a.m.</p> <ul style="list-style-type: none">• <i>sp_help_resource_limit</i> reports on all resource limits, limits for a given login or application, limits in effect at a given time or day of the week, or limits with a given scope or action.• For more information on resource limits, see the System Administration Guide.
Permissions	<p>Any user can execute <i>sp_help_resource_limit</i> to list his or her own resource limits. Only a System Administrator can execute <i>sp_help_resource_limit</i> to list limits that apply to other users.</p>
See also	<p><i>System procedures</i> – <i>sp_add_resource_limit</i>, <i>sp_drop_resource_limit</i>, <i>sp_modify_resource_limit</i></p>

sp_helpprotect

Description	<p>Reports on permissions for database objects, users, groups, or roles.</p>
Syntax	<pre>sp_helpprotect [<i>name</i> [, <i>username</i> [, "grant" [,"none" "granted" "enabled" "role_name"]]]]</pre>
Parameters	<p><i>name</i> is either the name of the table, view, stored procedure, SQLJ stored procedure, SQLJ function, or the name of a user, user-defined role, or group in the current database. If you do not provide a name, <i>sp_helpprotect</i> reports on all permissions in the database.</p> <p><i>username</i> – is a user’s name in the current database.</p> <p>grant – displays the privileges granted to <i>name</i> with grant option.</p> <p>none – ignores roles granted to the user when determining permissions granted.</p> <p>granted – includes information on all roles granted to the user when determining permissions granted.</p> <p>enabled – includes information on all roles activated by the user when determining permissions granted.</p>

role_name

– displays permission information for the specified role only, regardless of whether this role has been granted to the user.

Examples

Example 1

```
grant select on titles to judy
grant update on titles to judy
revoke update on titles(price) from judy
grant select on publishers to judy
with grant option
```

After this series of grant and revoke statements, executing sp_helpprotect titles results in this display:

grantor	grantee	type	action	object	column	grantable
dbo	judy	Grant	Select	titles	All	FALSE
dbo	judy	Grant	Update	titles	advance	FALSE
dbo	judy	Grant	Update	titles	notes	FALSE
dbo	judy	Grant	Update	titles	pub_id	FALSE
dbo	judy	Grant	Update	titles	pubdate	FALSE
dbo	judy	Grant	Update	titles	title	FALSE
dbo	judy	Grant	Update	titles	title_id	FALSE
dbo	judy	Grant	Update	titles	total_sales	FALSE
dbo	judy	Grant	Update	titles	type	FALSE
dbo	judy	Grant	Select	publishers	all	TRUE

Example 2

```
grant select, update on titles(price, advance)
to mary
with grant option
go
sp_helpprotect titles
```

After issuing this grant statement, sp_helpprotect displays the following:

grantor	grantee	type	action	object	column	grantable
dbo	mary	Grant	Select	titles	advance	TRUE
dbo	mary	Grant	Select	titles	price	TRUE
dbo	mary	Grant	Update	titles	advance	TRUE
dbo	mary	Grant	Update	titles	price	TRUE

Example 3

```
sp_helpprotect judy
```

Displays all the permissions that “judy” has in the database.

Example 4

```
sp_helpprotect sysusers, csmith, null, doctor,
"grant"
```

Displays any permissions that “csmith” has on the `sysusers` table, as well as whether “csmith” has with grant option which allows “csmith” to grant permissions to other users.

grantor	grantee	type	action	object	column	grantable
dbo	doctor	Grant	Delete	sysusers	All	FALSE
dbo	doctor	Grant	Insert	sysusers	All	FALSE
dbo	doctor	Grant	References	sysusers	All	FALSE
dbo	doctor	Grant	Select	sysattributes	All	FALSE

```
(1 row affected)
(return status = 0)
```

Example 5

```
sp_helpprotect doctor_role
```

Displays information about the permissions that the doctor role has in the database.

grantor	grantee	type	action	object	column	grantable
dbo	doctor	Grant	Delete	sysusers	All	FALSE
dbo	doctor	Grant	Insert	sysusers	All	FALSE
dbo	doctor	Grant	References	sysusers	All	FALSE
dbo	doctor	Grant	Select	sysattributes	All	FALSE

```
(1 row affected)
(return status = 0)
```

Example 6

```
sp_helpprotect sysusers, csmith, null, doctor_role,
"granted"
```

Displays information on all roles granted to “csmith”.

grantor	grantee	type	action	object	column	grantable
dbo	csmith	Grant	Update	sysusers	All	FALSE
dbo	doctor	Grant	Delete	sysusers	All	FALSE
dbo	doctor	Grant	Insert	sysusers	All	FALSE
dbo	doctor	Grant	References	sysusers	All	FALSE

(1 row affected)
 (return status = 0)

Example 7

```
sp_helpprotect sysattributes, rpillai, null, intern,
"enabled"
```

Displays information on all active roles granted to “rpillai”.

grantor	grantee	type	action	object	column	grantable
dbo	public	Grant	Select	sysattributes	All	FALSE

(1 row affected)
 (return status = 0)

Example 8

```
sp_helpprotect function_sqlj
```

Advises that SQLJ function access is public.

Implicit grant to public for SQLJ functions.

Usage

- sp_helpprotect reports permissions on a database object. If you supply the *username* parameter, only that user’s permissions on the database object are reported. If *name* is not an object, sp_helpprotect checks to see if it is a user, a group, or a role. If it is, sp_helpprotect lists the permissions for the user, group, or role.
- sp_helpprotect looks for objects and users in the current database only.
- If you do not specify an optional value such as granted, enabled, none, or *role_name*, Adaptive Server returns information on all roles activated by the current specified user.
- If the specified user is not the current user, Adaptive Server returns information on all roles granted to the specified user.
- Displayed information always includes permissions granted to the group in which the specified user is a member.

- In granting permissions, a System Administrator is treated as the object owner. If a System Administrator grants permission on another user's object, the owner's name appears as the grantor in *sp_helpprotect* output.

Permissions

Any user can execute *sp_helpprotect* to view his or her own permissions. Only a System Security Officer can execute *sp_helpprotect* to view permissions granted to other users.

See also

Commands – grant, revoke

System procedures – *sp_activeroles*, *sp_displayroles*

Procedures: *sp_helpsegment* – *sp_helpuser*

sp_helpsegment

Description	Reports information about a particular segment or about all segments in the current database.																					
Syntax	<i>sp_helpsegment</i> [<i>segname</i>]																					
Parameters	<p><i>segname</i></p> <p>– is the name of the segment about which you want information. If you omit this parameter, information about all segments in the current database appears.</p>																					
Examples	<p>Example 1</p> <pre> sp_helpsegment segment name status ----- 0 system 0 1 default 1 2 logsegment 0 </pre> <p>Reports information about all segments in the current database.</p> <p>Example 2</p> <pre> sp_helpsegment order_seg segment name status ----- 3 order_seg 0 </pre> <table> <thead> <tr> <th>device</th> <th>size</th> <th>free_pages</th> </tr> </thead> <tbody> <tr> <td>tpcd_data1</td> <td>25.0MB</td> <td>8176</td> </tr> <tr> <td>tpcd_data2</td> <td>25.0MB</td> <td>8512</td> </tr> <tr> <td>tpcd_data3</td> <td>25.0MB</td> <td>8392</td> </tr> <tr> <td>tpcd_data4</td> <td>25.0MB</td> <td>8272</td> </tr> <tr> <td>tpcd_data5</td> <td>25.0MB</td> <td>8448</td> </tr> <tr> <td>tpcd_data6</td> <td>25.0MB</td> <td>8512</td> </tr> </tbody> </table>	device	size	free_pages	tpcd_data1	25.0MB	8176	tpcd_data2	25.0MB	8512	tpcd_data3	25.0MB	8392	tpcd_data4	25.0MB	8272	tpcd_data5	25.0MB	8448	tpcd_data6	25.0MB	8512
device	size	free_pages																				
tpcd_data1	25.0MB	8176																				
tpcd_data2	25.0MB	8512																				
tpcd_data3	25.0MB	8392																				
tpcd_data4	25.0MB	8272																				
tpcd_data5	25.0MB	8448																				
tpcd_data6	25.0MB	8512																				

table_name	index_name	indid
orders	orders	0

total_size	total_pages	free_pages	used_pages
150.0MB	76800	50312	26488

Reports information about the segment named order_seg, including which database tables and indexes use that segment and the total number of pages, free pages and used pages on the segment.

Example 3

```
sp_helpsegment "default"
```

Reports information about the default segment. The keyword default must be enclosed in quotes.

Example 4

```
sp_helpsegment logsegment
```

segment name	status
2 logsegment	0

device	size	free_pages
tpcd_log1	20.0MB	10200

table_name	index_name	indid
syslogs	syslogs	0

total_size	total_pages	free_pages	used_pages
20.0MB	10240	10200	40

Reports information about the segment on which the transaction log is stored.

Usage

- sp_helpsegment displays information about the specified segment, when *segname* is given, or about all segments in the current database, when no argument is given.
- When you first create a database, Adaptive Server automatically creates the system, default, and logsegment segments. Use sp_addsegment to add segments to the current database.

- The system, default, and logsegment segments are numbered 0, 1, and 2, respectively.
- The “status” column indicates which segment is the default pool of space. Use *sp_placeobject* or the *on segment_name* clause of the *create table* or *create index* command to place objects on specific segments.
- The “indid” column is 0 if the table does not have a clustered index and is 1 if the table has a clustered index.

Permissions

Any user can execute *sp_helpsegment*.

See also

Commands – *create index*, *create table*

System procedures – *sp_addsegment*, *sp_dropsegment*, *sp_extendsegment*, *sp_helpdb*, *sp_helpdevice*, *sp_placeobject*

sp_helpserver

Description

Reports information about a particular remote server or about all remote servers.

Syntax

sp_helpserver [*server*]

Parameters

server

– is the name of the remote server about which you want information.

Examples

Example 1

```
sp_helpserver GATEWAY
```

Displays information about the remote server GATEWAY.

Example 2

```
sp_helpserver SYB_BACKUP
```

name	network_name	status	id
SYB_BACKUP	SYB_BACKUP	timeouts, no net password encryption	1

Displays information about the local Backup Server.

Example 3

```
sp_helpserver
```


On an 8-bit terminal, it appears as follows:

Sort Order Description

```
-----
Character Set = 1, iso_1
  ISO 8859-1 (Latin-1) - Western European 8-bit character set.
Sort Order = 50, bin_iso_1
  Binary sort order for the ISO 8859/1 character set (iso_1).
Characters, in Order
-----
! " # $ % & ` ( ) * + , - . / 0 1 2 3 4 5 6 7 8 9 : ; < = > ?
@ A B C D E F G H I J K L M N O P Q R S T U V W X Y Z [ \ ] ^ _
` a b c d e f g h i j k l m n o p q r s t u v w x y z { | } ~
ı ç £ ¤ ¥ ¦ § ¨ © ª « ¬ ® ¯ ° ± ² ³ ´ µ ¶ · ¸ ¹ º » ¼ ½ ¾ ¿ À
Á Â Ã Ä Å Æ Ç È É Ê Ë Ì Í Î Ï Ð Ñ Ò Ó Ô Õ Ö × Ø Ù Ú Û Ü Ý Þ à
á â ã ä å æ ç è é ê ë ì í î ï ñ ò ó ô õ ö ÷ ø ù ú û ü ý þ ÿ
```

For a Class 2 (multibyte) character set, the characters are not listed, but a description of the character set is included. For example:

Sort Order Description

```
-----
Character Set = 140, euc_jis
  Japanese. Extended Unix Code mapping for JIS-X0201
  (hankaku katakana) and JIS-X0208 (double byte) roman,
  kana, and kanji.
Class 2 character set
Sort Order = 50, bin_eucjis
  Binary sort order for Japanese using the EUC JIS
  character set as a basis.
```

- Usage
- Binary sort order is the default.
- Permissions
- Any user can execute *sp_helpsort*.

sp_helptext

- Description
- Displays the **source text** of a **compiled object**.
- Syntax
- sp_helptext objname*
- Parameters
- objname*
- is the name of the compiled object for which the source text is to be displayed. The compiled object must be in the current database.

Examples

Example 1

```
sp_helptext pub_idrule

# Lines of Text
-----
1
text
-----
create rule pub_idrule
as @pub_id in ("1389", "0736", "0877",
              "1622", "1756")
   or @pub_id like "99[0-9][0-9]"
```

Displays the source text of pub_idrule. Since this rule is in the pubs2 database, execute this command from pubs2.

Example 2

```
sp_helptext sp_helptext
```

Displays the source text of sp_helptext. Since system procedures are stored in sysystemprocs, execute this command from sysystemprocs.

Usage

- sp_helptext prints out the number of rows in syscomments (255 characters long each) that are occupied by the compiled object, followed by the source text of the compiled object.
- sp_helptext looks for the source text in the syscomments table in the current database.
- Encrypt the source text with sp_hidetext.
- A System Security Officer can prevent the source text of compiled objects from being displayed to most users who execute sp_helptext. To restrict select permission on the text column of the syscomments table to the object owner or a System Administrator, use sp_configure to set the select on syscomments.text column parameter to 0. This restriction is required to run Adaptive Server in the evaluated configuration. For more information about the evaluated configuration, see the System Administration Guide.

Permissions

Any user can execute sp_helptext.

See also

System procedures – sp_checksourc, sp_hidetext

sp_helpthreshold

Description	Reports the segment, free-space value, status, and stored procedure associated with all thresholds in the current database or all thresholds for a particular segment.
Syntax	<code>sp_helpthreshold [segname]</code>
Parameters	<i>segname</i> – is the name of a segment in the current database.
Examples	<p>Example 1</p> <pre>sp_helpthreshold logsegment</pre> <p>Shows all thresholds on the log segment.</p> <p>Example 2</p> <pre>sp_helpthreshold</pre> <p>Shows all thresholds on all segments in the current database.</p> <p>Example 3</p> <pre>sp_helpthreshold "default"</pre> <p>Shows all thresholds on the default segment. Note the use of quotes around the reserved word “default”.</p>
Usage	<ul style="list-style-type: none"> • <code>sp_helpthreshold</code> displays threshold information for all segments in the current database. If you provide the name of a segment, <code>sp_helpthreshold</code> lists all thresholds in that segment. • The <code>status</code> column is 1 for the last-chance threshold and 0 for all other thresholds. Databases that do not store their transaction logs on a separate segment have no last-chance threshold.
Permissions	Any user can execute <code>sp_helpthreshold</code> .
See also	<i>System procedures</i> – <code>sp_addthreshold</code> , <code>sp_droptreshold</code> , <code>sp_helpsegment</code> , <code>sp_modifythreshold</code> , <code>sp_thresholdaction</code>

sp_helpuser

Description	Reports information about a particular user, group, or alias, or about all users, in the current database.
Syntax	<code>sp_helpuser [name_in_db]</code>

Parameters

name_in_db
– is the user’s name in the current database.

Examples

Example 1

```
sp_helpuser

Users_name  ID_in_db      Group_name    Login_name
-----
ann         4             hackers      ann
dbo         1             public       sa
guest       2             public       NULL
judy        3             hackers      judy
```

Displays information about all users in the current database.

Example 2

```
sp_helpuser dbo

Users_name      ID_in_db      Group_name    Login_name
-----
dbo             1             public       sa

Users aliased to user.
Login_name
-----
andy
christa
howard
linda
```

Displays information about the Database Owner (user name “dbo”).

Usage

- sp_helpuser reports information about all users of the current database. If you specify a *name_in_db*, sp_helpuser reports information on the specified user only.
- If the specified user is not listed in the current database’s sysusers table, sp_helpuser checks to see if the user is aliased to another user or is a group name.

Permissions

Any user can execute sp_helpuser.

See also

System procedures – sp_adduser, sp_dropuser, sp_helpgroup

sp_hidetext

Description	Hides the source text for the specified compiled object .
Syntax	<code>sp_hidetext [objname [, tabname [, username]]]</code>
Parameters	<i>objname</i> – specifies the compiled object for which to hide the source text. <i>tabname</i> – specifies the name of the table or view for which to hide the source text. <i>username</i> – specifies the name of the user who owns the compiled object for which to hide the source text.
Examples	Example 1 <pre>sp_hidetext</pre> Hides the source text of all compiled objects in the current database. Example 2 <pre>sp_hidetext @objname = "sp_sort_table", @username = "Mary"</pre> Hides the source text of the user-defined stored procedure, <code>sp_sort_table</code> , that is owned by Mary. Example 3 <pre>sp_hidetext "pr_phone_list"</pre> Hides the source text of the stored procedure <code>pr_phone_list</code> . Example 4 <pre>sp_hidetext @tabname = "my_tab"</pre> Hides the source text of all check constraints, defaults, and triggers defined on the table <code>my_tab</code> .

Example 5

```
sp_hidetext "my_vu", "my_tab"
```

Hides the source text of the view my_vu and all check constraints, defaults, and triggers defined on the table my_tab.

Example 6

```
sp_hidetext @username = "Tom"
```

Hides the source text of all compiled objects that are owned by Tom.

Usage

- sp_hidetext hides the **source text** for the specified **compiled object**.

Warning! Before executing sp_hidetext, make sure you have a backup of the source text. The results of executing sp_hidetext are not reversible.

- If you do not provide any parameters, sp_hidetext hides the source text for all compiled objects in the current database.
- For more information about hiding source text, see the Transact-SQL User's Guide.

Permissions

Any user can use sp_hidetext to hide the source text of his or her own compiled objects. Only a Database Owner or a System Administrator can hide the source text of compiled objects that are owned by another user or use sp_hidetext with no parameters.

See also

System procedures – sp_checksourc

Procedures: *sp_import_qpgroup* – *sp_logiosize*

sp_import_qpgroup

Description	Imports abstract plans from a user table into an abstract plan group.
Syntax	<code>sp_import_qpgroup <i>tab</i>, <i>usr</i>, <i>group</i></code>
Parameters	<p><i>tab</i></p> <p>– is the name of a table from which to copy the plans. You can specify a database name, but not an owner name, in the form <i>dbname..tablename</i>. The total length must be 30 characters or less.</p> <p><i>usr</i></p> <p>– is the name of the user whose ID should be assigned to the abstract plans when they are imported.</p> <p><i>group</i></p> <p>– is the name of the abstract plan group that contains the plans to be imported.</p>
Examples	<pre>sp_import_qpgroup moveplans, dbo, new_plans</pre> <p>Copies plans from the table <code>moveplans</code> to the <code>new_plans</code> group, giving them the user ID for the Database Owner.</p>
Usage	<ul style="list-style-type: none">• <code>sp_import_qpgroup</code> copies plans from a user table to an abstract plan group in <code>sysqueryplans</code>. With <code>sp_export_qpgroup</code>, it can be used to copy abstract plan groups between servers and databases, or to copy plans belonging to one user and assign them the ID of another user.• <code>sp_import_qpgroup</code> creates the abstract plan group if it does not exist when the procedure is executed.• If an abstract plan group exists when <code>sp_import_qpgroup</code> is executed, it cannot contain any plans for the specified user. <code>sp_import_qpgroup</code> does not check the query text to determine whether queries already exist in the group. If you need to import plans for a user into a group where some plans for the user already exist:<ul style="list-style-type: none">• Use <code>sp_import_qpgroup</code> to import the plans into a new plan group.

- Use `sp_copy_all_qplans` to copy the plans from the newly-created group to the destination group. `sp_copy_all_qplans` does check queries to be sure that no duplicate plans are created.
- If you no longer need the group you created for the import, drop the plans in the group with `sp_drop_all_qplans`, then drop the group with `sp_drop_qpgroup`.
- To create an empty table in order to bulk copy abstract plans, use:

```
select * into load_table
from sysqueryplans
where 1 = 2
```

Permissions Only a System Administrator or the Database Owner can execute `sp_import_qpgroup`.

See also *Commands* – create plan

System procedures – `sp_copy_all_qplans`, `sp_copy_qplan`, `sp_export_qpgroup`, `sp_help_qpgroup`

sp_indsuspect

Description Checks user tables for indexes marked as suspect during recovery following a sort order change.

Syntax `sp_indsuspect [tab_name]`

Parameters *tab_name*
– is the name of the user table to be checked.

Examples `sp_indsuspect newacct`

Checks the table `newacct` for indexes marked as suspect.

- Usage
- `sp_indsuspect` with no parameter creates a list of all tables in the current database that have indexes that need to be rebuilt as a result of a sort order change. With a *tab_name* parameter, `sp_indsuspect` checks the specified table for indexes marked as suspect during recovery following a sort order change.
 - Use `sp_indsuspect` to list all suspect indexes. The table owner or a System Administrator can use `dbcc reindex` to check the integrity of the listed indexes and to rebuild them if necessary.

Permissions Any user can execute *sp_indsuspect*.
See also *Commands* – *dbcc*

sp_listsuspect_db

Description Lists all databases that currently have offline pages because of corruption detected on recovery.

Syntax *sp_listsuspect_db*

Parameters None.

Examples

```
sp_listsuspect_db
```


Lists the databases that have suspect pages.

Usage

- *sp_listsuspect_db* lists the database name, number of suspect pages, and number of objects containing suspect pages.
- Use *sp_listsuspect_page* to identify the suspect pages.

Permissions Any user can execute *sp_listsuspect_db*.

See also *System procedures* – *sp_listsuspect_page*, *sp_setsuspect_granularity*, *sp_setsuspect_threshold*

sp_listsuspect_object

Description Lists all indexes in a database that are currently offline because of corruption detected on recovery.

Syntax *sp_listsuspect_object* [*dbname*]

Parameters *dbname*
– is the name of the database.

Examples **Example 1**

```
sp_listsuspect_object
```


Lists the suspect indexes in the current database.

Example 2

```
sp_listsuspect_object pubs2
```

Lists the suspect indexes in the pubs2 database.

Usage

- If an index on a data-only-locked table has suspect pages, the entire index is taken offline during recovery. Offline indexes are not considered by the query optimizer.
- Use the system procedure sp_forceonline_object to bring an offline index online for repair.
- Indexes on allpages-locked tables are not taken completely offline during recovery; only individual pages of these indexes are taken offline. These pages can be brought online with sp_forceonline_page.
- sp_listsuspect_object lists the database name, object ID, object name, index ID, and access status for every suspect index in the specified database or, if *dbname* is omitted, in the current user database.
- A value of SA_ONLY in the access column means that the index has been forced online for System Administrator use only. A value of BLOCK_ALL means that the index is offline for everyone.
- For more information on recovery fault isolation, see the System Administration Guide.

Permissions

Any user can execute sp_listsuspect_object.

See also

System procedures – sp_forceonline_object

sp_listsuspect_page

Description

Lists all pages in a database that are currently offline because of corruption detected on recovery.

Syntax

```
sp_listsuspect_page [dbname]
```

Parameters

dbname
– is the name of the database.

Examples

Example 1

```
sp_listsuspect_page
```

Lists the suspect pages in the current database.

Example 2

```
sp_listsuspect_page pubs2
```

Lists the suspect pages in the pubs2 database.

Usage

- *sp_listsuspect_page* lists the database name, page ID, object, index ID, and access status for every suspect page in the specified database or, if *dbname* is omitted, in the current user database.
- A value of SA_ONLY in the “access” column indicates that the page has been forced online for System Administrator use only. A value of BLOCK_ALL indicates that the page is offline for everyone.

Permissions

Any user can execute *sp_listsuspect_page*.

See also

System procedures – *sp_listsuspect_db*, *sp_setsuspect_granularity*, *sp_setsuspect_threshold*

sp_lock

Description

Reports information about processes that currently hold locks.

Syntax

```
sp_lock [spid1 [, spid2]]
```

Parameters

spid1

– is the Adaptive Server process ID number from the master.dbo.sysprocesses table. Run *sp_who* to get the *spid* of the locking process.

spid2

– is another Adaptive Server process ID number to check for locks.

Examples

Example 1

sp_lock

The class column will display the cursor name for locks associated with a cursor for the current user and the cursor id for other users.

fid	spid	locktype	table_id	page	dbname	class	context
0	7	Sh_intent	480004741	0	master	Non Cursor Lock	NULL
0	18	Ex_intent	16003088	0	pubtune	Non Cursor Lock	NULL
0	18	Ex_page	16003088	587	pubtune	Non Cursor Lock	NULL
0	18	Ex_page	16003088	590	pubtune	Non Cursor Lock	NULL

0	18	Ex_page	16003088	1114	pubtune	Non	Cursor	Lock	NULL
0	18	Ex_page	16003088	1140	pubtune	Non	Cursor	Lock	NULL
0	18	Ex_page	16003088	1283	pubtune	Non	Cursor	Lock	NULL
0	18	Ex_page	16003088	1362	pubtune	Non	Cursor	Lock	NULL
0	18	Ex_page	16003088	1398	pubtune	Non	Cursor	Lock	NULL
0	18	Ex_page-blk	16003088	634	pubtune	Non	Cursor	Lock	NULL
0	18	Update_page	16003088	1114	pubtune	Non	Cursor	Lock	NULL
0	18	Update_page-blk	16003088	634	pubtune	Non	Cursor	Lock	NULL
0	23	Sh_intent	16003088	0	pubtune	Non	Cursor	Lock	NULL
0	23	Sh_intent	176003658	0	pubtune	Non	Cursor	Lock	NULL
0	23	Ex_intent	208003772	0	pubtune	Non	Cursor	Lock	NULL
1	1	Sh_intent	176003658	0	tpcd	Non	Cursor	Lock	Sync-pt
duration request									
1	1	Sh_intent-blk	208003772	0	tpcd	Non	Cursor	Lock	Sync-pt
duration request									
1	8	Sh_page	176003658	41571	tpcd	Non	Cursor	Lock	NULL
1	9	Sh_page	176003658	41571	tpcd	Non	Cursor	Lock	NULL
1	10	Sh_page	176003658	41571	tpcd	Non	Cursor	Lock	NULL
11	11	Sh_intent	176003658	0	tpcd	Non	Cursor	Lock	Sync-pt
duration request									
11	12	Sh_page	176003658	41571	tpcd	Non	Cursor	Lock	NULL
11	13	Sh_page	176003658	41571	tpcd	Non	Cursor	Lock	NULL
11	14	Sh_page	176003658	41571	tpcd	Non	Cursor	Lock	NULL

This example shows the lock status of serial processes with spids 7, 18, and 23 and two families of processes. The family with fid 1 has the coordinating processes with spid 1 and worker processes with spids 8, 9, and 10. The family with fid 11 has the coordinating processes with spid 11 and worker processes with spids 12, 13, and 14.

Example 2

sp_lock 7

The class column will display the cursor name for locks associated with a cursor for the current user and the cursor id for other users.

fid	spid	locktype	table_id	page	dbname	class	context		
----	----	-----	-----	-----	-----	-----	-----	-----	-----
0	7	Sh_intent	480004741	0	master	Non	Cursor	Lock	NULL

Displays information about the locks currently held by spid 7.

Usage

- sp_lock with no parameters reports information on all processes that currently hold locks.
- The only user control over locking is through the use of the holdlock keyword in the select statement.

- Use the `object_name` system function to derive a table's name from its ID number.
- `sp_lock` output is ordered by `fid` and then `spid`.
- The `loid` column identifies unique lock owner ID of the blocking transaction. Even `loid` values indicate that a local transaction owns the lock. Odd values indicate that an external transaction owns the lock.
- The `locktype` column indicates whether the lock is a shared lock (“Sh” prefix), an exclusive lock (“Ex” prefix) or an update lock, and whether the lock is held on a table (“table” or “intent”) or on a page (“page”).

A “blk” suffix in the “locktype” column indicates that this process is blocking another process that needs to acquire a lock. As soon as this process completes, the other process(es) moves forward. A “demand” suffix in the “locktype” column indicates that the process is attempting to acquire an exclusive lock. For more information about lock types, see the Performance and Tuning Guide.

- The `class` column indicates whether a lock is associated with a cursor. It displays one of the following:
 - “Non Cursor Lock” indicates that the lock is not associated with a cursor.
 - “Cursor Id *number*” indicates that the lock is associated with the cursor ID number for that Adaptive Server process ID.
 - A cursor name indicates that the lock is associated with the cursor *cursor_name* that is owned by the current user executing `sp_lock`.
- The `fid` column identifies the family (including the coordinating process and its worker processes) to which a lock belongs. Values for `fid` are:
 - A zero value indicates that the task represented by the `spid` is executed serially. It is not participating in parallel execution.
 - A nonzero value indicates that the task (`spid`) holding the lock is a member of a family of processes (identified by `fid`) executing a statement in parallel. If the value is equal to the `spid`, it indicates that the task is the coordinating process in a family executing a query in parallel.

- The context column identifies the context of the lock. Worker processes in the same family have the same context value. Legal values for “context” are as follows:
 - “NULL” means that the task holding this lock is either a query executing serially, or is a query executing in parallel in transaction isolation level 1.
 - “Sync-pt duration request” means that the task holding the lock will hold the lock until the query is complete.

A lock’s context may be “Sync-pt duration request” if the lock is a table lock held as part of a parallel query, if the lock is held by a worker process at transaction isolation level 3, or if the lock is held by a worker process in a parallel query and must be held for the duration of the transaction.

- “Ind pg” indicates locks on index pages (allpages-locked tables only)
- “Inf key” indicates an infinity key lock (for certain range queries at transaction isolation level 3 on data-only-locked tables)
- “Range” indicates a range lock (for range queries at transaction isolation level 3 on data-only-locked tables)

These new values may appear in combination with “Fam dur” (which replaces “Sync pt duration”) and with each other, as applicable.

- The row column displays the row number for row-level locks.
- sp_lock output also displays the following lock types:
 - “Sh_row” indicates shared row locks
 - “Update_row” indicates update row locks
 - “Ex_row” indicates exclusive row locks

Permissions

Any user can execute sp_lock.

See also

Commands – kill, select

System procedures – sp_familylock on page 215, sp_who

sp_locklogin

Description	Locks an Adaptive Server account so that the user cannot log in or displays a list of all locked accounts.
Syntax	<code>sp_locklogin [loginname, "{lock unlock}"]</code>
Parameters	<p><i>loginname</i></p> <ul style="list-style-type: none"> – is the name of the account to be locked or unlocked. <p>lock unlock</p> <ul style="list-style-type: none"> – specifies whether to lock or unlock the account.
Examples	<p>Example 1</p> <pre>sp_locklogin charles, "lock"</pre> <p>Locks the login account for the user “charles.”</p> <p>Example 2</p> <pre>sp_locklogin</pre> <p>Displays a list of all locked accounts.</p>
Usage	<ul style="list-style-type: none"> • Locking an Adaptive Server login account prevents that user from logging in. Use <code>sp_locklogin</code> instead of <code>sp_droplogin</code> for the following reasons: <ul style="list-style-type: none"> • You cannot drop a login who is a user in any database, and you cannot drop a user from a database if the user owns any objects in that database or has granted any permissions on objects to other users. • Adaptive Server may reuse the dropped login account’s server user ID (suid) when the next login account is created. This occurs only when the dropped login holds the highest suid in syslogins; however, it could compromise accountability if execution of <code>sp_droplogin</code> is not being audited. In addition, it is possible that the user with the reused suid will actually be able to access database objects that were authorized for the old suid. • You cannot drop the last remaining System Security Officer’s or System Administrator’s login account. • <code>sp_locklogin</code> with no parameters returns a list of all the locked accounts. • You can lock an account that is currently logged in. The user receives a warning that his or her account has been locked, but is not locked out of the account until he or she logs out.

- A locked account can be specified as a Database Owner and can own objects in any database.
- Locking an account that is already locked or unlocking an unlocked account has no effect.
- When locking a System Security Officer's login account, sp_locklogin verifies that at least one other unlocked System Security Officer's account exists. Similarly, sp_locklogin verifies that there is always an unlocked System Administrator's account. An attempt to lock the last remaining unlocked System Administrator or System Security Officer account causes sp_locklogin to return an error message and fail.

Permissions Only a System Administrator or a System Security Officers can execute sp_locklogin.

See also *System procedures* – sp_addlogin, sp_modifylogin, sp_password

sp_logdevice

Description Moves the transaction log of a database with log and data on the same device to a separate database device.

Syntax sp_logdevice *dbname*, *devname*

Parameters *dbname*
– is the name of the database whose syslogs table, which contains the transaction log, to put on a specific logical device.

devname
– is the logical name of the device on which to put the syslogs table. This device must be a database device associated with the database (named in create database or alter database). Run sp_helpdb for a report on the database's devices.

Examples **Example 1**

```
create database products on default = "10M", logs = "2M"  
go  
sp_logdevice products, logs  
go
```

Creates the database products and puts the table products.syslogs on the database device logs.

Example 2

```
alter database test log on logdev
go
sp_logdevice test, logdev
go
```

For the database test with log and data on the same device, places the log for test on the log device logdev.

Usage

- The *sp_logdevice* procedure affects only future allocations of space for syslogs. This creates a window of vulnerability during which the first pages of your log remain on the same device as your data. Therefore, the preferred method of placing a transaction log on a separate device is the use of the *log on* option to create database, which immediately places the entire transaction log on a separate device.
- Place transaction logs on separate database devices, for both recovery and performance reasons.

A very small, noncritical database could keep its log together with the rest of the database. Such databases use *dump database* to back up the database and log and *dump transaction with truncate_only* to truncate the log.

- *dbcc checkalloc* and *sp_helplog* show some pages for syslogs still allocated on the database device until after the next dump transaction. After that, the transaction log is completely transferred to the device named when you executed *sp_logdevice*.
- The size of the device required for the transaction log varies, depending on the amount of update activity and the frequency of transaction log dumps. As a rule, allocate to the log device 10 percent to 25 percent of the space you allocate to the database itself.
- Use *sp_logdevice* only for a database with log and data on the same device. Do not use *sp_logdevice* for a database with log and data on separate devices.
- To increase the amount of storage allocated to the transaction log use *alter database*. If you used the *log on* option to create database to place a transaction log on a separate device, use:

```
sp_extendsegment segname, devname
```

to increase the size of the log segment. If you did not use *log on*, execute *sp_logdevice*.

The device or segment on which you put syslogs is used *only* for the syslogs table. To increase the amount of storage space allocated for the rest of the database, specify any device other than the log device when you issue the alter database command.

- Use the disk init command to format a new database device for databases or transaction logs.
- For more information, see the *System Administration Guide*.

Permissions Only the Database Owner or a System Administrator can execute sp_logdevice.

See also *Commands* – alter database, create database, dbcc, disk init, dump database, dump transaction, select

System procedures – sp_extendsegment, sp_helpdevice

sp_loginconfig

Description *Windows NT only* – Displays the value of one or all integrated security parameters.

Syntax sp_loginconfig ["parameter_name"]

Parameters *parameter_name*
 – is the name of the integrated security parameter you want to examine. Values are: login mode, default account, default domain, set host, key _, key \$, key @, and key #.

Examples **Example 1**

```
sp_loginconfig

name                config_item
-----
login mode          standard
default account     NULL
default domain      NULL
set host            false
key _               domain separator
key $               space
key @               space
key #               -
```

Displays the values of all integrated security parameters.

Example 2

```
sp_loginconfig "login mode"
name                config_item
-----
login mode          standard
```

Displays the value of the login mode security parameter.

Usage

- The values of integrated security parameters are stored in the Windows NT Registry. See the chapter on login security in *Configuring Adaptive Server for Windows NT* for instructions on changing the parameters.
- *sp_loginconfig* displays the *config_item* values that were in effect when you started Adaptive Server. If you changed the Registry values after starting Adaptive Server, those values are not reflected in the *sp_loginconfig* output.

Permissions

Only a System Administrator can execute *sp_loginconfig*.

See also

System procedures – *sp_revokelogin*

sp_logininfo

Description

Windows NT only – Displays all roles granted to Windows NT users and groups with *sp_grantlogin*.

Syntax

sp_logininfo ["*login_name*" | "*group_name*"]

Parameters

login_name
 – is the network login name of the Windows NT user.

group_name
 – is the Windows NT group name.

Examples

Example 1

```
sp_logininfo regularjoe
account name  mapped login name  type          privilege
-----
HAZE\regularjoe HAZE_regularjoe    user          'oper_role'
```

Displays the permissions granted to the Windows NT user “regularjoe.”

Example 2

```
sp_logininfo
account name          mapped login name
type
privilege
-----
BUILTIN\Administrators  BUILTIN\Administrators
group
'sa_role sso_role oper_role sybase_ts_role navigator_role
replication_role'
HAZE\regularjoe        HAZE_regularjoe
user
'oper_role'
PCSRE\randy            PCSRE_alexander
user
'default'
```

Displays all permissions that were granted to Windows NT users and groups with sp_grantlogin.

Usage

- sp_logininfo displays all roles granted to Windows NT users and groups with sp_grantlogin.
- You can omit the domain name and domain separator (\) when specifying the Windows NT user name or group name.

Permissions

Only a System Administrator can execute sp_logininfo.

See also

Commands – grant, setuser
System procedures – sp_displaylogin, sp_grantlogin, sp_revokelogin, sp_role, sp_who

sp_logiosize

Description

Changes the log I/O size used by Adaptive Server to a different memory pool when doing I/O for the transaction log of the current database.

Syntax

sp_logiosize ["default" | "size" | "all"]

Parameters

default

– sets the log I/O size for the current database to Adaptive Server’s default value (two logical pages), if a memory pool that is two logical pages is available in the cache. Otherwise, Adaptive Server sets the log I/O size to one logical page. Since default is a keyword, the quotes are required when specifying this parameter.

size

– is the size to set the log I/O for the current database. Values are multiples of the logical page size, up to four times the amount. You must enclose the value in quotes.

all

– displays the log I/O size configured for all databases grouped by the cache name.

Examples

Example 1

```
sp_logiosize
```

The transaction log for database 'master' will use I/O size of 2 Kbytes.

Displays the log I/O size configured for the current database.

Example 2

```
sp_logiosize "8"
```

Changes the log I/O size of the current database to use the 8K memory pool. If the database’s transaction log is bound to a cache that does not have an 8K memory pool, Adaptive Server returns an error message indicating that such a pool does not exist, and the current log I/O size does not change.

Example 3

```
sp_logiosize "default"
```

Changes the log I/O size of the current database to Adaptive Server’s default value (one logical page size). If a memory pool the size of the logical page size does not exist in the cache used by the transaction log, Adaptive Server uses the 2K memory pool.

Example 4

```
sp_logiosize "all"
```

```
Cache name: default data cache
Data base                               Log I/O Size
-----
master                                   2 Kb
```

tempdb	2 Kb
model	2 Kb
sybsystemprocs	2 Kb
pubs3	2 Kb
pubtune	2 Kb
dbccdb	2 Kb
sybsyntax	2 Kb

Displays the log I/O size configured for all databases.

Usage

- sp_logiosize displays or changes the log I/O size for the current database. Any user can execute sp_logiosize to display the configured log I/O size. Only a System Administrator can change the log I/O size.
- If you specify sp_logiosize with no parameters, Adaptive Server displays the log I/O size of the current database.
- When you change the log I/O size, it takes effect immediately. Adaptive Server records the new I/O size for the database in the sysattributes table.
- Any value you specify for sp_logiosize must correspond to an existing memory pool configured for the cache used by the database's transaction log. Specify these pools using the sp_poolconfig system procedure.

Adaptive Server defines the default log I/O size of a database as two logical pages, if a memory pool the size of two logical pages is available in the cache. Otherwise, Adaptive Server sets the log I/O size to one logical page (a memory pool of one logical page is always present in any cache). For most work loads, a log I/O size of two logical pages performs much better than one of one logical page, so each cache used by a transaction log should have a memory pool the size of a logical page. For more information about configuring caches and memory pools, see the *System Administration Guide* and the *Performance and Tuning Guide*.

- If the transaction logs for one or more databases are bound to a cache of type logonly, any memory pools in that cache that have I/O sizes larger than the log I/O size defined for those databases will *not* be used.

For example, on a 2K server, assume that only two databases have their transaction logs bound to a “log only” cache containing 2K, 4K, and 8K memory pools. By default, *sp_logiosize* sets the log I/O size for these parameters at 4K, and the 8K pool is not used. Therefore, to avoid wasting cache space, be cautious when configuring the log I/O size.

- During recovery, only the logical page size memory pool of the default cache is active, regardless of the log I/O size configured for a database. Transactions logs are read into this pool of the default cache, and all transactions that must be rolled back, or rolled forward, read data pages into the default data cache.

Permissions

Only a System Administrator can execute *sp_logiosize* to change the log I/O size for the current database. Any user can execute *sp_logiosize* to display the log I/O size values.

See also

System procedures – *sp_cacheconfig*, *sp_poolconfig*

Procedures: *sp_modifylogin* – *sp_object_stats*

sp_modifylogin

Description Modifies the default database, default language, default role activation, or full name for an Adaptive Server login account. Changes the password expiration interval, the minimum password length, and the maximum number of failed logins allowed for a specified login.

Syntax *sp_modifylogin* {*loginame*|all overrides}, *option*, *value*

Parameters *loginame*
 – is the login account to be modified.

"all overrides"
 – removes the system overrides that were set using the "passwd expiration", "min passwd length", or "max failed_logins" parameters. To remove all the login-specific values, specify:

```
sp_modifylogin "all overrides" "option" -1
```

option
 – specifies the name of the option to be changed. The options are:

Option	Definition
defdb	The "home" database to which the user is connected when he or she logs in.
deflanguage	The official name of the user's default language.
fullname	The user's full name.
"add default role"	The role or roles to be activated by default at login.
"drop default role"	The role or roles to be dropped from the list of roles activated by default at login. This option affects only user-defined roles, not system roles.
"passwd expiration"	The password expiration interval in days. It can be anyvalue between 0 and 32767, inclusive.
"min passwd length"	The minimum password length required for the specified login. It can be any value between 0 and 30, inclusive. 0 specifies that no password is required. The default is 6.

Option	Definition
"max failed_logins"	The number of allowable failed login attempts for the specified login. It can be any value between 0 and 32767, inclusive.

value

– is the value of the option you specified for the *option* parameter. The *value* parameter is a character datatype; therefore, quotes are required for positive and negative numeric values.

Examples

Example 1

```
sp_modifylogin sarah, defdb, "pubs2"
```

Changes the default database for “sarah” to pubs2.

Example 2

```
sp_modifylogin claire, deflanguage, "french"
```

Sets the default language for “claire” to French.

Example 3

```
sp_modifylogin clemens, fullname, "Samuel Clemens"
```

Changes the full name of user “clemens” to “Samuel Clemens.”

Example 4

```
sp_modifylogin csmith, "add default role",  
specialist_role
```

Adds the specialist role to the list of roles activated by default when user csmith logs in.

Example 5

```
sp_modifylogin hpillai, "drop default role",  
intern_role
```

Drops the intern role from the list of roles activated by default when user “hpillai” logs in.

Example 6

```
sp_modifylogin "joe", "max failed_logins", "40"
```

Changes the maximum number of failed login attempts for the login “joe” to 40.

Example 7

```
sp_modifylogin "all overrides", "max failed_logins",  
"3"
```

Changes the overrides for maximum failed login attempts of all logins to 3.

Example 8

```
sp_modifylogin "all overrides", "max failed_logins",  
"-1"
```

Removes the overrides for maximum failed logins option for all logins.

Usage

- Set a default database, language, or full name either with `sp_modifylogin` or with `sp_addlogin` when first adding the user's login to Adaptive Server.
 - If you do not specify a default database, the user's default is master.
 - If you do not specify a language, the user's default language is set to the server's default language.
 - If you do not specify a full name, that column in `syslogins` remains blank.
- For more information about password expiration interval, minimum password length, and maximum number of failed logins, see "User-Defined Login Security" in the *System Administration Guide*.

Changing a user's default database

- After `sp_modifylogin` is executed to change the user's default database, the user is connected to the new *defdb* the next time he or she logs in. However, the user cannot access the database until the Database Owner gives the user access through `sp_adduser` or `sp_addalias`, or unless there is a "guest" user in the database's `sysusers` table. If the user does not have access to the database by any of these means, she or he is connected to master and an error message appears.
- If a user's default database is dropped, or if the user is dropped from the database, the user is connected to master on his or her next login, and an error message appears.
- If a user's default language is dropped from the server, the server-wide default language is used as the initial language setting, and a message appears.

Changing a user's role activation

- Use `sp_modifylogin` to set a role to be activated by default at login or to drop a role from those activated by default at login.

Permissions

Only a System Administrator can execute `sp_modifylogin` to change the default database, default language, or full name of another user. Only a System Security Officer can execute `sp_modifylogin` to activate another user's roles by default at login. Any user can execute `sp_modifylogin` to change his or her own login account.

See also

System procedures – `sp_activeroles`, `sp_addlogin`, `sp_displaylogin`, `sp_displayroles`, `sp_helprotect`

sp_modify_resource_limit

Description

Changes a resource limit by specifying a new limit value, or the action to take when the limit is exceeded, or both.

Syntax

```
sp_modify_resource_limit {name, appname }  
    , rangename , limittype , limitvalue , enforced  
    , action , scope
```

Parameters

name

– is the Adaptive Server login to which the limit applies. You must specify either a *name* or an *appname* or both. To modify a limit that applies to all users of a particular application, specify a *name* of null.

appname

– is the name of the application to which the limit applies. You must specify either a *name* or an *appname* or both. If the limit applies to all applications used by *name*, specify an *appname* of null. If the limit governs a particular application, specify the application name that the client program passes to the Adaptive Server in the login packet.

rangename

– is the time range during which the limit is enforced. You cannot modify this value, but you must specify a non-null value to uniquely identify the resource limit.

limittype

– is the type of resource to which the limit applies. You cannot modify this value, but you must specify a non-null value to uniquely identify the resource limit. The value must be one of the following:

Limit Type	Description
row_count	Limits the number of rows a query can return
elapsed_time	Limits the number of seconds in wall-clock time that a query batch or transaction can run
io_cost	Limits either the actual cost, or the optimizer's cost estimate, for processing a query

limit_value

– is the maximum amount of the server resource that the login or application can use before Adaptive Server enforces the limit. This must be a positive integer less than or equal to 2^{31} or null to retain the existing value. The following table indicates what value to specify for each limit type:

Limit Type	Limit Value
row_count	The maximum number of rows a query can return before the limit is enforced
elapsed_time	The maximum number of seconds in wall-clock time that a query batch or transaction can run before the limit is enforced
io_cost	A unitless measure derived from optimizer's costing formula

enforced

– determines whether the limit is enforced prior to or during query execution. You cannot modify this value. Use null as a placeholder.

action

– is the action to take when the limit is exceeded. The following codes apply to all limit types:

Action Code	Description
1	Issues a warning
2	Aborts the query batch
3	Aborts the transaction
4	Kills the session
null	Retains the existing value

scope

– is the scope of the limit. You cannot modify this value. You can use null as a placeholder.

Examples

Example 1

```
sp_modify_resource_limit robin, NULL, weekends,  
row_count, 3000, NULL, 1, NULL
```

Modifies a resource limit that applies to all applications used by “robin” during the *weekends* time range. The limit issues a warning when a query is expected to return more than 3000 rows.

Example 2

```
sp_modify_resource_limit NULL, acctg,  
"at all times", elapsed_time, 45, 2, 2, 6
```

Modifies a resource limit that applies to the *acctg* application on all days of the week and at all times of the day. The limit aborts the query batch when estimated query processing time exceeds 45 seconds.

Usage

- You cannot change the login or application to which a limit applies or specify a new time range, limit type, enforcement time, or scope.
- The modification of a resource limit causes the limits for each session for that login and/or application to be rebound at the beginning of the next query batch for that session.
- For more information, see the System Administration Guide.

Permissions

Only a System Administrator can execute `sp_modify_resource_limit`.

See also

System procedures – `sp_add_resource_limit`, `sp_drop_resource_limit`, `sp_help_resource_limit`

sp_modify_time_range

Description

Changes the start day, start time, end day, and/or end time associated with a named time range.

Syntax

```
sp_modify_time_range name, startday, endday, starttime, endtime
```

Parameters

name

– is the name of the time range. This must be the name of a time range stored in the `systemtimeranges` system table of the master database.

startday

– is the day of the week on which the time range begins. This must be the full weekday name for the default server language, as stored in the `syslanguages` system table of the master database, or null to keep the existing *startday*.

endday

– is the day of the week on which the time range ends. This must be the full weekday name for the default server language, as stored in the `syslanguages` system table of the master database, or null to keep the existing end day. The *endday* can fall either earlier or later in the week than the *startday*, or it can be the same day as the *startday*.

starttime

– is time of day at which the time range begins. Specify the *starttime* in terms of a twenty-four hour clock, with a value between 00:00 and 23:59. Use the following form:

"HH:MM"

or null to keep the existing *starttime*.

endtime

– is the time of day at which the time range ends. Specify the *endtime* in terms of a twenty-four hour clock, with a value between 00:00 (midnight) and 23:59. Use the following form:

"HH:MM"

or null to keep the existing *endtime*. The *endtime* must occur later in the day than the *starttime*, unless *endtime* is 00:00.

Note For time ranges that span the entire day, specify a start time of “00:00” and an end time of “23:59”.

Examples

Example 1

```
sp_modify_time_range business_hours, NULL,  
Saturday, NULL, NULL
```

Changes the end day of the *business_hours* time range from Friday to Saturday. Retains the existing start day, start time, and end time.

Example 2

```
sp_modify_time_range before_hours, Monday,  
Saturday, NULL, "08:00"
```

Usage	<p>Specifies a new end day and end time for the <i>before_hours</i> time range.</p> <ul style="list-style-type: none">• You cannot modify the “at all times” time range.• It is possible to modify a time range so that it overlaps with one or more other time ranges.• The modification of time ranges through the system stored procedures does not affect the active time ranges for sessions currently in progress.• Changes to a resource limit that has a transaction as its scope does not affect any transactions currently in progress.• For more information, see the System Administration Guide.
Permissions	Only a System Administrator can execute <code>sp_modify_time_range</code> .
See also	<i>System procedures</i> – <code>sp_add_resource_limit</code> , <code>sp_add_time_range</code> , <code>sp_drop_time_range</code>

sp_modifythreshold

Description	Modifies a threshold by associating it with a different threshold procedure, free-space level, or segment name. You <i>cannot</i> use <code>sp_modifythreshold</code> to change the amount of free space or the segment name for the last-chance threshold.
Syntax	<code>sp_modifythreshold <i>dbname</i>, <i>segname</i>, <i>free_space</i> [, <i>new_proc_name</i>] [, <i>new_free_space</i>] [, <i>new_segname</i>]</code>
Parameters	<p><i>dbname</i></p> <ul style="list-style-type: none">– is the database for which to change the threshold. This must be the name of the current database. <p><i>segname</i></p> <ul style="list-style-type: none">– is the segment for which to monitor free space. Use quotes when specifying the “default” segment. <p><i>free_space</i></p> <ul style="list-style-type: none">– is the number of free pages at which the threshold is crossed. When free space in the segment falls below this level, Adaptive Server executes the associated stored procedure.

new_proc_name

– is the new stored procedure to execute when the threshold is crossed. The procedure can be located in any database on the current Adaptive Server or on an Open Server. Thresholds cannot execute procedures on remote Adaptive Servers.

new_free_space

– is the new number of free pages to associate with the threshold. When free space in the segment falls below this level, Adaptive Server executes the associated stored procedure.

new_segname

– is the new segment for which to monitor free space. Use quotes when specifying the “default” segment.

Examples

Example 1

```
sp_modifythreshold mydb, "default", 200, NULL, 175
```

Modifies a threshold on the “default” segment of the mydb database to execute when free space on the segment falls below 175 pages instead of 200 pages. NULL is a placeholder indicating that the procedure name is not being changed.

Example 2

```
sp_modifythreshold mydb, data_seg, 250, new_proc
```

Modifies a threshold on the data_seg segment of mydb so that it executes the new_proc procedure.

Usage

- For more information, see the System Administration Guide.

Crossing a threshold

- When a threshold is crossed, Adaptive Server executes the associated stored procedure. Adaptive Server uses the following search path for the threshold procedure:
 - If the procedure name does not specify a database, Adaptive Server looks in the database in which the threshold was crossed.
 - If the procedure is not found in this database and the procedure name begins with “sp_”, Adaptive Server looks in the sybssystemprocs database.

If the procedure is not found in either database, Adaptive Server sends an error message to the error log.

- Adaptive Server uses a *hysteresis value*, the global variable @@thresh_hysteresis, to determine how sensitive thresholds are to variations in free space. Once a threshold executes its procedure, it is deactivated. The threshold remains inactive until the amount of free space in the segment rises to @@thresh_hysteresis pages above the threshold. This prevents thresholds from executing their procedures repeatedly in response to minor fluctuations in free space.

The last-chance threshold

- By default, Adaptive Server monitors the free space on the segment where the log resides and executes sp_thresholdaction when the amount of free space is less than that required to permit a successful dump of the transaction log. This amount of free space, the **last-chance threshold**, is calculated by Adaptive Server and cannot be changed by users.
- If the last-chance threshold is crossed before a transaction is logged, Adaptive Server suspends the transaction until log space is freed. Use sp_dboption to change this behavior for a particular database. Setting the abort tran on log full option to true causes Adaptive Server to roll back all transactions that have not yet been logged when the last-chance threshold is crossed.
- You cannot use sp_modifythreshold to change the free-space value or segment name associated with the last-chance threshold.
- Only databases that store their logs on a separate segment can have a last-chance threshold. Use sp_logdevice to move the transaction log to a separate device.

Other thresholds

- Each database can have up to 256 thresholds, including the last-chance threshold.
- Each threshold must be at least 2 times @@thresh_hysteresis pages from the next closest threshold.
- Use sp_helpthreshold for information about existing thresholds.
- Use sp_dropthreshold to drop a threshold from a segment.

Creating threshold procedures

- Any user with create procedure permission can create a threshold procedure in a database. Usually, a System Administrator creates sp_thresholdaction in the master database, and Database Owners create threshold procedures in user databases.

- *sp_modifythreshold* does not verify that the specified procedure exists. It is possible to associate a threshold with a procedure that does not yet exist.
- *sp_modifythreshold* checks to ensure that the user modifying the threshold procedure has been directly granted the “sa_role”. All system roles active when the threshold procedure is modified are entered in *systhresholds* as valid roles for the user writing the procedure. However, only directly granted system roles are activated when the threshold fires. Indirectly granted system roles and user-defined roles are not activated.
- Adaptive Server passes four parameters to a threshold procedure:
 - *@dbname*, *varchar(30)*, which identifies the database
 - *@segment_name*, *varchar(30)*, which identifies the segment
 - *@space_left*, *int*, which indicates the number of free pages associated with the threshold
 - *@status*, *int*, which has a value of 1 for last-chance thresholds and 0 for other thresholds

These parameters are passed by position rather than by name; your threshold procedure can use other names for them, but the procedure must declare them in the order shown and with the correct datatypes.

- It is not necessary to create a different procedure for each threshold. To minimize maintenance, create a single threshold procedure in the *sysystemprocs* database that can be executed by all thresholds.
- Include *print* and *raiserror* statements in the threshold procedure to send output to the error log.

Executing threshold procedures

- Tasks that are initiated when a threshold is crossed execute as background tasks. These tasks do not have an associated terminal or user session. If you execute *sp_who* while these tasks are running, the *status* column shows “background”.
- Adaptive Server executes the threshold procedure with the permissions of the user who modified the threshold, at the time he or she executed *sp_modifythreshold*, minus any permissions that have since been revoked.
- Each threshold procedure uses one user connection, for as long as it takes to execute the procedure.

Disabling free-space accounting

Warning! System procedures cannot provide accurate information about space allocation when free-space accounting is disabled.

- Use the no free space acctg option of sp_dboption to disable free-space accounting on non-log segments.
- You cannot disable free-space accounting on log segments.

Permissions Only the Database Owner or a System Administrator can execute sp_modifythreshold.

See also *Commands* – create procedure, dump transaction

System procedures – sp_addthreshold, sp_dboption, sp_droptreshold, sp_helpthreshold, sp_thresholdaction

sp_monitor

Description Displays statistics about Adaptive Server.

Syntax sp_monitor

Parameters None.

Examples

```
sp_monitor
last_run          current_run      seconds
-----
Jan 29 1987 10:11AM  Jan 29 1987 10:17AM  314

cpu_busy          io_busy         idle
-----
4250(215)-68%    67(1)-0%       109(100)-31%

packets_received  packets_sent    packet_errors
-----
781(15)          10110(9596)    0(0)

total_read        total_write    total_errors    connections
-----
394(67)          5392(53)      0(0)            15(1)
```

Reports information about how busy Adaptive Server has been.

Usage

- Adaptive Server keeps track of how much work it has done in a series of global variables. *sp_monitor* displays the current values of these global variables and how much they have changed since the last time the procedure executed.
- For each column, the statistic appears in the form *number(number)-number%* or *number(number)*.
 - The first number refers to the number of seconds (for “*cpu_busy*”, “*io_busy*”, and “*idle*”) or the total number (for the other columns) since Adaptive Server restarted.
 - The number in parentheses refers to the number of seconds or the total number since the last time *sp_monitor* was run. The percent sign indicates the percentage of time since *sp_monitor* was last run.

For example, if the report shows “*cpu_busy*” as “4250(215)-68%”, it means that the CPU has been busy for 4250 seconds since Adaptive Server was last started, 215 seconds since *sp_monitor* last ran, and 68 percent of the total time since *sp_monitor* was last run.

For the “*total_read*” column, the value 394(67) means there have been 394 disk reads since Adaptive Server was last started, 67 of them since the last time *sp_monitor* was run.

- Table 20-1 describes the columns in the *sp_monitor* report, the equivalent global variables, if any, and their meanings. With the exception of “*last_run*”, “*current_run*” and “*seconds*”, these column headings are also the names of global variables—except that all global variables are preceded by @@. There is also a difference in the units of the numbers reported by the global variables—the numbers reported by the global variables are not milliseconds of CPU time, but machine ticks.

Table 20-1: Columns in the *sp_monitor* report

Column Heading	Equivalent Variable	Meaning
<i>last_run</i>		Clock time at which the <i>sp_monitor</i> procedure last ran.
<i>current_run</i>		Current clock time.
<i>seconds</i>		Number of seconds since <i>sp_monitor</i> last ran.
<i>cpu_busy</i>	@@ <i>cpu_busy</i>	Number of seconds in CPU time that Adaptive Server’s CPU was doing Adaptive Server work.
<i>io_busy</i>	@@ <i>io_busy</i>	Number of seconds in CPU time that Adaptive Server has spent doing input and output operations.

Column Heading	Equivalent Variable	Meaning
idle	@@idle	Number of seconds in CPU time that Adaptive Server has been idle.
packets_received	@@pack_received	Number of input packets read by Adaptive Server.
packets_sent	@@pack_sent	Number of output packets written by Adaptive Server.
packet_errors	@@packet_errors	Number of errors detected by Adaptive Server while reading and writing packets.
total_read	@@total_read	Number of disk reads by Adaptive Server.
total_write	@@total_write	Number of disk writes by Adaptive Server.
total_errors	@@total_errors	Number of errors detected by Adaptive Server while reading and writing.
connections	@@connections	Number of logins or attempted logins to Adaptive Server.

- The first time sp_monitor runs after Adaptive Server start-up, the number in parentheses is meaningless.
- Adaptive Server’s housekeeper task uses the server’s idle cycles to write changed pages from cache to disk. This process affects the values of the “cpu_busy”, “io_busy”, and “idle” columns reported by sp_monitor. To disable the housekeeper task and eliminate these effects, set the housekeeper free write percent configuration parameter to 0:

```
sp_configure "housekeeper free write percent", 0
```

Permissions

Only a System Administrator can execute sp_monitor.

See also

System procedures – sp_who

sp_monitorconfig

Description

Displays cache usage statistics regarding metadata descriptors for indexes, objects, and databases. sp_monitorconfig also reports statistics on auxiliary scan descriptors used for referential integrity queries, and usage statistics for transaction descriptors and DTX participants.

Syntax

sp_monitorconfig "configname"

Parameters

configname

– is all or part of the configuration parameter name whose monitoring information is being queried. Valid configuration parameters are open indexes, open objects, open databases, aux scan descriptors, txn to pss ratio, number of dtx participants, and all. Specifying all displays descriptor help information for all indexes, objects, databases, and auxiliary scan descriptors in the server.

Examples

Example 1

```
sp_monitorconfig "open"
```

Configuration option is not unique.

option_name	config_value	run_value
curread change w/ open cursors	1	1
number of open databases	12	12
number of open indexes	500	500
number of open objects	500	500
open index hash spinlock ratio	100	100
open index spinlock ratio	100	100
open object spinlock ratio	100	100

Example 2

```
sp_monitorconfig "open objects"
```

Usage information at date and time: Jan 14 1997 8:54AM.

Name	# Free	# Active	% Active	# Max Ever Used	Re-used
number of open objects	217	283	56.60	300	No

In this example, there are 283 active object metadata descriptors, with 217 free. The maximum used at a peak period since Adaptive Server was last started is 300. You can then reset the size to 330, for example, to accommodate the 300 maximum used metadata descriptors, plus space for 10 percent more:

```
sp_configure "number of open objects", 330
```

Example 3

```
sp_monitorconfig "open indexes"
```

Usage information at date and time: Jan 14 1997 8:55AM.

Name	# Free	# Active	% Active	# Max Ever Used	Re-used
number of open indexes	556	44	7.33	44	No

In this example, the maximum number of index metadata descriptors is 44. You can reset the size to 100, the minimum acceptable value:

```
sp_configure "number of open indexes", 100
```

Example 4

```
sp_monitorconfig "aux scan descriptors"
```

Usage information at date and time: Jan 14 1997 8:56AM.

Name	# Free	# Active	% Active	# Max Ever Used	Re-used
number of aux scan descriptors	170	30	15.00	32	NA

In this example, the number of active scan descriptors is 30, though Adaptive Server is configured to use 200. Use the number of aux scan descriptors configuration parameter to reset the value to at least 32. A safe setting is 36, to accommodate the 32 scan descriptors, plus space for 10 percent more.

Example 5

```
sp_monitorconfig "number of open databases"
```

Usage information at date and time: Jan 14 1997 8:57AM.

Name	# Free	# Active	% Active	# Max Ever Used	Re-used
number of open databases	0	5	100.00	5	Yes

In this example, Adaptive Server is configured for 5 open databases, all of which have been used in the current session. However, as indicated by the Re-used column, an additional database needs to be opened. If all 5 databases are in use, an error may result, unless the descriptor for a database that is not in use can be reused. To prevent an error, reset number of open databases to a higher value.

Example 6

```
sp_monitorconfig "txn to pss ratio"
```

Usage information at date and time: Jun 18 1999 8:54AM.

Name	# Free	# Active	% Active	# Max Ever Used	Re-used
txn to pss NA ratio	784	80	10.20	523	

In this example, only 10.2 percent of the transaction descriptors are currently being used. However, the maximum number of transaction descriptors used at a peak period since Adaptive Server was last started is 523.

Usage

- *sp_monitorconfig* displays cache usage statistics regarding metadata descriptors for indexes, objects, and databases, such as the number of metadata descriptors currently in use by the server.
- *sp_monitorconfig* also reports the number of auxiliary scan descriptors in use. A scan descriptor manages a single scan of a table when queries are run on the table.
- The columns in the *sp_monitorconfig* output provide the following information:
 - # Free specifies the number of available metadata or auxiliary scan descriptors not currently used.
 - # Active specifies the number of metadata or auxiliary scan descriptors installed in cache (that is, active).
 - % Active specifies the percentage of cached or active metadata or auxiliary scan descriptors.
 - # Max Ever Used specifies the maximum number of metadata or auxiliary scan descriptors that have been in use since the server was started.
 - Re-used specifies whether a metadata descriptor was reused in order to accommodate an increase in indexes, objects, or databases in the server. The returned value is *Yes*, *No* or *NA* (for configuration parameters that do not support the reuse mechanism, such as the number of aux scan descriptors).
- Use the value in the # Max Ever Used column as a basis for determining an appropriate number of descriptors; be sure to add about 10 percent for the final setting. For example, if the maximum number of index metadata descriptors used is 142, you might set the number of open indexes configuration parameter to 157.

- If the *Re-used* column states *Yes*, reset the configuration parameter to a higher value. When descriptors need to be reused, there can be performance problems, particularly with open databases. An open database contains a substantial amount of metadata information, which means that to fill up an open database, Adaptive Server needs to access the metadata on the disk many times; the server can also have a spinlock contention problem. To check for spinlock contention, use the system procedure *sp_sysmon*. For more information, see the *Performance and Tuning Guide*. To find the current number of indexes, objects, or databases, use *sp_countmetadata*.
- To get an accurate reading, run *sp_monitorconfig* during a normal Adaptive Server peak time period. You can run *sp_monitorconfig* several times during the peak period to ensure that you are actually finding the maximum number of descriptors used.

Permissions

Only a System Administrator can execute *sp_monitorconfig*.

See also

System procedures – *sp_configure*, *sp_countmetadata*, *sp_helpconfig*, *sp_helpconstraint*

sp_object_stats

Description

Shows lock contention, lock wait-time, and deadlock statistics for tables and indexes.

Syntax

```
sp_object_stats interval [, top_n  
[, dbname, objname [, rpt_option ]]]
```

Parameters

interval

– specifies the time period for the sample. It must be in HH:MM:SS form, for example “00:20:00”.

top_n

– the number of objects to report, in order of contention. The default is 10.

dbname

– the name of the database to report on. If no database name is given, contention on objects in all databases is reported.

objname

– the name of a table to report on. If a table name is specified, the database name must also be specified.

rpt_option

– must be either *rpt_locks* or *rpt_objlist*.

Examples

Example 1

```
sp_object_stats "00:20:00"
```

Reports lock statistics on the top 10 objects server-wide.

Example 2

```
sp_object_stats "00:20:00", 5, pubtune
```

Reports only on tables in the *pubtune* database, and lists the five tables that experienced the highest contention.

Example 3

```
sp_object_stats "00:15:00", @rpt_option =
"rpt_objlist"
```

Shows only the names of the tables that had the highest locking activity, even if contention and deadlocking does not take place.

Usage

- *sp_object_stats* reports on the shared, update, and exclusive locks acquired on tables during a specified sample period. The following reports shows the titles tables:

```
Object Name: pubtune..titles (dbid=7,
objid=208003772,lockscheme=Datapages)
```

Page Locks	SH_PAGE	UP_PAGE	EX_PAGE\$
-----	-----	-----	-----
Grants:	94488	4052	4828
Waits:	532	500	776
Deadlocks:	4	0	24
Wait-time:	20603764 ms	14265708 ms	2831556 ms
Contention:	0.56%	10.98%	13.79%

*** Consider altering *pubtune..titles* to *Datarows* locking.

- Table 20-2 shows the meaning of the values.

Table 20-2: Output of sp_object_stats

Output Row	Value
Grants	The number of times the lock was granted immediately.
Waits	The number of times the task needing a lock had to wait.
Deadlocks	The number of deadlocks that occurred.
Wait-times	The total number of milliseconds that all tasks spent waiting for a lock.
Contention	The percentage of times that a task had to wait or encountered a deadlock.

- sp_object_stats recommends changing the locking scheme when total contention on a table is more than 15 percent, as follows:
 - If the table uses allpages locking, it recommends changing to datapages locking
 - If the table uses datapages locking, it recommends changing to datarows locking.
- rpt_option specifies the report type:
 - rpt_locks reports grants, waits, deadlocks and wait times for the tables with the highest contention. rpt_locks is the default.
 - rpt_objlist reports only the names of the objects that had the highest level of lock activity
- sp_object_stats creates a table named tempdb..syslkstats. This table is not dropped when the stored procedure completes, so it can be queried by a System Administrator using Transact-SQL.
- Only one user at a time should execute sp_object_stats. If more than one user tries to run sp_object_stats simultaneously, the second command may be blocked, or the results may be invalid.
- The tempdb..syslkstats table is dropped and re-created each time sp_object_stats is executed.
- The structure of tempdb..syslkstats is described in Table 20-3.

Table 20-3: Columns in the tempdb..syslkstats table

Column Name	Datatype	Description
dbid	smallint	Database ID
objid	int	Object ID
lockscheme	smallint	Integer values 1–3: Allpages = 1, Datapages = 2, Datarows = 3
page_type	smallint	Data page = 0, or index page = 1

Column Name	Datatype	Description
stat_name	char(30)	The statistics represented by this row
stat_value	float	The number of grants, waits or deadlocks, or the total wait time

The values in the `stat_name` column are composed of three parts:

- The first part is “ex” for exclusive lock, “sh” for shared lock, or “up” for update lock.
- The second part is “pg” for page locks, or “row” for row locks.
- The third part is “grants” for locks granted immediately, “waits” for locks that had to wait for other locks to be released, “deadlocks” for deadlocks, and “waittime” for the time waited to acquire the lock.
- If you specify a table name, `sp_object_stats` displays all tables by that name. If more than one user owns a table with the specified name, output for these tables displays the object ID, but not the owner name.

Permissions

Only a System Administrator can execute `sp_object_stats`.

See also

Commands – alter table

Procedures: *sp_passthru* – *sp_procxmode*

sp_passthru

Description	<i>Component Integration Services only</i> – Allows the user to pass a SQL command buffer to a remote server.
Syntax	<code>sp_passthru server, command, errcode, errmsg, rowcount [, arg1, arg2, ... argn]</code>
Parameters	<p><i>server</i> – is the name of a remote server to which the SQL command buffer will be passed. The class of this server must be a supported, non-local server class.</p> <p><i>command</i> – is the SQL command buffer. It can hold up to 255 characters.</p> <p><i>errcode</i> – is the error code returned by the remote server, if any. If no error occurred at the remote server, the value returned is 0.</p> <p><i>errmsg</i> – is the error message returned by the remote server. It can hold up to 255 characters. This parameter is set only if <i>errcode</i> is a nonzero number; otherwise NULL is returned.</p> <p><i>rowcount</i> – is the number of rows affected by the last command in the command buffer. If the command was an insert, delete, or update, this value represents the number of rows affected even though none were returned. If the last command was a query, this value represents the number of rows returned from the external server.</p> <p><i>arg1</i> ... <i>argn</i> – receives the results from the last row returned by the last command in the command buffer. You can specify up to 250 <i>arg</i> parameters. All must be declared as output parameters.</p>
Examples	<code>sp_passthru ORACLE, "select date from dual", @errcode output, @errmsg output, @rowcount output,</code>

@oradate output

Returns the date from the Oracle server in the output parameter *@oradate*. If an Oracle error occurs, the error code is placed in *@errcode* and the corresponding message is placed in *@errmsg*. The *@rowcount* parameter will be set to 1.

Usage

- *sp_passthru* allows the user to pass a SQL command buffer to a remote server. The syntax of the SQL statement or statements being passed is assumed to be the syntax native to the class of server receiving the buffer. No translation or interpretation is performed. Results from the remote server are optionally placed in output parameters.

Use *sp_passthru* only when Component Integration Services is installed and configured.

- You can include multiple commands in the command buffer. For some server classes, the commands must be separated by semicolons. Refer to the *Component Integration Services User's Guide* for a more complete discussion of query buffer handling in passthru mode.

Return Parameters

- The output parameters *arg1 ... argn* will be set to the values of corresponding columns from the last row returned by the last command in the command buffer. The position of the parameter determines which column's value the parameter will contain. *arg1* receives values from column 1, *arg2* receives values from column 2, and so on.
- If there are fewer optional parameters than there are returned columns, the excess columns are ignored. If there are more parameters than columns, the remaining parameters are set to NULL.
- An attempt is made to convert each column to the datatype of the output parameter. If the datatypes are similar enough to permit *implicit* conversion, the attempt will succeed. For information on implicit conversion, see "Datatype conversion functions." For information on which datatype represents the datatypes from each server class when in passthru mode, see the *Component Integration Services User's Guide*.

Permissions

Any user can execute *sp_passthru*.

See also

System procedures – *sp_autoconnect*, *sp_remotesql*

sp_password

Description	Adds or changes a password for an Adaptive Server login account.
Syntax	<code>sp_password caller_passwd, new_passwd [, loginame]</code>
Parameters	<p><i>caller_passwd</i></p> <ul style="list-style-type: none"> – is your password. When you are changing your own password, this is your old password. When a System Security Officer is using <code>sp_password</code> to change another user’s password, <i>caller_passwd</i> is the System Security Officer’s password. <p><i>new_passwd</i></p> <ul style="list-style-type: none"> – is the new password for the user, or for <i>loginame</i>. It must be at least 6 bytes long. Enclose passwords that include characters besides A-Z, a-z, or 0-9 in quotation marks. Also enclose passwords that begin with 0-9 in quotes. <p><i>loginame</i></p> <ul style="list-style-type: none"> – the login name of the user whose account password is being changed by the System Security Officer.
Examples	<p>Example 1</p> <pre>sp_password "3blindmice", "2mediumhot"</pre> <p>Changes your password from password from “3blindmice” to “2mediumhot.” (Enclose the passwords in quotes because they begin with numerals.)</p> <p>Example 2</p> <pre>sp_password "2tomato", sesame1, victoria</pre> <p>A System Security Officer whose password is “2tomato” has changed Victoria’s password to “sesame1.”</p> <p>Example 3</p> <pre>sp_password null, "16tons"</pre> <p>Changes your password from NULL to “16tons.” Notice that NULL is not enclosed in quotes. (NULL is not a permissible new password.)</p> <p>Example 4</p> <pre>PRODUCTION...sp_password figaro, lilacs</pre> <p>Changes your password on the PRODUCTION server from “figaro” to “lilacs.”</p>
Usage	<ul style="list-style-type: none"> • Any user can change his or her password with <code>sp_password</code>.

- New passwords must be at least 6 characters long. They cannot be NULL.
- The encrypted text of *caller_passwd* must match the existing encrypted password of the caller. If it does not, sp_password returns an error message and fails. master.dbo.syslogins lists passwords in encrypted form.
- If a client program requires users to have the same password on remote servers as on the local server, users must change their passwords on all the remote servers before changing their local passwords. Execute sp_password as a remote procedure call on each remote server. See example 4.
- You can set the systemwide password expiration configuration parameter to establish a password expiration interval that forces all Adaptive Server login accounts to change passwords on a regular basis. For more information, see the *System Administration Guide*.

Permissions

Only a System Security Officer can execute sp_password to change another user's password. Any user can execute sp_password to change his or her own password.

See also

System procedures – sp_addlogin, sp_adduser

sp_placeobject

Description

Puts future space allocations for a table or index on a particular segment.

Syntax

sp_placeobject *segname*, *objname*

Parameters

segname

– is the name of the segment on which to locate the table or index.

objname

– is the name of the table or index for which to place subsequent space allocation on the segment *segname*. Specify index names in the form “*tablename.indexname*”.

Examples

Example 1

```
sp_placeobject segment3, authors
```

This command places all subsequent space allocation for the table authors on the segment named “segment3”.

Example 2

```
sp_placeobject indexes, 'employee.employee_nc'
```

This command places all subsequent space allocation for the employee table's index named `employee_nc` on the segment named `indexes`.

Usage

- You cannot change the location of future space allocations for system tables.
- Placing a table or an index on a particular segment does not affect the location of any existing table or index data. It affects only future space allocation. Changing the segment used by a table or an index can spread the data among multiple segments.
- If you use `sp_placeobject` with a clustered index, the table moves with the index.
- You can specify a segment when you create a table or an index with `create table` or `create index`. If you do not specify a segment, the data goes on the default segment.
- When `sp_placeobject` splits a table or an index across more than one disk fragment, the diagnostic command `dbcc` displays messages about the data that resides on the fragments that were in use for storage before `sp_placeobject` executed. Ignore those messages.
- You cannot use `sp_placeobject` on a partitioned table.

Permissions

Only the table owner, Database Owner, or System Administrator can execute `sp_placeobject`.

See also

Commands – `alter table`, `dbcc`

System procedures – `sp_addsegment`, `sp_dropsegment`, `sp_extendsegment`, `sp_helpindex`, `sp_helpsegment`

sp_plan_dbccdb

Description

Recommends suitable sizes for new `dbccdb` and `dbccalt` databases, lists suitable devices for `dbccdb` and `dbccalt`, and suggests a cache size and a suitable number of worker processes for the target database.

Syntax

```
sp_plan_dbccdb [dbname]
```

Parameters

dbname

– specifies the name of the target database. If *dbname* is not specified, `sp_plan_dbccdb` makes recommendations for all databases in `master.sysdatabases`.

Examples

Example 1

```
sp_plan_dbccdb master
```

Recommended size for dbccdb is 4MB.

dbccdb database already exists with size 8MB.

Recommended values for workspace size, cache size and process count are:

dbname	scan ws	text ws	cache	process count
master	64K	64K	640K	1

Returns configuration recommendations for creating a dbccdb database suitable for checking the master database. The dbccdb database already existed at the time this command was run, so the size of the existing database is provided for comparison.

Example 2

```
sp_plan_dbccdb
```

Recommended minimum size for dbccdb is 4MB.

Recommended values for workspace size, cache size and process count are:

dbname	scan ws	text ws	cache	process count
master	64K	64K	640K	1
tempdb	64K	64K	640K	1
model	64K	64K	640K	1
sybsystemprocs	272K	80K	640K	1
dbccdb	128K	64K	640K	1

Returns configuration recommendations for creating a dbccdb database suitable for checking all databases in the server. No dbccdb database existed at the time this command was run.

Example 3

```
sp_plan_dbccdb pubs2
```

Recommended size for dbccdb is 4MB.

Recommended devices for dbccdb are:

Logical Device Name	Device Size	Physical Device Name
sprocdev	28672	/remote/sybase/devices/srv_sprocs_dat
tun_dat	8192	/remote/sybase/devices/srv_tun_dat
tun_log	4096	/remote/sybase/devices/srv_tun_log

Recommended values for workspace size, cache size and process count are:

dbname	scan ws	text ws	cache	process count
pubs2	64K	64K	640K	1

Returns configuration recommendations for creating a dbccdb database suitable for checking pubs2.

Usage

- sp_plan_dbccdb recommends suitable sizes for creating new dbccdb and dbccalt databases, lists suitable devices for the new database, and suggests cache size and a suitable number of worker processes for the target database.
- If you specify dbccdb, sp_plan_dbccdb recommends values for dbccalt, the alternate database. If you specify dbccalt, sp_plan_dbccdb recommends values for dbccdb.
- sp_plan_dbccdb does not report values for existing dbccdb and dbccalt databases. To gather configuration parameters for an existing dbccdb or dbccalt database, use sp_dbcc_evaluatedb.
- For information on the dbcc stored procedures for maintaining dbccdb and for generating reports from dbccdb, see Chapter 27, “dbcc Stored Procedures.”

Permissions

Only the System Administrator or Database Owner can execute sp_plan_dbccdb. Only the System Administrator can execute sp_plan_dbccdb without specifying a database name.

See also

Commands – dbcc
System procedures – sp_dbcc_evaluatedb

sp_poolconfig

Description	Creates, drops, resizes, and provides information about memory pools within data caches.
Syntax	<p>To create a memory pool in an existing cache, or to change pool size:</p> <pre>sp_poolconfig cache_name [, "mem_size [P K M G]", "config_poolK" [, "affected_poolK"]]</pre> <p>To change a pool's wash size:</p> <pre>sp_poolconfig cache_name, "io_size ", "wash=size[P K M G]"</pre> <p>To change a pool's asynchronous prefetch percentage:</p> <pre>sp_poolconfig cache_name, "io_size ", "local async prefetch limit=percent "</pre>
Parameters	<p><i>cache_name</i></p> <ul style="list-style-type: none">– is the name of an existing data cache. <p><i>mem_size</i></p> <ul style="list-style-type: none">– is the size of the memory pool to be created or the new total size for an existing pool, if a pool already exists with the specified I/O size. The minimum size of a pool is 512K. Specify size units with P for pages, K for kilobytes, M for megabytes, or G for gigabytes. The default is kilobytes. <p><i>config_pool</i></p> <ul style="list-style-type: none">– is the I/O size performed in the memory pool where the memory is to be allocated or removed. <p>Valid I/O sizes are multiples of the logical page size, up to four times the amount.</p> <p><i>affected_pool</i></p> <ul style="list-style-type: none">– is the size of I/O performed in the memory pool where the memory is to be deallocated. If <i>affected_pool</i> is not specified, the memory is taken from the logical page size memory pool. <p><i>io_size</i></p> <ul style="list-style-type: none">– is the size of I/O performed in the memory pool where the wash size is to be reconfigured. The combination of cache name and I/O size uniquely identifies a memory pool. <p><i>wash=size</i></p> <ul style="list-style-type: none">– Changes the wash size (the point in the cache at which Adaptive Server writes dirty pages to disk) for a memory pool.

local async prefetch limit=*percent*

- sets the percentage of buffers in the pool that can be used to hold buffers that have been read into cache by asynchronous prefetch, but that have not yet been used.

Examples

Example 1

```
sp_poolconfig pub_cache, "10M", "16K"
```

Creates a 16K pool in the data cache `pub_cache` with 10MB of space. All space is taken from the default 2K memory pool.

Example 2

```
sp_poolconfig pub_cache, "16M", "32K", "64K"
```

Moves 16MB of space to the 32K pool from the 64K pool of `pub_cache`.

Example 3

```
sp_poolconfig "pub_cache"
```

Reports the current configuration of `pub_cache`.

Example 4

```
sp_poolconfig pub_cache, "0K", "16K"
```

Removes the 16K memory pool from `pub_cache`, placing all of the memory assigned to it in the 2K pool.

Example 5

```
sp_poolconfig pub_cache, "2K", "wash=508K"
```

Changes the wash size of the 2K pool in `pubs_cache` to 508K.

Example 6

```
sp_poolconfig pub_cache, "2K",  
"local async prefetch limit=15"
```


Changes the asynchronous prefetch limit for the 2K pool to 15 percent.

Usage

- When you create a data cache with `sp_cacheconfig`, all space is allocated to the logical page size memory pool. `sp_poolconfig` divides the data cache into additional pools with larger I/O sizes.
- If no large I/O memory pools exist in a cache, Adaptive Server performs I/O in 2K units, the size of a data page, for all of the objects bound to the cache. You can often enhance performance by configuring pools that perform large I/O. A 16K memory pool reads and writes eight data pages in a single I/O operation.

- The combination of cache name and I/O size must be unique. In other words, you can have only one pool of a given I/O size in a particular data cache.
- Only one sp_poolconfig command can be active on a single cache at one time. If a second sp_poolconfig command is issued before the first one completes, it sleeps until the first command completes.
- Figure 21-1 shows a data cache on a server that uses 2K logical pages with:
 - The default data cache with a 2K pool and a 16K pool
 - A user cache with a 2K pool and a 16K pool
 - A log cache with a 2K pool and a 4K pool

Figure 21-1: Data cache with default and user-defined caches



- You can create pools with I/O sizes up to 16K in the default data cache.

- The minimum size of a memory pool is 512K. You cannot reduce the size of any memory pool in any cache to less than 512K by transferring memory to another pool.
- Two circumstances can create pool less than 512K:
 - If you attempt to delete a pool by setting its size to zero, and some of the pages are in use, *sp_poolconfig* reduces the pool size as much as possible, and prints a warning message. The status for the pool is set to “Unavailable/deleted”.
 - If you attempt to move buffers to create a new pool, and enough buffers cannot be moved to the new pool, *sp_poolconfig* moves as many buffers as it can, and the cache status is set to “Unavailable/too small.”

In both of these cases, you can retry to command at a later time. The pool will also be deleted or be changed to the desired size when the server is restarted.

- You can create memory pools while Adaptive Server is active; no restart is needed for them to take effect. However, Adaptive Server can move only “free” buffers (buffers that are not in use or that do not contain changes that have not been written to disk). When you configure a pool or change its size, Adaptive Server moves as much memory as possible to the pool and prints an informational message showing the requested size and the actual size of the pool. After a restart of Adaptive Server, all pools are created at the configured size.
- The following commands perform only 2K I/O: create database, alter database, some dbcc commands, disk init, and drop table. dbcc checktable can perform large I/O, and dbcc checkdb performs large I/O on tables and 2K I/O on indexes. Also, recovery uses only the 2K memory pool: all pages are read into and changed in the 2K pool of the default cache. Be sure that your default 2K pool is large enough for these activities.
- Most Adaptive Servers perform best with I/O configured for transactions logs that is twice the logical page size. Adaptive Server uses the default I/O size of twice the logical page size if the default cache or a cache with a transaction log bound to it is configured with a memory pool twice the logical page size. Otherwise, it uses the logical page size memory pool.

- You can increase the default log I/O size for a database using the `sp_logiosize` system procedure. However, the I/O size you specify must have memory pools of the same size in the cache bound to the transaction log. If not, Adaptive Server uses the logical page size memory pools.

Wash percentage

- The default value for the wash size is computed as follows:
 - If the pool size is less than 300MB, the default wash size is set to 20 percent of the buffers in the pool
 - If the pool size is greater than 300MB, the default wash size is 20 percent of the number of buffers in 300MB
- The minimum setting for the wash size is 10 buffers, and the maximum setting is 80 percent of the size of the pool.
- Each memory pool contains a wash area at the least recently used (LRU) end of the chain of buffers in that pool. Once dirty pages (pages that have been changed while in cache) move into the wash area, Adaptive Server initiates asynchronous writes on these pages. The wash area must be large enough so that pages can be written to disk before they reach the LRU end of the pool. Performance suffers when Adaptive Server needs to wait for clean buffers.

The default percentage, placing 20 percent of the buffers in the wash area, is sufficient for most applications. If you are using an extremely large memory pool, and your applications have a very high data modification rate, you may want to increase the size to 1 or 2 percent of the pool. Contact Sybase Technical Support for more information about choosing an effective wash size.

Local asynchronous prefetch percentage

- The default value for a pool's asynchronous prefetch percentage is set by the configuration parameter `global async prefetch limit`. The pool limit always overrides the global limit.
- To disable prefetch in a pool (if the global limit is a nonzero number), set the pool's limit to 0.
- For information on the performance impact of changes to the asynchronous prefetch limit, see the Performance and Tuning Guide.

Permissions

Only a System Administrator can execute `sp_poolconfig` to reconfigure memory pools within data caches. Any user can use `sp_poolconfig` to get information about memory pools.

See also *System procedures* – *sp_cacheconfig*, *sp_helpcache*, *sp_unbindcache*, *sp_unbindcache_all*

sp_primarykey

Description	Defines a primary key on a table or view.
Syntax	<code>sp_primarykey <i>tablename</i>, <i>col1</i> [, <i>col2</i>, <i>col3</i>, ..., <i>col8</i>]</code>
Parameters	<i>tablename</i> – is the name of the table or view on which to define the primary key. <i>col1</i> – is the name of the first column that makes up the primary key. The primary key can consist of from one to eight columns.
Examples	Example 1 <pre>sp_primarykey authors, au_id</pre> Defines the <i>au_id</i> field as the primary key of the table <i>authors</i> . Example 2 <pre>sp_primarykey employees, lastname, firstname</pre> Defines the combination of the fields <i>lastname</i> and <i>firstname</i> as the primary key of the table <i>employees</i> .
Usage	<ul style="list-style-type: none">• Executing <i>sp_primarykey</i> adds the key to the <i>syskeys</i> table. Only the owner of a table or view can define its primary key. <i>sp_primarykey</i> does not enforce referential integrity constraints; use the primary key clause of the <i>create table</i> or <i>alter table</i> command to enforce a primary key relationship.• Define keys with <i>sp_primarykey</i>, <i>sp_commonkey</i>, and <i>sp_foreignkey</i> to make explicit a logical relationship that is implicit in your database design. An application program can use the information.• A table or view can have only one primary key. To display a report on the keys that have been defined, execute <i>sp_helpkey</i>.• The installation process runs <i>sp_primarykey</i> on the appropriate columns of the system tables.
Permissions	Only the owner of the specified table or view can execute <i>sp_primarykey</i> .

See also *Commands* – alter table, create table, create trigger
System procedures – sp_commonkey, sp_dropkey, sp_foreignkey, sp_helpjoins, sp_helpkey

sp_processmail

Description *Windows NT only* – Reads, processes, sends, and deletes messages in the Adaptive Server message inbox, using the xp_findnextmsg, xp_readmail, xp_sendmail, and xp_deletemail system extended stored procedures (ESPs).

Syntax sp_processmail [subject] [, originator [, dbuser
[, dbname [, filetype [, separator]]]]]

Parameters

subject

– is the subject header of the message. If you specify a *subject* but not an *originator*, sp_processmail processes all unread messages in the inbox that has the specified subject header. If you specify both *subject* and *originator*, sp_processmail processes all unread messages with the specified subject header sent by the specified originator. If you do not specify either *subject* or *originator*, sp_processmail processes all the unread messages in the Adaptive Server message inbox.

originator

– is the sender of an incoming message. If you specify an *originator* and do not specify a *subject*, sp_processmail processes all unread messages in the inbox sent by the specified originator.

dbuser

– specifies the Adaptive Server login name to use for the user context for executing the query in the message. The default is “guest.”

dbname

– specifies the database name to use for the database context for executing the query in the message. The default is “master.”

filetype

– specifies the file extension of the attached file that contains the results of the query. The default is “.txt”.

separator

– specifies the character to use as a column separator in the query results. It is the same as the /s option of isql. The default is the tab character.

Examples

Example 1

```
sp_processmail @subject="SQL REPORT",  
@originator="janet", @dbuser="sa",  
@dbname="salesdb", @filetype="res", @separator=";"
```

Processes all unread messages in the Adaptive Server inbox with the subject header “SQL Report” submitted by mail user “janet”, processes the received queries in the salesdb database as user “sa”, and returns the query results to “janet” in a .res file attached to the mail message. The columns in the returned results are separated by semicolons.

Example 2

```
sp_processmail @dbuser="sa"
```

Processes all unread messages in the Adaptive Server inbox as user “sa” in the master database and returns the query results in .txt files, which are attached to the mail messages. The columns in the returned results are separated by tab characters.

Usage

- *sp_processmail* reads, processes, sends, and deletes messages in the Adaptive Server message inbox, using the *xp_findnextmsg*, *xp_readmail*, *xp_sendmail*, and *xp_deletemail* system ESPs.
- *sp_processmail* sends outgoing mail to the originator of the incoming mail message being processed.
- *sp_processmail* uses the default parameters when invoking the ESPs, except for the *dbuser*, *dbname*, *attachname*, and *separator* parameters to *xp_sendmail*, which can be overridden by the parameters to *sp_processmail*.
- *sp_processmail* processes all messages as Adaptive Server queries. It reads messages from the Adaptive Server inbox and returns query results to the sender of the message and all its cc'd and bcc'd recipients in an attachment to an Adaptive Server message. *sp_processmail* generates a name for the attached file consisting of “syb” followed by five random digits, followed by the extension specified by the *filetype* parameter; for example, “syb84840.txt.”
- *sp_processmail* deletes messages from the inbox after processing them.

- The *subject* and *originator* parameters specify which messages should be processed. If neither of these parameters is supplied, sp_processmail processes all the unread messages in the Adaptive Server message inbox.
- sp_processmail does not process attachments to incoming mail. The query must be in the body of the incoming message.

Permissions

Only a System Administrator can execute sp_processmail.

See also

Extended stored procedures – xp_deletemail, xp_findnextmsg, xp_readmail, xp_sendmail, xp_startmail

Utility – isql

sp_procqmode

Description

Displays the query processing mode of a stored procedure, view, or trigger.

Syntax

sp_procqmode [*object_name* [, detail]]

Parameters

object_name

– is the name of the stored procedure, view, or trigger whose query processing mode you are examining. If you do not specify an *object_name*, sp_procqmode reports on all procedures, views, and triggers in the current database.

detail

– returns information about whether the object contains a subquery, and whether there is information about the object in syscomments.

Examples

Example 1

```
sp_procqmode
```

Object Owner.name	Object Type	Processing Mode
dbo.au_info	stored procedure	pre-System 11
dbo.titleview	view	System 11 or later

Displays the query processing mode for all stored procedures in the current database.

Example 2

```
sp_procqmode old_sproc, detail
```


Object	Owner.Name	Object Type	Processing Mode	Subq	Text
dbo.au_info		stored procedure	pre-System 11	no	yes

Displays the query processing mode of the stored procedure `old_sproc`, reports whether `old_sproc` contains any subqueries, and reports whether `syscomments` has information about `old_sproc`.

Example 3

```
sp_procqmode null, detail
```

Displays detailed reports for all objects in the database.

Usage

- The processing mode identifies whether the object was created in SQL Server release 10.0 or earlier. Objects created on release 10.x (or earlier) servers are “pre-System 11” objects. Objects created on release 11.0 or later servers are “System 11 or later” objects.
- Subqueries in “pre-System 11” objects use a different processing mode than subqueries in “System 11 or later” objects. Upgrading to release 11.0 or later does not automatically change the processing mode of the subquery.

In general, the “System 11 or later” processing mode is faster than “pre-System 11” processing mode. To change the processing mode to “System 11 or later”, drop and re-create the object. You cannot create an object with “pre-System 11” processing on the current release of Adaptive Server, so you may want to create the object with another name and test it before dropping the version that uses “pre-System 11” processing mode.

- The processing mode displayed for a given object is independent of whether that object actually includes a subquery, and pertains only to the specified object, not to any dependent objects. You must check each object separately.
- The detailed report shows if the object contains a subquery, and reports if text is available in `syscomments` (for `sp_helptext` to report, or for the `defncopy` utility to copy out). `sp_procqmode` does not check that the text in `syscomments` is valid or complete.

Permissions

Only the Database Owner or object owner can execute `sp_procqmode`.

See also

Stored Procedures – `sp_helptext`
Utilities – `defncopy`

sp_procxmode

Description Displays or changes the transaction modes associated with stored procedures.

Syntax sp_procxmode [*procname* [, *tranmode*]]

Parameters

- procname*
– is the name of the stored procedure whose transaction mode you are examining or changing.
- tranmode*
– is the new transaction mode for the stored procedure. Values are "chained", "unchained", and "anymode".

Examples

Example 1

```
sp_procxmode

procedure name      user name      transaction mode
-----
byroyalty           dbo            Unchained
discount_proc      dbo            Unchained
history_proc        dbo            Unchained
insert_sales_proc   dbo            Unchained
insert_detail_proc  dbo            Unchained
storeid_proc        dbo            Unchained
storename_proc      dbo            Unchained
title_proc          dbo            Unchained
titleid_proc        dbo            Unchained
```

Displays the transaction mode for all stored procedures in the current database.

Example 2

```
sp_procxmode byroyalty

procedure name      transaction mode
-----
byroyalty           Unchained
```

Displays the transaction mode of the stored procedure byroyalty.

Example 3

```
sp_procxmode byroyalty, "chained"
```

Changes the transaction mode for the stored procedure byroyalty in the pubs2 database from “unchained” to “chained”.

Usage

- To change the transaction mode of a stored procedure, you must be the owner of the stored procedure, the owner of the database containing the stored procedure, or the System Administrator. The Database Owner or System Administrator can change the mode of another user's stored procedure by qualifying it with the database and user name. For example:

```
sp_procxmode "otherdb.otheruser.newproc", "chained"
```

- To use *sp_procxmode*, turn off chained transaction mode using the *chained* option of the *set* command. By default, this option is turned off.
- When you use *sp_procxmode* with no parameters, it reports the transaction modes of every stored procedure in the current database.
- To examine a stored procedure's transaction mode (without changing it), enter:

```
sp_procxmode procname
```

- To change a stored procedure's transaction mode, enter:

```
sp_procxmode procname, tranmode
```

- When you create a stored procedure, Adaptive Server tags it with the current session's transaction mode. This means:
 - You can execute "chained" stored procedures only in sessions using chained transaction mode.
 - You can execute "unchained" stored procedures only in sessions using unchained transaction mode.

To execute a particular stored procedure in either chained or unchained sessions, set its transaction mode to "anymode".

- If you attempt to run a stored procedure under the wrong transaction mode, Adaptive Server returns a warning message, but the current transaction, if any, is not affected.

Permissions

Only a System Administrator, the Database Owner, or the owner of a procedure can execute *sp_procxmode* to change the transaction mode. Any user can execute *sp_procxmode* to display the transaction mode.

See also

Commands – *begin transaction*, *commit*, *save transaction*, *set*

Procedures: *sp_recompile* – *sp_role*

sp_recompile

Description	Causes each stored procedure and trigger that uses the named table to be recompiled the next time it runs.
Syntax	<code>sp_recompile <i>objname</i></code>
Parameters	<i>objname</i> – is the name of a table in the current database.
Examples	<code>sp_recompile titles</code> Recompiles each trigger and stored procedure that uses the table <code>titles</code> the next time the trigger or stored procedure is run.
Usage	<ul style="list-style-type: none">• The queries used by stored procedures and triggers are optimized only once, when they are compiled. As you add indexes or make other changes to your database that affect its statistics, your compiled stored procedures and triggers may lose efficiency. By recompiling the stored procedures and triggers that act on a table, you can optimize the queries for maximum efficiency.• <code>sp_recompile</code> looks for <i>objname</i> only in the current database and recompiles triggers and stored procedures only in the current database. <code>sp_recompile</code> does not affect objects in other databases that depend on the table.• You cannot use <code>sp_recompile</code> on system tables.
Permissions	Any user can execute <code>sp_recompile</code> .
See also	<i>Commands</i> – create index, update statistics

sp_remap

Description	Remaps a stored procedure, trigger, rule, default, or view from releases later than 4.8 and prior to 10.0 to be compatible with releases 10.0 and later. Use sp_remap on pre-existing objects that the upgrade procedure failed to remap.
Syntax	sp_remap <i>objname</i>
Parameters	<i>objname</i> – is the name of a stored procedure, trigger, rule, default, or view in the current database.
Examples	Example 1 <pre>sp_remap myproc</pre> Remaps a stored procedure called myproc. Example 2 <pre>sp_remap "my_db..default_date"</pre> Remaps a rule called default_date. Execute a use my_db statement to open the my_db database before running this procedure.
Usage	<ul style="list-style-type: none">• If sp_remap fails to remap an object, drop the object from the database and re-create it. Before running sp_remap on an object, it is a good idea to copy its definition into an operating system file with the defncopy utility. For more information about defncopy, see the <i>Utilities Guide</i>.• sp_remap can cause your transaction log to fill rapidly. Before running sp_remap, use the dump transaction command to dump the transaction log, as needed.• You can use sp_remap only on objects in the current database.• sp_remap makes no changes to objects that were successfully upgraded to the current release.
Permissions	Only a System Administrator or the owner of an object can execute sp_remap.
See also	<i>Commands</i> – dump transaction <i>System procedures</i> – sp_helptext <i>Utility programs</i> – defncopy

sp_remoteoption

Description	Displays or changes remote login options.
Syntax	<code>sp_remoteoption [remoteserver [, loginame [, remotename [, optname [, optvalue]]]]]</code>
Parameters	<i>remoteserver</i> – is the name of the server that will be executing RPCs on this server.

Note This manual page uses the term "local server" to refer to the server that is executing the remote procedures that are run from a "remote server".

loginame

– is the login name that identifies the local login for the *remoteserver*, *loginame*, *remotename* combination.

remotename

– is the remote user name that identifies the remote login for the *remoteserver*, *loginame*, *remotename* combination.

optname

– is the name of the option to change. Currently, there is only one option, *trusted*, which means that the local server accepts remote logins from other servers without user-access verification for the particular remote login. The default is to use password verification. Adaptive Server understands any unique string that is part of the option name. Use quotes around the option name if it includes embedded blanks.

optvalue

– is either true or false. true turns the option on, false turns it off.

Examples

Example 1

```
sp_remoteoption
Settable remote login options.
remotelogin_option
-----
trusted
```

Displays a list of the remote login options.

Example 2

```
sp_remoteoption GATEWAY, churchy, pogo, trusted,
true
```

Defines the remote login from the remote server GATEWAY to be trusted (that is, the password is not checked).

Example 3

```
sp_remotoption GATEWAY, churchy, pogo, trusted,  
false
```

Defines the remote login “pogo” from the remote server GATEWAY as a login that is not trusted (that is, the password is checked).

Example 4

```
sp_remotoption GATEWAY, albert, NULL, trusted, true
```

Defines all logins from GATEWAY that map to login “albert” on the local server to be trusted.

Usage

- To display a list of the remote login options, execute sp_remotoption with no parameters.
- If you have used sp_addremotlogin to map all users from a remote server to the same local name, specify trusted for those users. For example, if all users from server GOODSRV that are mapped to “albert” are trusted, specify:

```
sp_remotoption GOODSRV, albert, NULL, trusted,  
true
```

If the logins are not specified as trusted, they cannot execute RPCs on the local server unless they specify local server passwords when they log into the remote server. When they use Open Client Client-Library, users can specify a password for server-to-server connections with the routine ct_remote_pwd. isql and bcp do not permit users to specify a password for RPC connections.

If users are logged into the remote server using “unified login”, the logins must also be trusted on the local server, or they must specify passwords for the server when they log into the remote server.

For more information about setting up servers for remote procedure calls and for using “unified login”, see the System Administration Guide.

Permissions

Only a System Security Officer can execute sp_remotoption.

See also

System procedures – sp_addremotlogin, sp_droptremotlogin, sp_helpremotlogin

Utility – isql

sp_remotesql

Description	<i>Component Integration Services only</i> – Establishes a connection to a remote server, passes a query buffer to the remote server from the client, and relays the results back to the client.
Syntax	<code>sp_remotesql server, query [, query2, ... , query254]</code>
Parameters	<p><i>server_name</i> – is the name of a remote server defined with <i>sp_addserver</i>.</p> <p><i>query</i> – is a query buffer a with maximum length of 255 characters.</p> <p><i>query2</i> through <i>query254</i> – is a query buffer with a maximum length of 255 characters. If supplied, these arguments are concatenated with the contents of <i>query1</i> into a single query buffer.</p>
Examples	<p>Example 1</p> <pre>sp_remotesql FREDS_SERVER, "select @@version"</pre> <p>Passes the query buffer to FREDS_SERVER, which interprets select @@version and returns the result to the client. Adaptive Server does not interpret the result.</p> <p>Example 2</p> <pre>create procedure freds_version as exec sp_remotesql FREDS_SERVER, "select @@version" go exec freds_version go</pre> <p>Illustrates the use of <i>sp_remotesql</i> in a stored procedure. This example and example 1 return the same information to the client.</p> <p>Example 3</p> <pre>sp_remotesql DCO_SERVER, "insert into remote_table (numbercol,intcol, floatcol,datecol)", "values (109.26,75, 100E5,'10-AUG-85')"</pre>

```
select @@error
```

The server concatenates two query buffers into a single buffer, and passes the complete insert statement to the server DCO_SERVER for processing. The syntax for the insert statement is a format that DCO_SERVER understands. The returned information is not interpreted by the server. This example also examines the value returned in @@error.

Example 4

```
declare @servname varchar(30)
declare @querybuf varchar(200)
select @servname = "DCO_SERV"
select @querybuf = "select table_name
                  from all_tables
                  where owner = 'SYS'"
exec sp_remotesql @servname, @querybuf
```

Illustrates the use of local variables as parameters to sp_remotesql.

Usage

- sp_remotesql establishes a connection to a remote server, passes a query buffer to the remote server from the client, and relays the results back to the client. The local server does not intercept results.
- You can use sp_remotesql within another stored procedure.
- The query buffer parameters must be a character expression with a maximum length of 255 characters. If you use a query buffer that is not char or varchar, you will receive datatype conversion errors.
- sp_remotesql sets the global variable @@error to the value of the last error message returned from the remote server if the severity of the message is greater than 10.
- If sp_remotesql is issued from within a transaction, Adaptive Server verifies that a transaction has been started on the remote server before passing the query buffer for execution. When the transaction terminates, the remote server is directed to commit the transaction. The work performed by the contents of the query buffer is part of the unit of work defined by the transaction.

If transaction control statements are part of the query buffer, it is the responsibility of the client to ensure that the transaction commit and rollback occur as expected. Mixing Transact-SQL with transaction control commands in the query buffer can cause unpredictable results.

- The local server manages the connection to the remote server. Embedding connect to or disconnect commands in the query buffer causes results that require interpretation by the remote server. This is not required or recommended. Typically, the result is a syntax error.

Permissions	Any user can execute <i>sp_remoteseql</i> .
See also	<i>Commands</i> – connect to...disconnect <i>System procedures</i> – <i>sp_autoconnect</i> , <i>sp_passthru</i>

sp_rename

Description	Changes the name of a user-created object or user-defined datatype in the current database.
Syntax	<code>sp_rename objname, newname</code>
Parameters	<i>objname</i> – is the original name of the user-created object (table, view, column, stored procedure, index, trigger, default, rule, check constraint, referential constraint, or user-defined datatype). If the object to be renamed is a column in a table, <i>objname</i> must be in the form “ <i>table.column</i> ”. If the object is an index, <i>objname</i> must be in the form “ <i>table.indexname</i> ”. <i>newname</i> – is the new name of the object or datatype. The name must conform to the rules for identifiers and must be unique to the current database.
Examples	Example 1 <pre>sp_rename titles, books</pre> Renames the titles table to books. Example 2 <pre>sp_rename "books.title", bookname</pre> Renames the title column in the books table to bookname. Example 3 <pre>sp_rename "books.titleind", titleindex</pre> Renames the titleind index in the books table to titleindex. Example 4 <pre>sp_rename tid, bookid</pre> Renames the user-defined datatype tid to bookid.

Usage

- sp_rename changes the name of a user-created object or datatype. You can change only the name of an object or datatype in the database in which you issue sp_rename.
- When you are renaming a column or index, do not specify the table name in the *newname*. See examples 2 and 3.
- If a column and an index have the same name, sp_rename changes the column name. The syntax does not allow you to specify that the index should be changed instead of the column. In this case, to change the index name, drop the index and recreate it.
- You can change the name of an object referenced by a view. For example, if a view references the new_sales table and you rename new_sales to old_sales, the view will reference old_sales.
- You cannot change the names of system objects and system datatypes.

Warning! Procedures, triggers, and views that depend on an object whose name has been changed work until they are dropped and re-created. Also, the old object name appears in query results until the user changes and re-creates the procedure, trigger, or view. Change the definitions of any dependent objects when you execute sp_rename. Find dependent objects with sp_depends.

Permissions

Only the Database Owner or a System Administrator can use the setuser command to assume another database user's identity to rename objects owned by other users. All users can execute sp_rename to rename their own objects.

See also

System procedures – sp_depends, sp_rename

sp_renamedb

Description

Changes the name of a user database.

Syntax

sp_renamedb *dbname*, *newname*

Parameters

dbname

– is the original name of the database.

newname

– is the new name of the database. Database names must conform to the rules for identifiers and must be unique.

Examples

Example 1

```
sp_renamedb accounting, financial
```

Renames the accounting database to financial.

Example 2

```
sp_dboption work, single, true
go
use work
go
checkpoint
go
sp_renamedb work, workdb
go
use master
go
sp_dboption workdb, single, false
go
use workdb
go
checkpoint
go
```

Renames the database named `work`, which is a Transact-SQL reserved word, to `workdb`. This example shows how `sp_dboption` is used to place the `work` database in single-user mode before renaming it and restore it to multi-user mode afterward.

Usage

- `sp_renamedb` changes the name of a database. You *cannot* rename system databases or databases with external referential integrity constraints.
- The System Administrator must place a database in single-user mode with `sp_dboption` before renaming it and must restore it to multi-user mode afterward.
- `sp_renamedb` fails if any table in the database references, or is referenced by, a table in another database. Use the following query to determine which tables and external databases have foreign key constraints on primary key tables in the current database:

```
select object_name(tableid), db_name(frgndbid)
from sysreferences
where frgndbid is not null
```

Use the following query to determine which tables and external databases have primary key constraints for foreign key tables in the current database:

```
select object_name(reftabid), db_name(pmrydbid)
from sysreferences
where pmrydbid is not null
```

Use alter table to drop the cross-database constraints in these tables. Then, rerun sp_renamedb.

- When you change a database name:
 - Drop all stored procedures, triggers, and views that include the database name
 - Change the source text of the dropped objects to reflect the new database name
 - Re-create the dropped objects
 - Change all applications and SQL source scripts that reference the database, either in a use *database_name* command or as part of a fully qualified identifier (in the form *dbname.[owner].objectname*).
- If you use scripts to run dbcc commands or dump database and dump transaction commands on your databases, be sure to update those scripts.

Warning! Procedures, triggers, and views that depend on a database whose name has been changed work until they are re-created. Change the definitions of any dependent objects when you execute sp_renamedb. Find dependent objects with sp_depends.

Permissions

Only a System Administrator can execute sp_renamedb.

See also

Commands – create database

System procedures – sp_changedbowner, sp_dboption, sp_depends, sp_helpdb, sp_rename

sp_rename_qpgroup

Description	Renames an abstract plan group.
Syntax	<code>sp_rename_qpgroup old_name, new_name</code>
Parameters	<p><i>old_name</i></p> <ul style="list-style-type: none"> – is the current name of the abstract plan group. <p><i>new_name</i></p> <ul style="list-style-type: none"> – is the new name for the group. The specified <i>new_name</i> cannot be the name of an existing abstract plan group in the database.
Examples	<pre>sp_rename_qpgroup dev_plans, prod_plans</pre> <p>Changes the name of the group from <code>dev_plans</code> to <code>prod_plans</code>.</p>
Usage	<ul style="list-style-type: none"> • Use <code>sp_rename_qpgroup</code> to rename an abstract plan group. You cannot use the name of an existing plan group for the new name. • <code>sp_rename_qpgroup</code> does not affect the contents of the renamed group. IDs of existing abstract plans are not changed. • You cannot rename the default abstract plan groups, <code>ap_stdin</code> and <code>ap_stdout</code>. • <code>sp_rename_qpgroup</code> cannot be run in a transaction.
Permissions	Only a System Administrator or the Database Owner can execute <code>sp_rename_qpgroup</code> .
See also	<i>System procedures</i> – <code>sp_help_qpgroup</code>

sp_reportstats

Description	Reports statistics on system usage.
Syntax	<code>sp_reportstats [loginame]</code>
Parameters	<p><i>loginame</i></p> <ul style="list-style-type: none"> – is the login name of the user to show accounting totals for.

Examples

Example 1

```
sp_reportstats
```

Name	Since	CPU	Percent CPU	I/O	Percent I/O
-----	-----	----	-----	-----	-----
julie	jun 19 1993	10000	24.9962%	5000	24.325%

jason	jun 19 1993	10002	25.0013%	5321	25.8866%
ken	jun 19 1993	10001	24.9987%	5123	24.9234%
kathy	jun 19 1993	10003	25.0038%	5111	24.865%

```
      Total CPU    Total I/O
      -----    -
      40006        20555
```

Displays a report of current accounting totals for all Adaptive Server users.

Example 2

```
      sp_reportstats kathy
```

Name	Since	CPU	Percent CPU	I/O	Percent I/O
-----	-----	----	-----	-----	-----
kathy	Jul 24 1993	498	49.8998%	48392	9.1829%

```
      Total CPU    Total I/O
      -----    -
      998          98392
```

Displays a report of current accounting totals for user “kathy.”

Usage

- sp_reportstats prints out the current accounting totals for all logins, as well as each login’s individual statistics and percentage of the overall statistics. sp_reportstats accepts one parameter, the login name of the account to report. With no parameters, sp_reportstats reports on all accounts.
- sp_reportstats does not report statistics for any process with a system user ID (suid) of 0 or 1. This includes deadlock detection, checkpoint, housekeeper, network, auditing, mirror handlers, and all users with sa_role.
- The units reported for “CPU” are *machine* clock ticks, not Adaptive Server clock ticks.
- The “probe” user exists for the two-phase commit probe process, which uses a challenge-and-response mechanism to access Adaptive Server.

Permissions

Only a System Administrator can execute sp_reportstats.

See also

System procedures – sp_clearstats, sp_configure

sp_revokelogin

Description	<i>Windows NT only</i> – Revokes Adaptive Server roles and default permissions from Windows NT users and groups when Integrated Security mode or Mixed mode (with Named Pipes) is active.
Syntax	<code>sp_revokelogin {login_name group_name}</code>
Parameters	<p><i>login_name</i> – is the network login name of the Windows NT user.</p> <p><i>group_name</i> – is the Windows NT group name.</p>
Examples	<p>Example 1</p> <pre>sp_revokelogin jeanluc</pre> <p>Revokes all permissions from the Windows NT user named “jeanluc”.</p> <p>Example 2</p> <pre>sp_revokelogin Administrators</pre> <p>Revokes all roles from the Windows NT Administrators group.</p>
Usage	<ul style="list-style-type: none"> • Use <code>sp_revokelogin</code> only when Adaptive Server is running in Integrated Security mode or Mixed mode, when the connection is Named Pipes. If Adaptive Server is running in Standard mode, or in Mixed mode using a connection other than Named Pipes, use the <code>revoke</code> command. • If you revoke a user’s roles and default privileges with <code>sp_revokelogin</code>, that user can no longer log into Adaptive Server over a trusted connection.
Permissions	Only a System Administrator can execute <code>sp_revokelogin</code> .
See also	<p><i>Commands</i> – <code>grant</code>, <code>revoke</code>, <code>setuser</code></p> <p><i>System procedures</i> – <code>sp_droplogin</code>, <code>sp_dropuser</code>, <code>sp_logininfo</code></p>

sp_role

Description	Grants or revokes roles to an Adaptive Server login account.
Syntax	<code>sp_role {"grant" "revoke"}, rolename, loginame</code>

Parameters	<p>grant revoke</p> <ul style="list-style-type: none">– specifies whether to grant the role to or revoke the role from <i>loginame</i>. <p><i>rolename</i></p> <ul style="list-style-type: none">– is the role to be granted or revoked. <p><i>loginame</i></p> <ul style="list-style-type: none">– is the login account to or from which the role is to be granted or revoked.
Examples	<pre>sp_role "grant", sa_role, alexander</pre> <p>Grants the System Administrator role to the login account named “alexander”.</p>
Usage	<ul style="list-style-type: none">• sp_role grants or revokes roles to an Adaptive Server login account.• When you grant a role to a user, it takes effect the next time the user logs into Adaptive Server. Alternatively, the user can enable the role immediately by using the set role command. For example, the command: <pre>set role sa_role on</pre>enables the System Administrator role for the user.• You cannot revoke a role from a user while the user is logged in.• When users log in, all roles that have been granted to them are active (on). To turn a role off, use the set command. For example, to deactivate the System Administrator role, use the command: <pre>set role "sa_role" off</pre>
Permissions	<p>Only a System Administrator can execute sp_role to grant the System Administrator role to other users. Only a System Security Officer can execute sp_role to grant any role other than “sa” to other users.</p>
See also	<p><i>Commands</i> – grant, revoke, set</p> <p><i>Functions</i> – proc_role</p> <p><i>System procedures</i> – sp_displaylogin</p>

Procedures: *sp_sendmsg* – *sp_transactions*

sp_sendmsg

Description	Sends a message to a User Datagram Protocol (UDP) port.
Syntax	<code>sp_sendmsg ip_address, port_number, message</code>
Parameters	<p><i>ip_address</i></p> <ul style="list-style-type: none"> – is the IP address of the machine where the UDP application is running. <p><i>port_number</i></p> <ul style="list-style-type: none"> – is the port number of the UDP port. <p><i>message</i></p> <ul style="list-style-type: none"> – is the message to send. It can be up to 255 characters in length.
Examples	<code>sp_sendmsg "120.10.20.5", 3456, "Hello World"</code>
Usage	<ul style="list-style-type: none"> • <code>sp_sendmsg</code> is not supported on Windows NT. • To enable the use of UDP messaging, a System Security Officer must set the configuration parameter <code>allow_sendmsg</code> to 1. • No security checks are performed with <code>sp_sendmsg</code>. Sybase strongly recommends caution when using <code>sp_sendmsg</code> to send sensitive information across the network. By enabling this functionality, the user accepts any security problems which result from its use. • This sample C program listens on a port that you specify and echoes the messages it receives. For example, to receive the <code>sp_sendmsg</code> calls for example 1, use:

```

updmon 3456

#include <stdlib.h>
#include <stdio.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
#include <unistd.h>

```

```
#include <fcntl.h>

main(argc, argv)
int argc; char *argv[];
{
    struct sockaddr_in sadr;
    int portnum,sck,dummy,msglen;
    char msg[256];

    if (argc < 2) {
        printf("Usage: udpmon <udp portnum>\n");
        exit(1);
    }

    if ((portnum=atoi(argv[1])) < 1) {
        printf("Invalid udp portnum\n");
        exit(1);
    }

    if ((sck=socket(AF_INET,SOCK_DGRAM,IPPROTO_UDP)) < 0) {
        printf("Couldn't create socket\n");
        exit(1);
    }

    sadr.sin_family = AF_INET;
    sadr.sin_addr.s_addr = inet_addr("0.0.0.0");
    sadr.sin_port = portnum;

    if (bind(sck,&sadr,sizeof(sadr)) < 0) {
        printf("Couldn't bind requested udp port\n");
        exit(1);
    }

    for (;;)
    {
        if((msglen=recvfrom(sck,msg,sizeof(msg),0,NULL,&dummy))
< 0)
            printf("Couldn't recvfrom() from udp port\n");
        printf("%.*s\n", msglen, msg);
    }
}
```

Permissions Any user can execute sp_sendmsg.

See also *Function* – syb_sendmsg

sp_serveroption

Description	Displays or changes remote server options.
Syntax	<code>sp_serveroption [server, optname, optvalue]</code>
Parameters	<p><i>server</i></p> <ul style="list-style-type: none"> – is the name of the remote server for which to set the option. <p><i>optname</i></p> <ul style="list-style-type: none"> – is the name of the option to be set or unset. Table 23-1 lists the option names.

Table 23-1: *sp_serveroption* options

Option	Meaning
mutual authentication	<i>This option is valid for “rpc security model B” only.</i> This option specifies that 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.
external engine auto start	Specifies that EJB Server starts up each time Adaptive Server starts up. The default is true; starting Adaptive Server also starts up EJB Server.
net password encryption	Specifies whether to initiate connections with a remote server with the client side password encryption handshake or with the normal (unencrypted password) handshake sequence. The default is false, no network encryption.
readonly	Specifies that access to the server named is read only. This option is available only with Component Integration Services.
rpc security model A	The default model for handling RPCs. This model does not support mutual authentication, message integrity, or message confidentiality between the local server and the remote server.
rpc security model B	This model results in a single, secure physical connection established between the local and remote servers. Logical connections for each RPC that is executed are multiplexed over the single, secure, physical connection. This model supports mutual authentication, message confidentiality via encryption, and message integrity.
security mechanism	<i>This option is valid for “rpc security model B” only.</i> This option specifies the security mechanism for the remote server. You must set this option to true to use security model B.
timeouts	When unset (false), disables the normal timeout code used by the local server, so the site connection handler does not automatically drop the physical connection after one minute with no logical connection. The default is true.
use message confidentiality	<i>This option is valid for “rpc security model B” only.</i> This option specifies that messages are encrypted when sent to the remote server, and results from the remote server are encrypted.
use message integrity	<i>This option is valid for “rpc security model B” only.</i> This option specifies that messages between the servers are checked for tampering.

Adaptive Server accepts any unique string that is part of the option name. Use quotes around the option name if it includes embedded blanks.

optvalue

– is true (on) or false (off) for all options except the security mechanism option.

For the security mechanism option, specify the name of the security mechanism. To see the names of the security mechanisms available on a server, execute:

```
select * from syssecmechs
```

Examples

Example 1

```
sp_serveroption
Settable server options.
```

```
-----
mutual authentication
net password encryption
readonly
rpc security model A
rpc security model B
security mechanism
timeouts
use message confidentiality
use message integrity
timeouts
net password encryption
```

Displays a list of the server options.

Example 2

```
sp_serveroption GATEWAY, "timeouts", false
```

Tells the server not to time out inactive physical connections with the remote server GATEWAY.

Example 3

```
sp_serveroption GATEWAY,
"net password encryption", true
```

Specifies that when connecting to the remote server GATEWAY, GATEWAY sends back an encryption key to encrypt the password to send to it.

Example 4

```
sp_serveroption SYB_EJB,  
                "external engine auto start", true
```

Specifies that the EJB Server SYB_EJB starts up each time Adaptive Server starts up.

Example 5

```
sp_serveroption TEST3, "rpc security model B", true
```

Specifies that the security model for RPCs for the server “TEST3” is security model B.

Example 6

```
sp_serveroption TEST3, "security mechanism", dce
```

Specifies that the security model to use for RPCs for “TEST3” is DCE.

Example 7

```
sp_serveroption TEST3, "mutual authentication", true
```

Specifies that the local server will check the authenticity of the remote server “TEST3”. With security model B, the remote server will check the authenticity of the local server, whether or not this option is set.

Usage

- To display a list of server options that can be set by the user, use *sp_serveroption* with no parameters.
- Once *timeouts* is set to *false*, the site handlers will continue to run until one of the two servers is shut down.
- The *net password encryption* option allows clients to specify whether to send passwords in plain text or encrypted form over the network when initiating a remote procedure call. If *net password encryption* is *true*, the initial login packet is sent without passwords, and the client indicates to the remote server that encryption is desired. The remote server sends back an encryption key, which the client uses to encrypt its passwords. The client then encrypts its passwords, and the remote server uses the key to authenticate them when they arrive.
- To set network password encryption for a particular *isql* session, you can use a command line option for *isql*. For more information, see the *Utility Programs manual* for your platform.
- You cannot use the *net password encryption* option when connecting to a pre-release 10.0 SQL Server.

- The options security mechanism, mutual authentication, use message confidentiality, and use message integrity do *not* apply to security model A.
- To use security model B, both the local server and the remote server must use model B and both must use the same security mechanism.
- For more information on server options, see the System Administration Guide.

Permissions Only a System Administrator can execute sp_serveroption to set the timeouts option. Any user can execute sp_serveroption with no parameters to display a list of options.

Only a System Security Officer can set the net password encryption, security mechanism, mutual authentication, use message confidentiality, and use message integrity options.

See also *System procedures* – sp_helpserver, sp_password
Utility – isql

sp_setlangalias

Description Assigns or changes the alias for an alternate language.

Syntax sp_setlangalias *language*, *alias*

Parameters *language*
– is the official language name of the alternate language.
alias
– is the new local alias for the alternate language.

Examples sp_setlangalias french, français

This command assigns the alias name “français” for the official language name “french”.

Usage

- *alias* replaces the current value of syslanguages.alias for the official name.
- The set language command can use the new *alias* in place of the official language name.

Permissions Only a System Administrator can execute sp_setlangalias.

See also *Commands* – set
System procedures – *sp_addlanguage*, *sp_droplanguage*, *sp_helplanguage*

sp_setpglockpromote

Description Sets or changes the lock promotion thresholds for a database, for a table, or for Adaptive Server.

Syntax `sp_setpglockpromote {"database" | "table"}, objname, new_lwm, new_hwm, new_pct`
`sp_setpglockpromote server, NULL, new_lwm, new_hwm, new_pct`

Parameters

- `server`
 - sets server-wide values for the lock promotion thresholds.
- `"database" | "table"`
 - specifies whether to set the lock promotion thresholds for a database or table. “database” and “table” are Transact-SQL keywords, so the quotes are required.
- `objname`
 - is either the name of the table or database for which you are setting the lock promotion thresholds or null, if you are setting server-wide values.
- `new_lwm`
 - specifies the value to set for the low watermark (LWM) threshold. The LWM must be less than or equal to the high watermark (HWM). The minimum value for LWM is 2. This parameter can be null.
- `new_hwm`
 - specifies the value to set for the lock promotion HWM threshold. The HWM must be greater than or equal to the LWM. The maximum HWM is 2,147,483,647. This parameter can be null.
- `new_pct`
 - specifies the value to set for the lock promotion percentage (PCT) threshold. PCT must be between 1 and 100. This parameter can be null.

Examples

Example 1

```
sp_setpglockpromote "server", NULL, 200, 300, 50
```

Sets the server-wide lock promotion LWM to 200, the HWM to 300, and the PCT to 50.

Example 2

```
sp_setpglockpromote "database", master, 1000, 1100, 45
```

Sets lock promotion thresholds for the master database.

Example 3

```
sp_setpglockpromote "table", "pubs2..titles", 500, 700, 10
```

Sets lock promotion thresholds for the titles table in the pubs2 database. This command must be issued from the pubs2 database.

Example 4

```
sp_setpglockpromote "database", master, @new_hwm=1600
```

Changes the HWM threshold to 1600 for the master database. The thresholds were previously set with sp_setpglockpromote. This command must be issued from the master database.

Usage

- sp_setpglockpromote configures the lock promotion values for a table, for a database, or for Adaptive Server.

Adaptive Server acquires page locks on a table until the number of locks exceeds the lock promotion threshold. sp_setpglockpromote changes the lock promotion thresholds for an object, a database, or the server. If Adaptive Server is successful in acquiring a table lock, the page locks are released.

When the number of locks on a table exceeds the HWM threshold, Adaptive Server attempts to escalate to a table lock. When the number of locks on a table is below the LWM, Adaptive Server does not attempt to escalate to a table lock. When the number of locks on a table is between the HWM and LWM and the number of locks exceeds the PCT threshold, Adaptive Server attempts to escalate to a table lock.

- Lock promotion thresholds for a table override the database or server-wide settings. Lock promotion thresholds for a database override the server-wide settings.

- Lock promotion thresholds for Adaptive Server do not need initialization, but you must initialize database and table lock promotion thresholds by specifying LWM, HWM, and PCT with *sp_setpglockpromote*, which creates a row for the object in *sysattributes* when it is first run for a database or table. Once the thresholds have been initialized, then they can be modified individually, as in example 4.
- For a table or a database, *sp_setpglockpromote* sets LWM, HWM, and PCT in a single transaction. If *sp_setpglockpromote* encounters an error while updating any of the values, then all changes are aborted and the transaction is rolled back. For server-wide changes, one or more thresholds may fail to be updated while others are successfully updated. Adaptive Server returns an error message if any values fail to be updated.
- To view the server-wide settings for the lock promotion thresholds, use *sp_configure* "lock promotion" to see all three threshold values. To view lock promotion settings for a database, use *sp_helpdb*. To view lock promotion settings for a table, use *sp_help*.

Permissions

Only a System Administrator can execute *sp_setpglockpromote*.

See also

System procedures – *sp_configure*, *sp_droppglockpromote*, *sp_helpdb*

sp_setpsex

Description

Sets custom execution attributes for a session while the session is active.

Syntax

sp_setpsex *spid*, *exeattr*, *value*

Parameters

spid

– is the ID of the session for which to set execution variables. Use *sp_who* to see *spids*.

exeattr

– identifies the execution attribute to be set. Values are priority and enginegroup.

value

– is the new value of *exeattr*. Values for each attribute are as follows:

- If *exeattr* is *priority*, *value* is HIGH, MEDIUM, or LOW.
- If *exeattr* is *enginegroup*, *value* is the name of an existing engine group.

Examples

```
sp_setpsexex 1, "priority", "HIGH"
```

This statement sets the priority of the process with an ID of 1 to HIGH.

Usage

- Execution attribute values specified with `sp_setpsexex` are valid for the current session only and do not apply after the session terminates.
- Use `sp_setpsexex` with caution or it can result in degraded performance. Changing attributes “on the fly”, using `sp_setpsexex`, can help if the process is not getting CPU time; however, if the performance problem is due to something else, such as locks, changing execution attributes could make the problem worse.
- Because you can only set execution attributes for sessions, `sp_setpsexex` cannot be set for a worker process `spid`.
- Except for the housekeeper `spid`, you cannot set execution attributes for system `spids`.
- `sp_setpsexex` does not work if there are no online engines in the associated engine group.

Permissions

Only a System Administrator can execute `sp_setpsexex` without restriction. Any user can execute `sp_setpsexex` to lower the priority of a process owned by that user.

See also

System procedures – `sp_addexclass`, `sp_bindexclass`, `sp_dropexclass`, `sp_showexclass`

sp_set_qplan

Description

Changes the text of the abstract plan of an existing plan without changing the associated query.

Syntax

```
sp_set_qplan id, plan
```

Parameters

id

– is the ID of the abstract plan.

	<i>plan</i>
	– is a new abstract plan.
Examples	<pre>sp_set_qplan 563789159, "(g_join (scan t1) (scan t2))"</pre>
Usage	<ul style="list-style-type: none"> • Use <i>sp_set_qplan</i> to change the abstract plan of an existing plan. You can specify a maximum of 255 characters for a plan. If the abstract plan is longer than 255 characters, you can drop the old plan with <i>sp_drop_qplan</i> and then use <i>create plan</i> to create a new plan for the query. • When you change a plan with <i>sp_set_qplan</i>, plans are not checked for valid abstract plan syntax. Also, the plan is not checked for compatibility with the SQL text. All plans modified with <i>sp_set_qplan</i> should be immediately checked for correctness by running the query for the specified ID. • To find the ID of a plan, use <i>sp_help_qpgroup</i>, <i>sp_help_qplan</i>, or <i>sp_find_qplan</i>. Plan IDs are also returned by <i>create plan</i> and are included in <i>showplan</i> output.
Permissions	Any user can execute <i>sp_set_qplan</i> to change the text for a plan that he or she owns. Only the System Administrator or the Database Owner can change the text for a plan that belongs to another user.
See also	<i>Commands</i> – <i>create plan</i> <i>System procedures</i> – <i>sp_find_qplan</i> , <i>sp_help_qplan</i>

sp_setrowlockpromote

Description	Sets or changes row-lock promotion thresholds for a datarows-locked table, for all datarows-locked tables in a database, or for all datarows-locked tables on a server.
Syntax	<pre>sp_setrowlockpromote "server", NULL, new_lwm, new_hwm, new_pct sp_setrowlockpromote {"database" "table"}, objname, new_lwm, new_hwm, new_pct</pre>
Parameters	<p><i>server</i></p> <p>– sets server-wide values for the row lock promotion thresholds.</p>

"database"

| "table" – specifies whether to set the row-lock promotion thresholds for a database or table.

objname

– is either the name of the table or database for which you are setting the row-lock promotion thresholds or null, if you are setting server-wide values.

new_lwm

– specifies the value to set for the low watermark (LWM) threshold. The LWM must be less than or equal to the high watermark (HWM). The minimum value for LWM is 2. This parameter can be null.

new_hwm

– specifies the value to set for the high watermark (HWM) threshold. The HWM must be greater than or equal to the LWM. The maximum HWM is 2,147,483,647. This parameter can be null.

new_pct

– specifies the value to set for the lock promotion percentage (PCT) threshold. PCT must be between 1 and 100. This parameter can be null.

Examples

Example 1

```
sp_setrowlockpromote "database", engdb, 400, 400,95
```

Sets row lock promotion values for all datarows-locked tables in the engdb database.

Example 2

```
sp_setrowlockpromote "table", sales, 250, 250, 100
```

Sets row lock promotion values for the sales table.

Usage

- sp_setrowlockpromote sets or changes row-lock promotion thresholds for a table, a database, or Adaptive Server.

Adaptive Server acquires row locks on a datarows-locked table until the number of locks exceeds the lock promotion threshold. If Adaptive Server is successful in acquiring a table lock, the row locks are released.

When the number of row locks on a table exceeds the HWM, Adaptive Server attempts to escalate to a table lock. When the number of row locks on a table is below the LWM, Adaptive Server does not attempt to escalate to a table lock. When the number of row locks on a table is between the HWM and LWM, and the number of row locks exceeds the PCT threshold as a percentage of the number of rows in a table, Adaptive Server attempts to escalate to a table lock.

- Lock promotion is always two-tiered, that is, row locks are promoted to table locks. Adaptive Server does not promote from row locks to page locks.
- Lock promotion thresholds for a table override the database or server-wide settings. Lock promotion thresholds for a database override the server-wide settings.
- To change the lock promotion thresholds for a database, you must be using the master database. To change the lock promotion thresholds for a table in a database, you must be using the database where the table resides.
- Server-wide row lock promotion thresholds can also be set with *sp_configure*. When you use *sp_setrowlockpromote* to change the values server-wide, it changes the configuration parameters, and saves the configuration file. When you first install Adaptive Server, the server-wide row lock promotion thresholds set by the configuration parameters are:

row lock promotion HWM	200
row lock promotion LWM	200
row lock promotion PCT	100

For more information, see the System Administration Guide.

- The system procedure *sp_sysmon* reports on row lock promotions.
- Database-level row lock promotion thresholds are stored in the *master..sysattributes* table. If you dump a database, and load it only another server, you must set the row lock promotion thresholds on the new server. Object-level row lock promotion thresholds are stored in the *sysattributes* table in the user database, and are included in the dump.

Permissions

Only a System Administrator can execute *sp_setrowlockpromote*.

See also

System procedures – *sp_droprolockpromote*

sp_setsuspect_granularity

Description Displays or sets the recovery fault isolation mode for a user database, which governs how recovery behaves when it detects data corruption.

Syntax `sp_setsuspect_granularity [dbname [, "database" | "page" [, "read_only"]]]`

Parameters

- dbname*
 - is the name of the database for which to display or set the recovery fault isolation mode. For displaying, the default is the current database. For setting, you must be in the master database and specify the target *dbname*.
- database*
 - marks the entire database suspect, which makes it inaccessible, if the recovery process detects that any of its data is suspect.
- page*
 - marks only the corrupt pages suspect, making them inaccessible, if recovery detects corrupt data in the database. The rest of the data is accessible.
- read_only*
 - if specified, marks the entire database read only if recovery marks any pages suspect.

Examples

Example 1

```
sp_setsuspect_granularity
```

DB Name	Cur. Suspect Gran.	Cfg. Suspect Gran.	Online mode
pubs2	database	database	read/write

Displays the recovery fault isolation mode for the current database.

Example 2

```
sp_setsuspect_granularity pubs2
```

Displays the current and configured recovery fault isolation mode for the pubs2 database.

Example 3

```
sp_setsuspect_granularity pubs2, "page"
```

DB Name	Cur. Suspect Gran.	Cfg. Suspect Gran.
pubs2	database	database

`sp_setsuspect_granularity`: The new values will become effective during the next recovery of the database 'pubs2'.

The next time recovery runs in the pubs2 database, if any corrupt pages are detected, only the suspect pages will be taken offline and the rest of the database will be brought online.

Example 4

```
sp_setsuspect_granularity pubs2, "page",  
"read_only"
```

The next time recovery runs in the pubs2 database, if any corrupt pages are detected, only the suspect pages will be taken offline and the rest of the database will be brought online in read only mode.

Example 5

```
sp_setsuspect_granularity pubs2, "database"
```

The next time recovery runs in the pubs2 database, if any corrupt data is detected, the entire database will be marked suspect and taken offline.

Usage

- `sp_setsuspect_granularity` displays and sets the recovery fault isolation mode. This mode governs whether recovery marks an entire database or only the corrupt pages suspect when it detects that any data that it requires has been corrupted. For more information, see the System Administration Guide.
- The default recovery fault isolation mode of a user database is “database”. You can set the recovery fault isolation mode only for a user database, not for a system database.
- You must be in the master database to set the recovery fault isolation mode.
- Data marked suspect due to corruption persists across Adaptive Server start-ups. When certain pages have been marked suspect, they remain offline after you reboot the server.
- When part or all of a database is marked suspect, the suspect data is not accessible to users unless a System Administrator has made the suspect data accessible with the `sp_forceonline_db` and `sp_forceonline_page` procedures.
- General database corruption, such as a corrupt database log or the unavailability of another resource not specific to a page, causes the entire database to be marked suspect, even if the recovery fault isolation mode is “page”.

- If you do not specify page or database, Adaptive Server displays the current and configured settings. The current setting is the one that was in effect the last time recovery was executed in the database. The configured setting is the one that will be in effect the next time recovery is executed in the database.
- If the database comes online in read_only mode, no user can modify any of its data, including data that is unaffected by the suspect pages and is thus online. However, the system administrator can make the database writeable using the sp_dboption system procedure to set read only to false. In this case, users could then modify the online data, but the suspect data would remain inaccessible.

Permissions Only a System Administrator can execute sp_setsuspect_granularity to set the recovery fault isolation mode. Any user can execute sp_setsuspect_granularity to display the settings.

See also *Commands* – dump database, dump transaction, load database
System procedures – sp_forceonline_db, sp_forceonline_page, sp_listsuspect_db, sp_listsuspect_page, sp_setsuspect_threshold

sp_setsuspect_threshold

Description Displays or sets the maximum number of suspect pages that Adaptive Server allows in a database before marking the entire database suspect.

Syntax sp_setsuspect_threshold [*dbname* [, *threshold*]]

Parameters *dbname*
– is the name of the database for which you want to display or set the suspect escalation threshold. The default is the current database.
threshold
– indicates the maximum number of suspect data pages that recovery will allow before marking the entire database suspect. The default is 20 pages. The minimum is 0.

Examples **Example 1**

```
sp_setsuspect_threshold pubs2, 5
```

Sets the maximum number of suspect pages to five. If there are more than five suspect pages, recovery will mark the entire database suspect.

Example 2

```
sp_setsuspect_threshold pubs2
```

Displays the current and configured settings for the suspect escalation threshold for the pubs2 database.

Example 3

```
sp_setsuspect_threshold
```

Displays the current and configured settings for the recovery fault isolation threshold for the current user database.

Usage

- You must be in the master database to set the suspect escalation threshold with `sp_setsuspect_threshold`.
- If you do not specify the number of pages, Adaptive Server displays the current and configured settings. The current setting is the one that was in effect the last time recovery was executed in the database. The configured setting is the one that will be in effect the next time recovery is executed in the database.

Permissions

Only a System Administrator can execute `sp_setsuspect_threshold` to set the escalation threshold. Any user can execute `sp_setsuspect_threshold` to display the current settings.

See also

System procedures – `sp_forceonline_db`, `sp_forceonline_page`, `sp_listsuspect_db`, `sp_listsuspect_page`, `sp_setsuspect_granularity`

sp_showcontrolinfo

Description

Displays information about engine group assignments, bound client applications, logins, and stored procedures.

Syntax

```
sp_showcontrolinfo [object_type, object_name, spid]
```

Parameters

object_type

– is AP for application, LG for login, PR for stored procedure, EG for engine group, or PS for process. If you do not specify an *object_type* or specify an *object_type* of null, `sp_showcontrolinfo` displays information about all types.

object_name

– is the name of the application, login, stored procedure, or engine group. Do not specify an *object_name* if you specify PS as the *object_type*. If you do not specify an *object_name* (or specify an *object_name* of null), sp_showcontrolinfo displays information about all object names.

spid

– is the Adaptive Server process ID. Specify an *spid* only if you specify PS as the *object_type*. If you do not specify an *spid* (or specify an *spid* of null), sp_showcontrolinfo displays information for all spids. Use sp_who to see spids.

Examples

Example 1

```
sp_showcontrolinfo
```

Shows all user-assigned execution class-to-object bindings.

Example 2

```
sp_showcontrolinfo 'AP', 'isql'
```

Displays the execution class of the isql application.

Example 3

```
sp_showcontrolinfo 'PS'
```

Displays the execution class for all processes assigned to engine groups.

Example 4

```
sp_showcontrolinfo 'PS', null, 7
```

Displays the execution class for spid 7.

Usage

- When used with no parameters, sp_showcontrolinfo displays information about all user-assigned engine group assignments, bound client applications, logins, and stored procedures. When used with the object_type parameter, sp_showcontrolinfo provides information on an individual basis about application, login, or stored procedure bindings to an execution class, engine group compositions, and session-level attribute bindings. For more information, see the Performance and Tuning Guide.
- Unless object_type is PR, execute sp_showcontrolinfo from the master database. If object_type is PR, execute sp_showcontrolinfo from the database in which the procedure resides.

- If *object_type* is null, *sp_showcontrolinfo* displays execution class information for objects that match the other parameters.
- If *object_name* is null, *sp_showcontrolinfo* displays the binding information for all applications, logins, and stored procedures.
- If *spid* is null, *sp_showcontrolinfo* displays execution class information for objects that match the other parameters.

Permissions

Any user can execute *sp_showcontrolinfo*.

See also

System procedures – *sp_addexclass*, *sp_bindexclass*, *sp_clearpsex*, *sp_dropengine*, *sp_dropexclass*, *sp_showexclass*, *sp_showpsex*, *sp_unbindexclass*

Utility – *isql*

sp_showexclass

Description

Displays the execution class attributes and the engines in any engine group associated with the specified execution class.

Syntax

sp_showexclass [*exclassname*]

Parameters

exclassname
– is the name of an execution class.

Examples

Example 1

```
sp_showexclass
```

classname	priority	engine_group	engines
EC1	HIGH	ANYENGINE	ALL
EC2	MEDIUM	ANYENGINE	ALL
EC3	LOW	LASTONLINE	0

Displays the priority and engine group attribute values for all execution classes.

Example 2

```
sp_showexclass 'EC1'
```

classname	priority	engine_group	engines
EC1	HIGH	ANYENGINE	ALL

	Displays the attribute values of execution class EC1.
Usage	<ul style="list-style-type: none">• sp_showexeclass displays the execution class attributes and the engines in any engine group associated with <i>execlassname</i>. For more information, see the Performance and Tuning Guide.• If <i>execlassname</i> is NULL or absent, sp_showexeclass displays the priority and engine group attribute values for all execution classes, including the attribute values of the system-defined classes EC1, EC2, and EC3.
Permissions	Any user can execute sp_showexeclass.
See also	<i>System procedures</i> – sp_addexeclass, sp_bindexeclass, sp_dropexeclass, sp_showcontrolinfo, sp_unbindexeclass

sp_showplan

Description	Displays the showplan output for any user connection for the current SQL statement or for a previous statement in the same batch.
Syntax	<p>sp_showplan <i>spid</i>, <i>batch_id</i> output, <i>context_id</i> output, <i>stmt_num</i> output</p> <p>To display the showplan output for the current SQL statement without specifying the <i>batch_id</i>, <i>context_id</i>, or <i>stmt_num</i>:</p> <pre>sp_showplan <i>spid</i>, null, null, null</pre>
Parameters	<p><i>spid</i></p> <p>– is the process ID for any user connection. Use sp_who to see spids.</p> <p><i>batch_id</i></p> <p>– is a unique, nonnegative number for a batch</p> <p><i>context_id</i></p> <p>– is a unique number for every procedure (or trigger) executed in a batch.</p> <p><i>stmt_num</i></p> <p>– is the number of the current statement within a batch. The <i>stmt_num</i> must be a positive number.</p>
Examples	<p>Example 1</p> <pre>declare @batch int declare @context int declare @statement int</pre>

```
exec sp_showplan 99, @batch output, @context output,  
@statement output
```

Displays the query plan for the current statement running in the user session with a *spid* value of 99, as well as values for the *batch_id*, *context_id*, and *statement_id* parameters. These values can be used to retrieve query plans in subsequent iterations of *sp_showplan* for the user session with a *spid* of 99.

Example 2

```
sp_showplan 99, null, null, null
```

Displays the showplan output for the current statement running in the user session with a *spid* value of 99.

Usage

- *sp_showplan* displays the showplan output for a currently executing SQL statement or for a previous statement in the same batch.
- To see the query plan for the previous statement within the same batch, execute *sp_showplan* again with the same parameter values, but subtract 1 from the statement number. Using this method, you can view all the statements in the statement batch back to query number one.
- *sp_showplan* can be run independently of Adaptive Server Monitor™ Server.
- If the *context_id* is greater than 0 for a SQL batch, the current statement is embedded in a stored procedure (or trigger) called from the original SQL batch. Select the sysprocesses row with the same *spid* value to display the procedure ID and statement ID.

Permissions

Only a System Administrator can execute *sp_showplan*.

See also

System procedures – *sp_who*

sp_showpsex

Description

Displays execution class, current priority, and affinity for all client sessions running on Adaptive Server.

Syntax

```
sp_showpsex [spid]
```

Parameters

spid

– is the Adaptive Server session ID for which you want a report. The *spid* must belong to the application or login executing `sp_showpsexex`. Use `sp_who` to list *spids*.

Examples

Example 1

```
sp_showpsexex
```

spid	appl_name	login_name	exec_class	current_priority	task_affinity
1	isql	sa	EC1	HIGH	NONE
5		NULL	NULL	LOW	NULL
7	ctisql	sa	EC2	MEDIUM	NONE
8	ctisql	sa	EC2	MEDIUM	NONE

Displays execution class, current priority, and affinity for all current client sessions.

Example 2

```
sp_showpsexex 5
```

Displays the application name, login name, current priority, and engine affinity of the process with *spid* 5.

Usage

- `sp_showpsexex` displays execution class, current priority, and affinity for all sessions (objects with an *spid*). For more information, see the Performance and Tuning Guide.
- If the *spid* is NULL or absent, `sp_showpsexex` reports on all sessions currently running on Adaptive Server.
- `sp_showpsexex` does not report information for the following system processes: deadlock, checkpoint, network, auditing, and mirror handlers. It does display information for the housekeeper *spid*.

Permissions

Any user can execute `sp_showpsexex`.

See also

System procedures – `sp_addengine`, `sp_addexclass`, `sp_bindexclass`, `sp_clearpsexex`, `sp_dropengine`, `sp_dropexclass`, `sp_showcontrolinfo`, `sp_showexclass`, `sp_unbindexclass`

sp_spaceused

Description Displays estimates of the number of rows, the number of data pages, the size of indexes, and the space used by a specified table or by all tables in the current database.

Syntax `sp_spaceused [objname [,1]]`

Parameters *objname*
 – is the name of the table on which to report. If omitted, a summary of space used in the current database appears.

1
 – prints separate information on the table's indexes and text/image storage.

Examples

Example 1

```

      sp_spaceused titles
name          rowtotal  reserved  data    index_size  unused
-----
titles       18         46 KB    6 KB    4 KB        36 KB

```

Reports on the amount of space allocated (reserved) for the titles table, the amount used for data, the amount used for index(es), and the available (unused) space.

Example 2

```

      sp_spaceused titles, 1
index_name    size          reserved  unused
-----
titleidind    2 KB          32 KB    24 KB
titleind      2 KB          16 KB    14 KB

name          rowtotal  reserved  data    index_size  unused
-----
titles       18         46 KB    6 KB    4 KB        36 KB

```

In addition to information on the titles table, prints information for each index on the table.

Example 3

```

      sp_spaceused blurbs,1
index_name    size          reserved  unused
-----
blurbs        0 KB          14 KB    12 KB

```

```
tblurbs                14 KB      16 KB      2 KB

name      rowtotal reserved      data      index_size unused
-----
tblurbs      6          30 KB      2 KB      14 KB      14 KB
```

Displays the space taken up by the text/image page storage separately from the space used by the table. The object name for text/image storage is “t” plus the table name.

Example 4

```
sp_spaceused

database_name  database_size
-----
master          5 MB

reserved      data          index_size  unused
-----
2176 KB       1374 KB       72 KB       730 KB
```

Prints a summary of space used in the current database.

Example 5

```
sp_spaceused syslogs

name      rowtotal  reserved  data  index_size  unused
-----
syslogs   Not avail. 32 KB    32 KB  0 KB        0 KB
```

Reports on the amount of space reserved and the amount of space available for the transaction log.

Usage

- sp_spaceused displays estimates of the number of data pages, space used by a specified table or by all tables in the current database, and the number of rows in the tables. sp_spaceused computes the rowtotal value using the rowcnt built-in function. This function uses a value for the average number of rows per data page based on a value in the allocation pages for the object. This method is very fast, but the results are estimates, and update and insert activity change actual values. The update statistics command, dbcc checktable, and dbcc checkdb update the rows-per-page estimate, so rowtotal is most accurate after one of these commands executes. Always use select count(*) if you need exact row counts.
- sp_spaceused reports on the amount of space affected by tables, clustered indexes, and nonclustered indexes.

- The amount of space allocated (reserved) reported by *sp_spaceused* is a total of the data, index size, and available (unused) space.
- Space used by text and image columns, which are stored as separate database objects, is reported separately in the *index_size* column and is included in the summary line for a table. The object name for text/image storage in the *index_size* column is “t” plus the table name.
- When used on syslogs, *sp_spaceused* reports *rowtotal* as “Not available”. See example 5.

Permissions

Any user can execute *sp_spaceused*.

See also

Catalog stored procedures – *sp_statistics*

Commands – create index, create table, drop index, drop table

System procedures – *sp_helpindex*

sp_ssladmin

Description

Adds, deletes, or displays a list of server certificates for Adaptive Server.

Syntax

```
sp_ssladmin {[addcert, certificate_path [, password|NULL]]
[dropcert, certificate_path]
[!scert]
[help]}
```

Parameters

addcert

– adds a certificate for the local server in the certificates file.

certificate_path

– specifies the absolute path to the certificates file on the local server.

password

– the password that is used to encrypt the private key when adding a new server certificate to the certificates file.

NULL

– used to require an attended start-up of Adaptive Server by requesting the password during start-up from the command line.

dropcert

– deletes the certificate from the certificate file.

- ls-cert
– lists the certificates in the certificate file.
- help
– displays online help for sp_ssladmin.

Examples

Example 1

```
sp_ssladmin addcert, "/sybase/ASE-12_5/certificates/Server1.crt",  
"mypassword"
```

This adds an entry for the local server, *Server1.crt*, in the certificates file in the absolute path to */sybase/ASE-12_5/certificates* (*x:\sybase\ASE-12_5\certificates* on Windows). The private key is encrypted with the password “*mypassword*”. The password should be the one specified when you created the private key.

Example 2

```
sp_ssladmin dropcert, "/sybase/ASE-12_5/certificates/Server1.crt"
```

Deletes the certificate, *Server1.crt* from the certificates file located in */sybase/ASE-12_5/certificates* (*x:\sybase\ASE-12_5\certificates* on Windows).

Example 3

```
1> sp_ssladmin ls-cert  
2> go  
  
certificate_path  
-----  
/sybase/ASE-12_5/certificates/Server1.crt
```

Lists of all server certificates on the local server.

Usage

- The Adaptive Server listener must present to the client a certificate. The common name in the certificate must match the common name used by the client in the interfaces file. If they do not match, the server authentication and login fail.
- When NULL is specified as the password, *dataserver* must be started with a *-y* flag. This flag prompts the administrator for the private-key password at the command line.
- The use of NULL as the password is intended to protect passwords during the initial configuration of SSL, before the SSL encrypted session begins.

After restarting Adaptive Server with an SSL connection established, use *sp_ssladmin* again, this time using the actual password. The password is then encrypted and stored by Adaptive Server. Any subsequent starts of Adaptive Server from the command line would use the encrypted password; you do not have to specify the password on the command line during start up.

- You can specify “localhost” as the *hostname* in the *interfaces* file (*sql.ini* on Windows) to prevent clients from connecting remotely. Only a local connection can be established, and the password is never transmitted over a network connection.

Permissions

You must have the System Security Officer role to use *sp_ssladmin*.

sp_syntax

Description

Displays the syntax of Transact-SQL statements, system procedures, utilities, and other routines for Adaptive Server, depending on which products and corresponding *sp_syntax* scripts exist on your server.

Syntax

sp_syntax word [, *mod*][, *language*]

Parameters

word

- is the name or partial name of a command or routine; for example, “help”, to list all system procedures providing help. To include spaces or Transact-SQL reserved words, enclose the word in quotes.

mod

- is the name or partial name of one of the modules such as “Transact-SQL” or “Utility”. Each *sp_syntax* installation script adds different modules. Use *sp_syntax* without any parameters to see which modules exist on your server.

language

- is the language of the syntax description to be retrieved. *language* must be a valid language name in the *syslanguages* table.

Examples**Example 1**

```
sp_syntax
```

```
sp_syntax provides syntax help for Sybase products.  
These modules are installed on this Server:
```

```
Module
-----
OpenVMS
Transact-SQL
UNIX Utility
System Procedure
```

Usage: sp_syntax command [, module [, language]]

Displays all sp_syntax modules available on your server.

Example 2

```
sp_syntax "disk"
```

Displays the syntax and functional description of all routines containing the word or word fragment “disk”. Since “disk” is a Transact-SQL reserved word, enclose it in quotes.

Usage

- The text for sp_syntax is in the database sybsyntax. Load sp_syntax and the sybsyntax database onto Adaptive Server with the installation script described in configuration documentation for your platform. If you cannot access sp_syntax, see your System Administrator for information about installing it on your server.
- You can use wildcard characters within the command name you are searching for. However, if you are looking for a command or function that contains the literal “_”, you may get unexpected results, since the underscore wildcard character represents any single character.

Permissions

Any user can execute sp_syntax.

Tables used

sybsyntax..sybsyntax

See also

System procedures – sp_helpdb

sp_sysmon

Description

Displays performance information.

Syntax

```
sp_sysmon begin_sample
sp_sysmon { end_sample | interval }
[ , section [, applmon] ]
sp_sysmon { end_sample | interval } [, applmon ]
```

Parameters

begin_sample

– starts sampling. You cannot specify a section when you specify *begin_sample*.

end_sample

– ends sampling and prints the report.

interval

– specifies the time period for the sample. It must be in HH:MM:SS form, for example “00:20:00”.

section

– is the abbreviation for one of the sections printed by *sp_sysmon*. Table 23-2 lists the values and corresponding names of the report sections.

Table 23-2: *sp_sysmon* report sections

Report Section	Parameter
Application Management	apmgmt
Data Cache Management	dcache
Disk I/O Management	diskio
ESP Management	esp
Index Management	indexmgmt
Kernel Utilization	kernel
Lock Management	locks
Memory Management	memory
Metadata Cache Management	mdcache
Monitor Access to Executing SQL	monaccess
Network I/O Management	netio
Parallel Query Management	parallel
Procedure Cache Management	pcache
Recovery Management	recovery
Task Management	taskmgmt
Transaction Management	xactmgmt
Transaction Profile	xactsum
Worker Process Management	wpm

applmon

– specifies whether to print application detail, application and login detail, or no application detail. The default is to omit the application detail. Valid values are listed in Table 23-3.

Table 23-3: Values for applmon parameter to sp_sysmon

Parameter	Information Reported
appl_only	CPU, I/O, priority changes and resource limit violations by application name.
appl_and_login	CPU, I/O, priority changes and resource limit violations by application name and login name.
no_appl	Skips the by application or by login section of the report. This is the default.

This parameter is only valid when printing the full report and when you specify apmgmt for the *section*.

Examples

Example 1

```
sp_sysmon "00:10:00"
```

Prints monitor information after 10 minutes.

Example 2

```
sp_sysmon "00:05:00", diskio
```

Prints only the “Disk Management” section of the sp_sysmon report after 5 minutes.

Example 3

```
sp_sysmon begin_sample
go
execute proc1
go
execute proc2
go
select sum(total_sales) from titles
go
sp_sysmon end_sample, dcache
go
```

Starts the sample, executes procedures and a query, ends the sample, and prints only the “Data Cache” section of the report.

Example 4

```
sp_sysmon "00:05:00", @applmon = appl_and_login
```

Prints the full report and includes application and login detail for each login.

Usage	<ul style="list-style-type: none">• <i>sp_sysmon</i> displays information about Adaptive Server performance. It sets internal counters to 0, then waits for the specified interval while activity on the server causes the counters to be incremented. When the interval ends, <i>sp_sysmon</i> prints information from the values in the counters. For more information, see the Performance and Tuning Guide.• To print only a single section of the report, use the values listed in Table 23-3 for the second parameter.• If you use <i>sp_sysmon</i> in batch mode, with <i>begin_sample</i> and <i>end_sample</i>, the time interval between executions must be at least one second. You can use <i>waitfor delay "00:00:01"</i> to lengthen the execution time of a batch.• During the sample interval, results are stored in signed integer values. Especially on systems with many CPUs and high activity, these counters can overflow. If you see negative results in your <i>sp_sysmon</i> output, reduce your sample time.
Permissions	Only a System Administrator can execute <i>sp_sysmon</i> .

sp_thresholdaction

Description	Executes automatically when the number of free pages on the log segment falls below the last-chance threshold, unless the threshold is associated with a different procedure. Sybase does not provide this procedure.
Syntax	When a threshold is crossed, Adaptive Server passes the following parameters to the threshold procedure by position: <pre>sp_thresholdaction @dbname, @segment_name, @space_left, @status</pre>
Parameters	<p><i>@dbname</i> – is the name of a database where the threshold was reached.</p> <p><i>@segment_name</i> – is the name of the segment where the threshold was reached.</p> <p><i>@space_left</i> – is the threshold size, in logical pages.</p>

@status

– is 1 for the last-chance threshold; 0 for all other thresholds.

Examples

```
create procedure sp_thresholdaction
    @dbname varchar(30),
    @segmentname varchar(30),
    @space_left int,
    @status int
as
    dump transaction @dbname to tapedump1
```

Creates a threshold procedure for the last-chance threshold that dumps the transaction log to a tape device.

Usage

- sp_thresholdaction must be created by the Database Owner (in a user database), or a System Administrator (in the sybsystemprocs database), or a user with create procedure permission.
- You can add thresholds and create threshold procedures for any segment in a database.
- When the last-chance threshold is crossed, Adaptive Server searches for the sp_thresholdaction procedure in the database where the threshold event occurs. If it does not exist in that database, Adaptive Server searches for it in sybsystemprocs. If it does not exist in sybsystemprocs, it searches master. If Adaptive Server does not find the procedure, it sends an error message to the error log.
- sp_thresholdaction should contain a dump transaction command to truncate the transaction log.
- By design, the last-chance threshold allows enough free space to record a dump transaction command. There may not be enough space to record additional user transactions against the database. Only commands that are not recorded in the transaction log (select, fast bcp, readtext, and writetext) and commands that might be necessary to free additional log space (dump transaction, dump database, and alter database) can be executed. By default, other commands are suspended and a message is sent to the error log. To abort these commands rather than suspend them, use the abort tran on log full option of sp_dboption followed by the checkpoint command.

Waking suspended processes

- Once the dump transaction command frees sufficient log space, suspended processes automatically awaken and complete.

- If fast bcp, writetext, or select into have resulted in unlogged changes to the database since the last backup, the last-chance threshold procedure cannot execute a dump transaction command. When this occurs, use dump database to make a copy of the database, then use dump transaction to truncate the transaction log.
- If this does not free enough space to awaken the suspended processes, it may be necessary to increase the size of the transaction log. Use the log on option of the alter database command to allocate additional log space.
- As a last resort, System Administrators can use sp_who to determine which processes are suspended, then use the kill command to kill them.

See also

Commands – create procedure, dump transaction

System procedures – sp_addthreshold, sp_dboption, sp_droptreshold, sp_helpsegment, sp_helpthreshold, sp_modifythreshold

sp_transactions

Description	Reports information about active transactions.
Syntax	sp_transactions ["xid", <i>xid_value</i>] ["state", {"heuristic_commit" "heuristic_abort" "prepared" "indoubt"} [, "xactname"]] ["gtrid", <i>gtrid_value</i>]
Parameters	<i>xid_value</i> – is a transaction name from the xactname column of master.dbo.systransactions. <i>gtrid_value</i> – is the global transaction ID name for a transaction coordinated by Adaptive Server.
Examples	Example 1

```

sp_transactions

xactkey                type                coordinator starttime
state                  connection dbid    spid    loid
failover                srvname                namelen
xactname
-----
-----

```


- sp_transactions with the xid keyword displays the gtrid, commit_node, and parent_node columns only for the specified transaction.
- sp_transactions with the state keyword displays information only for the active transactions in the specified state.

sp_transactions with both xid and xactname displays only the transaction names for transactions in the specified state.
- sp_transactions with the gtrid keyword displays information only for the transactions with the specified global transaction ID.
- sp_transactions replaces the sp_xa_scan_xact procedure provided with XA-Library and XA-Server products.
- For more information, see *Using Adaptive Server Distributed Transaction Management Features*.

Column descriptions for sp_transactions output

- The xactkey column shows the internal transaction key that Adaptive Server uses to uniquely identify the transaction.
- The type column indicates the type of transaction:
 - “Local” means that the transaction was explicitly started on the local Adaptive Server with a begin transaction statement.
 - “Remote” indicates a transaction executing on a remote Adaptive Server.
 - “External” means that the transaction has an external coordinator associated with it. For example, transactions coordinated by a remote Adaptive Server, MSDTC, or an X/Open XA transaction manager are flagged as “External.”
 - “Dtx_State” is a special state for distributed transactions coordinated by Adaptive Server. It indicates that a transaction on the local server was either committed or aborted, but Adaptive Server has been unable to resolve a branch of that transaction on a remote participant. This may happen in cases where Adaptive Server loses contact with a server it is coordinating.
- The coordinator column indicates the method or protocol used to manage a distributed transaction:

sp_transactions “coordinator” value	Meaning
None	Transaction is not a distributed transaction and does not require a coordinating protocol.

<i>sp_transactions</i> “coordinator” value	Meaning
ASTC	Transaction is coordinated using the Adaptive Server transaction coordination services.
XA	Transaction is coordinated by the X/Open XA-compliant transaction manager via the Adaptive Server XA-Library interface. Such transaction managers include Encina, CICS, and Tuxedo.
DTC	Transaction is coordinated by MSDTC.
SYB2PC	Transaction is coordinated using Sybase two-phase commit protocol.

- The starttime column indicates the time that the transaction started.
- The state column indicates the state of the transaction at the time *sp_transactions* ran:

<i>sp_transactions</i> “state” value	Meaning
Begun	Transaction has begun but no updates have been performed.
Done Command	Transaction completed an update command.
Done	X/Open XA transaction has finished modifying data.
Prepared	Transaction has successfully prepared.
In Command	Transaction is currently modifying data.
In Abort Cmd	Execution of the current command in the transaction has been aborted.
Committed	Transaction has successfully committed, and the commit log record has been written.
In Post Commit	Transaction has successfully committed, but is currently deallocating transaction resources.
In Abort Tran	The transaction is being aborted. This may happen either as a result of an explicit command, or because of a system failure.
In Abort Savept	The transaction is being rolled back to a savepoint.
Begun-Detached	Transaction has begun, but there is no thread currently attached to it.
Done Cmd-Detached	Transaction has finished modifying data, and no thread is currently attached to it.
Done-Detached	Transaction will modify no more data, and no thread is currently attached to it.

sp_transactions "state" value	Meaning
Prepared-Detached	Transaction has successfully prepared, and no thread is currently attached to it.
Heur Committed	Transaction has been heuristically committed using the dbcc complete_xact command.
Heur Rolledback	Transaction has been heuristically rolled back using the dbcc complete_xact command.

- The connection column indicates whether or not the transaction is currently associated with a thread:
 - “Attached” indicates that the transaction has an associated thread of control.
 - “Detached” indicates that there is no thread currently associated with the transaction. Some external transaction managers, such as CICS and TUXEDO, use the X/Open XA “suspend” and “join” semantics to associate different threads with the same transaction.
- The dbid column indicates the database ID of the database in which transaction started.
- The spid column indicates the server process ID associated with the transaction. If the transaction is “Detached,” the “spid” value is 0.
- The loid column indicates the unique lock owner ID from master.dbo.systransactions.
- The failover column indicates the failover state for the transaction:
 - “Resident Tx” indicates that the transaction started and is executing on the same server. “Resident Tx” is displayed under normal operating conditions, and on systems that do not utilize Adaptive Server high availability features.
 - “Failed-over Tx” is displayed after there has been a failover to a secondary companion server. “Failed-over Tx” means that a transaction originally started on a primary server and reached the prepared state, but was automatically migrated to the secondary companion server (for example, as a result of a system failure on the primary server). The migration of a prepared transaction occurs transparently to an external coordinating service.

- “Tx by Failover-Conn” indicates that there was an attempt to start the transaction on a designated server, but the transaction was instead started on the secondary companion server. This occurs when the original server has experienced a failover condition.
- The *srvname* column indicates the name of the remote server on which the transaction is executing. This column is only meaningful for remote transactions. For local and external transactions, *srvname* is null.
- The *namelen* column indicates the total length of the *xactname* value.
- *xactname* is the transaction name. For local transactions, the transaction name may be defined as part of the begin transaction command. External transaction managers supply unique transaction names in a variety of formats. For example, X/Open XA-compliant transaction managers supply a transaction ID (*xid*) consisting of a global transaction identifier and a branch qualifier, both of which are stored in *xactname*.
- For transactions coordinated by Adaptive Server, the *gtrid* column displays the global transaction ID. Transaction branches that are part of the same distributed transaction share the same *gtrid*. You can use a specific *gtrid* with the *sp_transactions* *gtrid* keyword to determine the state of other transaction branches in the same distributed transaction.

sp_transactions cannot display the *gtrid* for transactions that have an external coordinator. For transactions coordinated by an X/Open XA-compliant transaction manager, MSDTC, or SYB2PC, the *gtrid* column shows the full transaction name supplied by the external coordinator.
- For transactions coordinated by Adaptive Server, the *commit_node* column indicates the server that executes the outermost block of the distributed transaction. This outermost block ultimately determines the commit status of all subordinate transactions.

For transactions not coordinated by Adaptive Server, *commit_node* displays one of the values described in Table 23-4.

Table 23-4: Values for commit_node and parent_node

Value	Meaning
server_name	Commit or parent node is an Adaptive Server with the specified server_name.
XATM	Commit or parent node is an X/Open XA-compliant transaction manager.
MSDTCTM	Commit or parent node is MSDTC.
SYB2PCTM	Transaction is coordinated using SYB2PC protocol.

- For transactions coordinated by Adaptive Server, the parent_node column indicates the server that is coordinating the external transaction on the local server.

For transactions not coordinated by Adaptive Server, parent_node displays one of the values described in Table 23-4.

Note The values for commit_node and parent_node can be different, depending on the levels of hierarchy in the distributed transaction.

Permissions

Any user can execute sp_transactions.

See also

System procedures – sp_lock, sp_who

Procedures: *sp_unbindcache* – *sp_who*

sp_unbindcache

Description Unbinds a database, table, index, text object, or image object from a data cache.

Syntax `sp_unbindcache dbname [, [owner.]tablename
[, indexname | "text only"]]`

Parameters

dbname

– is the name of database to be unbound or the name of the database containing the objects to be unbound.

owner

– is the name of the table's owner. If the table is owned by the Database Owner, the owner name is optional.

tablename

– is the name of the table to be unbound from a cache or the name of a table whose index, text object, or image object is to be unbound from a cache.

indexname

– is the name of an index to be unbound from a cache.

text only

– unbinds text or image objects from a cache.

Examples

Example 1

```
sp_unbindcache pubs2, titles
```

Unbinds the titles table from the cache to which it is bound.

Example 2

```
sp_unbindcache pubs2, titles, titleidind
```

Unbinds the titleidind index from the from the cache to which it is bound.

Example 3

```
sp_unbindcache pubs2, au_pix, text
```

Unbinds the text or image object for the au_pix table from the cache to which it is bound.

Example 4

```
sp_unbindcache pubs2, syslogs
```

Unbinds the transaction log, syslogs, from its cache.

Usage

- When you unbind a database or database object from a cache, all subsequent I/O for the cache is performed in the default data cache. All dirty pages in the cache being unbound are written to disk, and all clean pages are cleared from the cache. For more information, see the Performance and Tuning Guide.
- Cache unbindings take effect immediately and do not require a restart of the server.
- When you drop a database, table, or index, its cache bindings are automatically dropped.
- To unbind a database, you must be using the master database. For tables, indexes, text objects, or image objects, you must be using the database where the objects are stored.
- To unbind any system tables in a database, you must be using the database, and the database must be in single-user mode. Use the command:

```
sp_dboption db_name, "single user", true
```

See sp_dboption for more information.

- The following procedures provide information about the bindings for their respective objects: sp_helpdb for databases, sp_help for tables, and sp_helpindex for indexes.
- sp_helpcache prints the names of objects bound to caches.
- sp_unbindcache needs to acquire an exclusive table lock when you are unbinding a table or its indexes to a cache. No pages can be read while the unbinding takes place. If a user holds locks on a table, and you issue sp_unbindcache on that object, the sp_unbindcache task sleeps until the locks are released.
- When you change the cache binding for an object with sp_bindcache or sp_unbindcache, the stored procedures that reference the object are recompiled the next time they are executed. When you change the binding for a database, the stored procedures that reference objects in the database are recompiled the next time they are executed.

- To unbind all objects from a cache, use the system procedure `sp_unbindcache_all`.

Permissions

Only a System Administrator can execute `sp_unbindcache`.

See also

System procedures – `sp_bindcache`, `sp_dboption`, `sp_helpcache`, `sp_unbindcache_all`

sp_unbindcache_all

Description

Unbinds all objects that are bound to a cache.

Syntax

`sp_unbindcache_all cache_name`

Parameters

cache_name

- is the name of the data cache from which objects are to be unbound.

Examples

```
sp_unbindcache_all pub_cache
```

Unbinds all databases, tables, indexes, text objects and image objects that are bound to `pub_cache`.

Usage

- When you unbind entities from a cache, all subsequent I/O for the cache is performed in the default cache.
- To unbind individual objects from a cache, use the system procedure `sp_unbindcache`.
- See `sp_unbindcache` for more information about unbinding caches.

Permissions

Only a System Administrator can execute `sp_unbindcache_all`.

See also

System procedures – `sp_bindcache`, `sp_helpcache`, `sp_unbindcache`

sp_unbindefault

Description

Unbinds a created default value from a column or from a user-defined datatype.

Syntax

`sp_unbindefault objname [, futureonly]`

Parameters

objname

– is the name of either the table and column or the user-defined datatype from which to unbind the default. If the parameter is not of the form “*table.column*”, then *objname* is assumed to be a user-defined datatype. When unbinding a default from a user-defined datatype, any columns of that type that have the same default as the user-defined datatype are also unbound. Columns of that type, whose default has already been changed, are unaffected.

futureonly

– prevents existing columns of the specified user-defined datatype from losing their defaults. It is ignored when unbinding a default from a column.

Examples

Example 1

```
sp_unbinddefault "employees.startdate"
```

Unbinds the default from the startdate column of the employees table.

Example 2

```
sp_unbinddefault ssn
```

Unbinds the default from the user-defined datatype named ssn and all columns of that type.

Example 3

```
sp_unbinddefault ssn, futureonly
```

Unbinds defaults from the user-defined datatype ssn, but does not affect existing columns of that type.

Usage

- Use sp_unbinddefault to remove defaults created with sp_bindefault. Use alter table to drop defaults declared using the create table or alter table statements.
- Columns of a user-defined datatype lose their current default unless the default has been changed or the value of the optional second parameter is futureonly.
- To display the text of a default, execute sp_helptext with the default name as the parameter.

Permissions

Only the object owner can execute sp_unbinddefault.

See also

Commands – create default, drop default

System procedures – sp_bindefault, sp_helptext

sp_unbindexclass

Description	Removes the execution class attribute previously associated with an client application, login, or stored procedure for the specified scope.
Syntax	<code>sp_unbindexclass object_name, object_type, scope</code>
Parameters	<p><i>object_name</i></p> <ul style="list-style-type: none"> – is the name of the application, login, or stored procedure for which to remove the association to the execution class. <p><i>object_type</i></p> <ul style="list-style-type: none"> – identifies the type of <i>object_name</i> as ap, lg, or pr for application, login, or stored procedure. <p><i>scope</i></p> <ul style="list-style-type: none"> – is the application name or the login name for which the unbinding applies for an application or login. It is the stored procedure owner name (user name) for stored procedures.
Examples	<pre>sp_unbindexclass 'sa', 'lg', 'isql'</pre> <p>Removes the association between “sa” login scoped to application isql and an execution class. “sa” automatically binds itself to another execution class, depending on other binding specifications, precedence, and scoping rules. If no other binding is applicable, the object binds to the default execution class, EC2.</p>
Usage	<ul style="list-style-type: none"> • The parameters must match an existing entry in the sysattributes system table. • If you specify a null value for scope, Adaptive Server unbinds the object for which the scope is null, if there is one. • A null value for scope does not indicate that unbinding should apply to all bound objects. • When unbinding a stored procedure from an execution class, you must use the name of the stored procedure owner (user name) for the scope parameter. • Stored procedures can be dropped before or after unbinding. • A user cannot be dropped from a database if the user owns a stored procedure that is bound to an execution class in that database. • Unbind objects of type PR before dropping them from the database.

- Unbinding will fail if the associated engine group has no online engines and active processes are bound to the associated execution class.
- Due to precedence and scoping rules, the execution class being unbound may or may not have been in effect for the object called `object_name`. The object automatically binds itself to another execution class, depending on other binding specifications and precedence and scoping rules. If no other binding is applicable, the object binds to the default execution class, EC2.

Permissions Only a System Administrator can execute `sp_unbindexclass`.

See also *System procedures* – `sp_addexclass`, `sp_bindexclass`, `sp_dropexclass`, `sp_showexclass`

Utility – `isql`

sp_unbindmsg

Description Unbinds a user-defined message from a constraint.

Syntax `sp_unbindmsg constrname`

Parameters *constrname*
– is the name of the constraint from which a message is to be unbound.

Examples `sp_unbindmsg positive_balance`
Unbinds a user-defined message from the constraint `positive_balance`.

Usage

- You can bind only one message to a constraint. To change the message bound to a constraint, use `sp_bindmsg`; the new message number replaces any existing bound message. It is not necessary to use `sp_unbindmsg` first.
- To retrieve message text from the `sysusermessages` table, execute `sp_getmessage`.

Permissions Only the object owner can execute `sp_unbindmsg`.

See also *System procedures* – `sp_addmessage`, `sp_bindmsg`, `sp_getmessage`

sp_unbindrule

Description	Unbinds a rule from a column or from a user-defined datatype.
Syntax	<code>sp_unbindrule <i>objname</i> [, futureonly]</code>
Parameters	<p><i>objname</i></p> <ul style="list-style-type: none"> – is the name of the table and column or of the user-defined datatype from which the rule is to be unbound. If the parameter is not of the form “<i>table.column</i>”, then <i>objname</i> is assumed to be a user-defined datatype. Unbinding a rule from a user-defined datatype also unbinds it from columns of the same type. Columns that are already bound to a different rule are unaffected. <p>futureonly</p> <ul style="list-style-type: none"> – prevents columns of the specified user-defined datatype from losing their rules. It is ignored when unbinding a rule from a column.
Examples	<p>Example 1</p> <pre>sp_unbindrule "employees.startdate"</pre> <p>Unbinds the rule from the startdate column of the employees table.</p> <p>Example 2</p> <pre>sp_unbindrule def_ssn</pre> <p>Unbinds the rule from the user-defined datatype named def_ssn and all columns of that type.</p> <p>Example 3</p> <pre>sp_unbindrule ssn, futureonly</pre> <p>The user-defined datatype ssn no longer has a rule, but existing ssn columns are unaffected.</p>
Usage	<ul style="list-style-type: none"> • Executing <code>sp_unbindrule</code> removes a rule from a column or from a user-defined datatype in the current database. If you do not want to unbind the rule from existing <i>objname</i> columns, use <code>futureonly</code> as the second parameter. • You cannot use <code>sp_unbindrule</code> to unbind a check constraint. Use <code>alter table</code> to drop the constraint. • To unbind a rule from a table column, specify the <i>objname</i> argument in the form “<i>table.column</i>”.

- The rule is unbound from all existing columns of the user-defined datatype unless the rule has been changed or the value of the optional second parameter is futureonly.
- To display the text of a rule, execute sp_helptext with the rule name as the parameter.

Permissions Only the object owner can execute sp_unbindrule.

See also *Commands* – create rule, drop rule

System procedures – sp_bindrule, sp_helptext

sp_volchanged

Description Notifies the Backup Server that the operator performed the requested volume handling during a dump or load.

Syntax `sp_volchanged session_id, devname, action
[, fname [, vname]]`

Parameters

session_id
– identifies the Backup Server session that requested the volume change. Use the @*session_id* parameter specified in the Backup Server’s volume change request.

devname
– is the device on which a new volume was mounted. Use the @*devname* parameter specified in the Backup Server’s volume change request. If the Backup Server is not located on the same machine as the Adaptive Server, use the form:
`device at backup_server_name`

action
– indicates whether the Backup Server should abort, proceed with, or retry the dump or load.

fname
– is the file to be loaded. If you do not specify a file name with sp_volchanged, the Backup Server loads the file =*filename* parameter of the load command. If neither sp_volchanged nor the load command specifies which file to load, the Backup Server loads the first file on the tape.

vname

– is the volume name that appears in the ANSI tape label. The Backup Server writes the volume name in the ANSI tape label when overwriting an existing dump, dumping to a brand new tape, or dumping to a tape whose contents are not recognizable. If you do not specify a *vname* with *sp_volchanged*, the Backup Server uses the *dumpvolume* value specified in the dump command. If neither *sp_volchanged* nor the dump command specifies a volume name, the Backup Server leaves the name field of the ANSI tape label blank.

During loads, the Backup Server uses the *vname* to confirm that the correct tape has been mounted. If you do not specify a *vname* with *sp_volchanged*, the Backup Server uses the *dumpvolume* specified in the load command. If neither *sp_volchanged* nor the load command specifies a volume name, the Backup Server does not check the name field of the ANSI tape label before loading the dump.

Examples

```
sp_volchanged 8, "/dev/nrmt4", RETRY
```

The following message from Backup Server indicates that a mounted tape's expiration date has not been reached:

```
Backup Server: 4.49.1.1: OPERATOR: Volume to be overwritten on
'/dev/rmt4' has not expired: creation date on this volume is Sunday, Nov.
15, 1992, expiration date is Wednesday, Nov. 25, 1992.
Backup Server: 4.78.1.1: EXECUTE sp_volchanged
@session_id = 8,
@devname = '/auto/remote/pubs3/SERV/Masters/testdump',
@action = { 'PROCEED' | 'RETRY' | 'ABORT' }
```

The operator changes the tape, then issues the command in the example.

Usage

- If the Backup Server detects a problem with the currently mounted volume, it requests a volume change:
 - On OpenVMS systems, the Backup Server sends volume change messages to the operator terminal on the machine on which it is running. Use the *with notify = client* option of the dump or load command to route other Backup Server messages to the terminal session on which the dump or load request initiated.
 - On UNIX systems, the Backup Server sends messages to the client that initiated the dump or load request. Use the *with notify = operator_console* option of the dump or load command to route messages to the terminal where the Backup Server was started.

- After mounting another volume, the operator executes sp_volchanged from any Adaptive Server that can communicate with the Backup Server performing the dump or load. The operator does not have to log into the Adaptive Server on which the dump or load originated.
- On OpenVMS systems, the operating system—not the Backup Server—requests a volume change when it detects the end of a volume or when the specified drive is offline. The operator uses the OpenVMS REPLY command to reply to these messages.
- On UNIX systems, the Backup Server requests a volume change when the tape capacity has been reached. The operator mounts another tape and executes sp_volchanged. Table 24-1 illustrates this process.

Table 24-1: Changing tape volumes on a UNIX system

Sequence	Operator, Using isql	Adaptive Server	Backup Server
1	Issues the dump database command		
2		Sends dump request to Backup Server	
3			Receives dump request message from Adaptive Server Sends message for tape mounting to operator Waits for operator's reply
4	Receives volume change request from Backup Server Mounts tapes Executes sp_volchanged		
5			Checks tapes If tapes are okay, begins dump When tape is full, sends volume change request to operator
6	Receives volume change request from Backup Server Mounts tapes Executes sp_volchanged		

Sequence	Operator, Using isql	Adaptive Server	Backup Server
7			Continues dump When dump is complete, sends messages to operator and Adaptive Server
8	Receives message that dump is complete Removes and labels tapes	Receives message that dump is complete Releases locks Completes the dump database command	

Permissions Any user can execute *sp_volchanged*.

See also *Commands* – dump database, dump transaction, load database, load transaction

Utility – isql

sp_who

Description Reports information about all current Adaptive Server users and processes or about a particular user or process.

Syntax *sp_who* [*loginame* | "*spid*"]

Parameters *loginame*
– is the Adaptive Server login name of the user you are requesting a report on.

spid
– is the number of the process you are requesting a report on. Enclose process numbers in quotes (Adaptive Server expects a char type).

Examples **Example 1**

```

sp_who

fid spid  status      loginame  origname  hostname  blk_spid  dbname
      cmd
-----
0      1  recv sleep  bird      bird      jazzy     0         master
      AWAITING COMMAND  0x0000ed92
0      2  sleeping NULL      NULL               0         master

```

```

NETWORK HANDLER      0x0000ed92
0  3  sleeping  NULL      NULL      0      master
MIRROR HANDLER      0x0000ed92
0  4  sleeping  NULL      NULL      0      master
AUDIT PROCESS        0x0000ed92
0  5  sleeping  NULL      NULL      0      master
CHECKPOINT SLEEP    0x0000ed92
0  6  recv sleep  rose      rose      petal   0      master
AWAITING COMMAND    0x0000ed92
0  7  sleeping  NULL      NULL      actor   0      sybssystemdb
ASTC HANDLER        0x0000ed92
0  8  running   robert    sa        helos   0      master
SELECT              0x0000ed92
0  9  send sleep  daisy     daisy     chain   0      pubs2
SELECT              0x0000ed92
0 10  alarm sleep  lily      lily      pond    0      master
WAITFOR             0x0000ed92
0 11  lock sleep  viola     viola     cello   8      pubs2
SELECT              0x0000ed92

```

Reports on the processes running on Adaptive Server. Process 11 (a select on a table) is blocked by process 8 (a begin transaction followed by an insert on the same table).

For process 8, the current *loginame* is “robert”, but the original *loginame* is “sa”. Login “sa” executed a set proxy command to impersonate the user “robert”.

Example 2

```
sp_who victoria
```

Reports on the processes being run by the user “victoria”.

Example 3

```
sp_who "17"
```

Reports what Adaptive Server process number 17 is doing.

Example 4

```

sp_who

```

fid	spid	status	loginame	origname	hostname	blk_spid	dbname
	cmd		block_xloid				
0	1	running	sa	sa	helos	0	master
	SELECT			0			

0	2	sleeping	NULL	NULL	0	master
		NETWORK HANDLER	0			
0	3	sleeping	NULL	NULL	0	master
		DEADLOCK TUNE	0			
0	4	sleeping	NULL	NULL	0	master
		MIRROR HANDLER	0			
0	5	sleeping	NULL	NULL	actor	0
		ASTC HANDLER	0			
0	6	sleeping	NULL	NULL	0	master
		CHECKPOINT SLEEP	0			
0	5	sleeping	NULL	NULL	0	master
		HOUSEKEEPER	0			

Reports on the processes running on Adaptive Server. Although no user processes other than *sp_who* are running, the server still shows activity. During idle cycles, the housekeeper task moves dirty buffers into the buffer wash region.

Usage

- *sp_who* reports information about a specified user or Adaptive Server process.
- Without parameters, *sp_who* reports which users are running what processes in all databases.
- The columns returned by *sp_who* are:
 - *fid* – identifies the family (including the coordinating process and its worker processes) to which a lock belongs. For more information, see *sp_familylock*.
 - *spid* – identifies the process number. A System Administrator can use this number with the Transact-SQL *kill* command to stop the process.
 - *status* – indicates whether the process is running or sleeping.
 - *loginame* – the login or alias of the user who started the process. For all system processes, *loginame* is *NULL*.
 - *origname* – If the *loginame* is an alias, *origname* shows the real login name. If not, *origname* shows the same information as *loginame*.
 - *hostname* – the name of the server on which the database resides.
 - *blk_spid* – contains the process IDs of the blocking process, if there is one. A blocking process (which may be infected or have an exclusive lock) is one that is holding resources needed by another process.

- `dbname` – indicates the name of the database on which the process is running.
- `cmd` – identifies the command or process currently being executed. Evaluation of a conditional statement, such as an `if` or `while` loop, returns `cond`.
- `block_xloid` – identifies the unique lock owner ID of a blocking transaction.
- Running `sp_who` on a single-engine server shows the `sp_who` process currently running and all other processes that are runnable or in one of the sleep states. In multi-engine servers, there can be a “running” process for each engine.
- If you enable mirrored disks or remote procedure calls, the mirror handler and the site handler also appear in the report from `sp_who`.

Permissions

Any user can execute `sp_who`.

See also

Commands – `kill`

System procedures – `sp_lock`

This chapter describes catalog stored procedures, which retrieve information from the system tables in tabular form.

Table 25-1 lists the catalog stored procedures that are covered in this chapter.

Table 25-1: Catalog stored procedures

Procedure	Description
sp_column_privileges	Returns permissions information for one or more columns in a table or view.
sp_columns	Returns information about the type of data that can be stored in one or more columns.
sp_databases	Returns a list of the databases in Adaptive Server.
sp_datatype_info	Returns information about a particular datatype or about all supported datatypes.
sp_fkeys	Returns information about foreign key constraints created in the current database with the create table or alter table command.
sp_pkeys	Returns information about primary key constraints created for a single table with the create table or alter table command.
sp_server_info	Returns a list of Adaptive Server attribute names and current values.
sp_special_columns	Returns the optimal set of columns that uniquely identify a row in a table or view; can also return a list of the columns that are automatically updated when any value in the row is updated by a transaction.
sp_proc_columns	Returns information about a stored procedure's input and return parameters.
sp_statistics	Returns a list of indexes on a single table.
sp_stored_procedures	Returns information about one or more stored procedures.
sp_table_privileges	Returns privilege information for all columns in a table or view.
sp_tables	Returns a list of objects that can appear in a from clause.

Overview

Catalog stored procedures retrieve information from the system tables in tabular form.

The catalog stored procedures, created by installmaster at installation, are located in the sybssystemprocs database and are owned by the System Administrator.

Many of them can be run from any database. If a catalog stored procedure is executed from a database other than sybssystemprocs, it retrieves information from the system tables in the database from which it was executed.

All catalog stored procedures execute at isolation level 1.

All catalog stored procedures report a return status. For example, this means that the procedure executed successfully. The examples in this book do not include the return status:

```
return status = 0
```

Specifying optional parameters

If a parameter value for a catalog stored procedure contains punctuation or embedded blanks, or is a reserved word, you must enclose it in single or double quotes. If the parameter is an object name qualified by a database name or owner name, enclose the entire name in single or double quotes.

Note Do not use delimited identifiers as catalog stored procedure parameters. Doing so may produce unexpected results.

In many cases, it is more convenient to supply parameters to the catalog stored procedures in the form:

```
@parametername = value
```

than to supply all the parameters. The parameter names in the syntax statements match the parameter names defined by the procedures.

For example, the syntax for sp_columns is:

```
sp_columns table_name [, table_owner]  
[, table_qualifier] [, column_name]
```

To use sp_columns to find information about a particular column, you can use:

```
sp_columns publishers, @column_name = "pub_id"
```

This provides the same information as the command with all of the parameters specified:

```
sp_columns publishers, "dbo", "pubs2", "pub_id"
```

You can also use “null” as a placeholder:

```
sp_columns publishers, null, null, "pub_id"
```

If you specify more parameters than the number of parameters expected by the system procedure, Adaptive Server ignores the extra parameters.

Pattern matching

Adaptive Server offers a wide range of pattern matching through regular expressions. However, for maximum interoperability, assume only SQL standards pattern matching (the % and _ wildcard characters).

System procedure tables

These catalog stored procedures	Use these catalog stored procedure tables
sp_columns	spt_datatype_info
sp_datatype_info	spt_datatype_info_ext
sp_special_columns	spt_server_info
sp_sproc_columns	

in the sybssystemprocs database to convert internal system values such as status bits into human-readable format.

The catalog stored procedures sp_column_privileges and sp_table_privileges create and then drop temporary tables.

ODBC datatypes

Table 25-2 and Table 25-3 list the datatype code numbers and matching datatype names returned by `sp_columns` and `sp_sproc_columns` in the “`data_type`” column. The source for the description is the Open Database Connectivity (ODBC) Application Programming Interface (API).

Table 25-2: Code numbers for ODBC datatypes

Datatype	Code #
char	1
decimal	3
double precision	8
float	6
integer	4
numeric	2
real	7
smallint	5
varchar	12

Table 25-3: Code numbers for extended datatypes

Datatype	Code #
bigint	-5
binary (bit datatype)	-2
bit	-7
date	9
long varbinary	-4
long varchar	-1
time	10
timestamp	11
tinyint	-6
varbinary (bit-varying datatype)	-3

sp_column_privileges

Description Returns permissions information for one or more columns in a table or view.

Syntax `sp_column_privileges table_name [, table_owner
[, table_qualifier [, column_name]]]`

Parameters

table_name
– is the name of the table. The use of wildcard characters in pattern matching is not supported.

table_owner
– is the name of the table owner. The use of wildcard characters in pattern matching is not supported. If you do not specify the table's owner, `sp_column_privileges` looks for a table owned by the current user and then for a table owned by the Database Owner.

table_qualifier
– is the name of the database. Values are the name of the current database and null.

column_name
– is the name of the column whose permissions you want to display. Use wildcard characters to request information for more than one column. If you do not specify a column name, permissions information for all columns in the specified table is returned.

Examples `sp_column_privileges discounts, null, null,
discounttype`

table_qualifier	table_owner	table_name	column_name
grantor	grantee	privilege	is_grantable
pubs2	dbo	discounts	discounttype
dbo	dbo	SELECT	YES
pubs2	dbo	discounts	discounttype
dbo	dbo	UPDATE	YES
pubs2	dbo	discounts	discounttype
dbo	dbo	REFERENCE	YES
pubs2	dbo	discounts	discounttype
dbo	guest	SELECT	NO
pubs2	dbo	discounts	discounttype
dbo	guest	UPDATE	NO
pubs2	dbo	discounts	discounttype
dbo	guest	REFERENCE	NO

Usage • Table 25-4 describes the results set:

Table 25-4: Results set for sp_column_privileges

Column	Datatype	Description
table_qualifier	varchar(32)	The name of the database in which the table specified for the <i>table_name</i> parameter is stored.
table_owner	varchar(32)	The table owner. If no value was specified for the <i>table_owner</i> parameter, this value is the current owner or the dbo.
table_name	varchar(32)	The name specified for the <i>table_name</i> parameter. This value cannot be NULL.
column_name	varchar(32)	The specified column name. If no column name was specified in the statement, the results include all columns in the specified table.
grantor	varchar(32)	The name of the database user who has granted permissions on <i>column_name</i> to <i>grantee</i> . This value cannot be NULL.
grantee	varchar(32)	The name of the database user who was granted permissions on <i>column_name</i> by <i>grantor</i> . This value cannot be NULL.
privilege	varchar(32)	Identifies the column privilege. May be one of the following: SELECT - The grantee is permitted to retrieve data for the column. UPDATE - The grantee is permitted to update data in the column. REFERENCE - The grantee is permitted to refer to the column within a constraint (for example, a unique, referential, or table check constraint).
is_grantable	varchar(3)	Indicates whether the grantee is permitted to grant the privilege to other users. The values are YES, NO, and NULL.

Permissions Any user can execute sp_column_privileges.

sp_columns

Description Returns information about the type of data that can be stored in one or more columns.

Syntax `sp_columns table_name [, table_owner]
 [, table_qualifier] [, column_name]`

Parameters *table_name*
 – is the name of the table or view. Use wildcard characters to request information about more than one table.

table_owner
 – is the owner of the table or view. Use wildcard characters to request information about tables owned by more than one user. If you do not specify a table owner, sp_columns looks for tables owned by the current user and then for tables owned by the Database Owner.

table_qualifier

– is the name of the database. This can be either the current database or NULL.

column_name

– is the name of the column for which you want information. Use wildcard characters to request information about more than one column.

Examples

```
sp_columns "publishers", null, null, "p%"
```

table_qualifier	table_owner	table_name	column_name	data_type	type_name	precision	length	scale	radix	nullable	remarks	ss_data_type	colid
pubs2	dbo	publishers	pub_id	1	char		4			0			
				NULL	NULL								
				NULL									
					47	1							
pubs2	dbo	publishers	pub_name	12	varchar		40			1			
				NULL	NULL								
				NULL									
					39	2							

Displays information about all columns in the publishers table that begin with “p”.

```
sp_columns "s%", null, null, "st%"
```

Displays information about all columns beginning with “st” in tables that begin with “s”.

Usage

- Table 25-5 shows the results set:

Table 25-5: Results set for sp_columns

Column	Datatype	Description
table_qualifier	varchar(32)	The name of the database in which the table specified for the <i>table_name</i> parameter is stored.
table_owner	varchar(32)	The table owner. If no value was specified for the <i>table_owner</i> parameter, this value is the current owner or the dbo.
table_name	varchar(32)	NOT NULL.
column_name	varchar(32)	NOT NULL.
data_type	smallint	Integer code for ODBC datatype. If this is a datatype that cannot be mapped into an ODBC type, it is NULL.
type_name	varchar(30)	String representing a datatype. The underlying DBMS presents this datatype name.
precision	int	Number of significant digits.
length	int	Length in bytes of a datatype.
scale	smallint	Number of digits to the right of the decimal point.
radix	smallint	Base for numeric datatypes.
nullable	smallint	The value 1 means NULL is possible; 0 means NOT NULL.
remarks	varchar(254)	
ss_data_type	smallint	An Adaptive Server datatype.
colid	tinyint	A column appended to the results set.
column_def	varchar(255)	
sql_data_type	smallint	
sql_datetime_sub	smallint	
char_octet_length	int	
ordinal_position	int	
is_nullable	varchar(3)	

- sp_columns reports the type_name as float, and data_type as 6 for columns defined as *double precision*. The Adaptive Server *double precision* datatype is a float implementation supports the range of values as specified in the ODBC specifications.

Permissions Any user can execute sp_columns.

sp_databases

Description Returns a list of databases in Adaptive Server.

Syntax sp_databases

Parameters None.

Examples

```

sp_databases

database_name      database_size      remarks
-----
master             5120              NULL
model              2048              NULL
mydb               2048              NULL
pubs2              2048              NULL
sybsecurity        5120              NULL
sybssystemprocs    16384             NULL
tempdb             2048              NULL

```

Usage

- Table 25-6 describes the results set:

Table 25-6: Results set for `sp_databases`

Column	Datatype	Description
database_name	char(32)	NOT NULL database name.
database_size	int	Size of database, in kilobytes.
remarks	varchar(254)	Adaptive Server always returns NULL.

Permissions Any user can execute `sp_databases`.

sp_datatype_info

Description Returns information about a particular ODBC datatype or about all ODBC datatypes.

Syntax `sp_datatype_info [data_type]`

Parameters *data_type*
 – is the code number for the specified ODBC datatype about which information is returned. Datatype codes are listed in Table 25-2 on page 420 and Table 25-3 on page 420 .

Usage

- Table 25-7 describes the results set:

Table 25-7: Results set for sp_datatype_info

Column	Datatype	Description
type_name	varchar(30)	A DBMS-dependent datatype name (the same as the type_name column in the sp_columns results set).
data_type	smallint	A code for the ODBC type to which all columns of this type are mapped.
precision	int	The maximum precision for the datatype on the data source. Zero is returned for datatypes where precision is not applicable.
literal_prefix	varchar(32)	Character(s) used to prefix a literal. For example, a single quotation mark (') for character types and 0x for binary.
literal_suffix	varchar(32)	Character(s) used to terminate a literal. For example, a single quotation mark (') for character types and nothing for binary.
create_params	varchar(32)	A description of the creation parameters for this datatype.
nullable	smallint	The value 1 means this datatype can be created allowing null values; 0 means it cannot.
case_sensitive	smallint	The value 1 means all columns of this type are case sensitive (for collations); 0 means they are not.
searchable	smallint	The value 1 means columns of this type can be used in a where clause.
unsigned_attribute	smallint	The value 1 means the datatype is unsigned; 0 means the datatype is signed.
money	smallint	The value 1 means it is a money datatype; 0 means it is not.
auto_increment	smallint	The value 1 means the datatype is automatically incremented; 0 means it is not.
local_type_name	varchar(128)	Localized version of the data source dependent name of the datatype.

Permissions Any user can execute sp_datatype_info.

sp_fkeys

Description Returns information about foreign key constraints created with the create table or alter table command in the current database.

Syntax `sp_fkeys pktable_name [, pktable_owner] [, pktable_qualifier] [, fktable_name] [, fktable_owner] [, fktable_qualifier]`

Parameters

pktable_name

– is the name of the primary key table. The use of wildcard characters in pattern matching is not supported. You must specify either the *pktable_name* or the *fktable_name*, or both.

pktable_owner

– is the name of the primary key table owner. The use of wildcard characters in pattern matching is not supported. If you do not specify the table owner, *sp_fkeys* looks for a table owned by the current user and then for a table owned by the Database Owner.

pktable_qualifier

– is the name of the database that contains the primary key table. This can be either the current database or NULL.

fktable_name

– is the name of the foreign key table. The use of wildcard characters in pattern matching is not supported. Either the *fktable_name* or the *pktable_name*, or both, must be given.

fktable_owner

– is the name of the foreign key table owner. The use of wildcard characters in pattern matching is not supported. If an *fktable_owner* is not specified, *sp_fkeys* looks for a table owned by the current user and then for a table owned by the Database Owner.

fktable_qualifier

– is the name of the database that contains the foreign key table. This can be either the current database or null.

Usage

- *sp_fkeys* returns information about foreign key constraints created with the *create table* or *alter table* command in the current database. A foreign key is a key column in a table that logically depends on a **primary key** column in another table.
- Table 25-8 describes the results set:

Table 25-8: Results set for *sp_fkeys*

Column	Datatype	Description
<i>pktable_qualifier</i>	<code>varchar(32)</code>	The database that contains the primary key table.
<i>pktable_owner</i>	<code>varchar(32)</code>	The owner of the primary key table.
<i>pktable_name</i>	<code>varchar(32)</code>	NOT NULL.
<i>pkcolumn_name</i>	<code>varchar(32)</code>	NOT NULL.
<i>fktable_qualifier</i>	<code>varchar(32)</code>	The database that contains the foreign key table.
<i>fktable_owner</i>	<code>varchar(32)</code>	The owner of the foreign key table.

Column	Datatype	Description
fktable_name	varchar(32)	NOT NULL.
fkcolumn_name	varchar(32)	NOT NULL.
key_seq	smallint	NOT NULL. The sequence number of the column in a multicolumn primary key.
update_rule	smallint	Action to be applied to the foreign key when the SQL operation is UPDATE. Zero is returned for this column.
delete_rule	smallint	Action to be applied to the foreign key when the SQL operation is DELETE. Zero is returned for this column.

- Both the primary key and foreign key must have been declared in a create table or alter table statement.
- If the primary key table name is supplied, but the foreign key table name is NULL, sp_fkeys returns all tables that include a foreign key to the given table. If the foreign key table name is supplied, but the primary key table name is NULL, sp_fkeys returns all tables that are related by a primary key/foreign key relationship to foreign keys in the foreign key table.
- sp_fkeys does not return information about keys declared with the sp_commonkey, sp_foreignkey or sp_primarykey system procedures.

Permissions

Any user can execute sp_fkeys.

sp_pkeys

Description

Returns information about primary key constraints created with the create table or alter table command for a single table.

Syntax

```
sp_pkeys table_name [, table_owner]
        [, table_qualifier]
```

Parameters

table_name

– is the name of the table. The use of wildcard characters in pattern matching is not supported.

table_owner

– is the name of the table owner. The use of wildcard characters in pattern matching is not supported. If *table_owner* is not specified, sp_pkeys looks for a table owned by the current user and then for a table owned by the Database Owner.

table_qualifier

– is the name of the database that contains the table. This can be either the current database or NULL.

Usage

- Table 25-9 describes the results set:

Table 25-9: Results set for *sp_pkeys*

Column	Datatype	Description
table_qualifier	varchar(32)	The database name. This field can be NULL.
table_owner	varchar(32)	The table owner. If no value was specified for the <i>table_owner</i> parameter, this value is the current owner or the dbo.
table_name	varchar(32)	NOT NULL.
column_name	varchar(32)	NOT NULL.
key_seq	smallint	NOT NULL. The sequence number of the column in a multicolumn primary key.

- Primary keys must have been declared with the *create table* or *alter table* statement, not with the *sp_primarykey* system procedure.
- The term *primary key* refers to a logical primary key for a table. Adaptive Server expects that every logical primary key has a unique index defined on it and that this unique index is also returned in *sp_statistics*.

Permissions

Any user can execute *sp_pkeys*.

sp_server_info

Description

Returns a list of Adaptive Server attribute names and current values.

Syntax

sp_server_info [*attribute_id*]

Parameters

attribute_id

– is the integer ID of the server attribute.

Examples

Example 1

```
sp_server_info 12
```

```
attribute_id attribute_name          attribute_value
```

```
-----
```

```
12 MAX_OWNER_NAME_LENGTH          0
```

Example 2

sp_server_info

Returns the list of server attributes, described by the mandatory rows, and their values.

Usage

- Table 25-10 describes the results set:

Table 25-10: Results set for sp_server_info

Column	Datatype	Description
attribute_id	int	NOT NULL.
attribute_name	varchar(60)	NOT NULL.
attribute_value	varchar(255)	

- Table 25-11 shows the mandatory rows in the results set:

Table 25-11: Mandatory results returned by sp_server_info

ID	Server Attribute Name	Description	Value
1	DBMS_NAME	Name of the DBMS.	SQL SERVER
2	DBMS_VER	Version of the DBMS.	@@version
6	DBE_NAME	Unused	
10	OWNER_TERM	Adaptive Server's term for a table owner (the second part of a three-part name).	owner
11	TABLE_TERM	Adaptive Server's term for a table (the third part of a three-part name).	table
12	MAX_OWNER_NAME_LENGTH	Maximum length of the name for a table owner (the second part of a three-part name).	30
13	TABLE_LENGTH	The maximum number of characters for a table name.	30
14	MAX_QUAL_LENGTH	Maximum length of the name for a table qualifier (the first part of a three-part table name).	30
15	COLUMN_LENGTH	The maximum number of characters for a column name.	30
16	IDENTIFIER_CASE	The case sensitivity of user-defined names (table names, column names, and stored procedure names) in the database (the case in which these objects are presented in the system catalogs).	MIXED
18	COLLATION_SEQ	The assumed ordering of the character set for this server.	
19	SAVEPOINT_SUPPORT	Does the underlying DBMS support named savepoints?	Y

ID	Server Attribute Name	Description	Value
20	MULTI_RESULT_SETS	Does the underlying DBMS or the gateway itself support multiple results sets (can multiple statements be sent through the gateway, with multiple results sets returned to the client)?	Y
22	ACCESSIBLE_TABLES	In the <code>sp_tables</code> stored procedure, does the gateway return only tables, views, and so on, that are accessible by the current user (that is, the user who has at least <code>select</code> privileges for the table)?	Y
100	USERID_LENGTH	The maximum number of characters for a user name.	30
101	QUALIFIER_TERM	Adaptive Server's term for a table qualifier (the first part of a three-part name).	database
102	NAMED_TRANSACTIONS	Does the underlying DBMS support named transactions?	Y
103	SPROC_AS_LANGUAGE	Can stored procedures be executed as language events?	Y
103	REMOTE_SPROC	Can stored procedures be executed through the remote stored procedure APIs in DB-Library?	Y
104	ACCESSIBLE_SPROC	In the <code>sp_stored_procedures</code> stored procedure, does the gateway return only stored procedures that are executable by the current user?	Y
105	MAX_INDEX_COLS	Maximum number of columns in an index for the DBMS.	32
106	RENAME_TABLE	Can tables be renamed?	Y
107	RENAME_COLUMN	Can columns be renamed?	Y
108	DROP_COLUMN	Can columns be dropped?	Y
109	INCREASE_COLUMN_LENGTH	Can column size be increased?	N
110	DDL_IN_TRANSACTION	Can DDL statements appear in transactions?	Y
111	DESCENDING_INDEXES	Are descending indexes supported?	Y
112	SP_RENAME	Can a stored procedure be renamed?	Y
500	SYS_SPROC_VERSION	The version of the catalog stored procedures currently implemented.	01.01.2822

Permissions

Any user can execute `sp_server_info`.

sp_special_columns

Description Returns the optimal set of columns that uniquely identify a row in a table or view; can also return a list of timestamp columns, whose values are automatically generated when any value in the row is updated by a transaction.

Syntax `sp_special_columns table_name [, table_owner] [, table_qualifier] [, col_type]`

Parameters

- table_name*
 - is the name of the table or view. The use of wildcard characters in pattern matching is not supported.
- table_owner*
 - is the name of the table or view owner. The use of wildcard characters in pattern matching is not supported. If you do not specify the table owner, sp_special_columns looks for a table owned by the current user and then for a table owned by the Database Owner.
- table_qualifier*
 - is the name of the database. This can be either the current database or NULL.
- col_type*
 - is R to return information about columns whose values uniquely identify any row in the table, or V to return information about timestamp columns, whose values are generated by Adaptive Server each time a row is inserted or updated.

Examples **Example 1**

```

                                sp_special_columns systypes
scope  column_name      data_type type_name      precision
      length          scale
-----
0 name          12 varchar          30
      30      NULL
    
```

Example 2

```

                                sp_special_columns @table_name=authors, @col_type=R
scope  column_name      data_type type_name      precision
      length          scale
-----
0 au_id          12 varchar          11
    
```


11 NULL

Usage

- Table 25-12 describes the results set:

Table 25-12: Results set for `sp_special_columns`

Column	Datatype	Description
scope	int	NOT NULL. Actual scope of the row ID. Adaptive Server always returns 0.
column_name	varchar(30)	NOT NULL. Column identifier.
data_type	smallint	The integer code for an ODBC datatype. If this datatype cannot be mapped to an ANSI/ISO type, the value is NULL. The native datatype name is returned in the <code>type_name</code> column. (See the ODBC datatypes Table 25-2.)
type_name	varchar(13)	The string representation of the datatype. This is the datatype name as presented by the underlying DBMS.
precision	int	The number of significant digits.
length	int	The length in bytes of the datatype.
scale	smallint	The number of digits to the right of the decimal point.

Permissions

Any user can execute `sp_special_columns`.

sp_sproc_columns

Description

Returns information about a stored procedure's input and return parameters.

Syntax

```
sp_sproc_columns procedure_name [, procedure_owner]
                [, procedure_qualifier] [, column_name]
```

Parameters

procedure_name

– is the name of the stored procedure. The use of wildcard characters in pattern matching is not supported.

procedure_owner

– is the owner of the stored procedure. The use of wildcard characters in pattern matching is not supported. If no owner is specified, `sp_sproc_columns` returns all columns.

procedure_qualifier

– is the name of the database. This can be either the current database or NULL.

column_name

– is the name of the parameter about which you want information. If you do not supply a parameter name, sp_sproc_columns returns information about all input and return parameters for the stored procedure.

Usage

- Table 25-13 describes the results set:

Table 25-13: Results set for sp_sproc_columns

Column	Datatype	Description
procedure_qualifier	varchar(30)	
procedure_owner	varchar(30)	
procedure_name	varchar(41)	NOT NULL.
column_name	varchar(30)	NOT NULL.
column_type	smallint	
data_type	smallint	The integer code for an ODBC datatype. If this datatype cannot be mapped to an ANSI/ISO type, the value is NULL. The native datatype name is returned in the type_name column.
type_name	char(30)	The string representation of the datatype. This is the datatype name as presented by the underlying DBMS.
precision	int	The number of significant digits.
length	int	The length in bytes of the datatype.
scale	smallint	The number of digits to the right of the decimal point.
radix	smallint	Base for numeric types.
nullable	smallint	The value 1 means this datatype can be created allowing null values; 0 means it cannot.
remarks	varchar(254)	NULL.
ss_data_type	tinyint	An Adaptive Server datatype.
colid	tinyint	An Adaptive Server specific column appended to the result set.

- sp_sproc_columns reports the type_name as float, and data_type as 6 for parameters defined as *double precision*. The Adaptive Server *double precision* datatype is a float implementation supports the range of values as specified in the ODBC specifications.

Permissions

Any user can execute sp_sproc_columns.

sp_statistics

Description Returns a list of indexes on a single table.

Syntax `sp_statistics table_name [, table_owner] [, table_qualifier] [, index_name] [, is_unique]`

Parameters

table_name
– is the name of the table. The use of wildcard character pattern matching is not supported.

table_owner
– is the owner of the table. The use of wildcard character pattern matching is not supported. If *table_owner* is not specified, `sp_statistics` looks for a table owned by the current user and then for a table owned by the Database Owner.

table_qualifier
– is the name of the database. This can be either the current database or NULL.

index_name
– is the index name. The use of wildcard character pattern matching is not supported.

is_unique
– is Y to return only unique indexes; otherwise, is N to return both unique and nonunique indexes.

Examples `sp_statistics publishers`

```

table_qualifier          table_owner
table_name              non_unique
index_qualifier         index_name
type    seq_in_index  column_name                collation
cardinality pages
-----
pubs2
publishers              NULL
NULL                   NULL
      0      NULL NULL                NULL
      3      1
pubs2
publishers              NULL
publishers              pubind

```

1 3 1 pub_id A

Usage

- Table 25-14 describes the results set:

Table 25-14: Results set for sp_statistics

Column	Datatype	Description
<i>table_qualifier</i>	varchar(32)	The database name. This field can be NULL.
<i>table_owner</i>	varchar(32)	
<i>table_name</i>	varchar(32)	NOT NULL.
<i>non_unique</i>	smallint	NOT NULL. The value 0 means unique, and 1 means not unique.
<i>index_qualifier</i>	varchar(32)	
<i>index_name</i>	varchar(32)	
<i>type</i>	smallint	NOT NULL. The value 0 means clustered, 2 means hashed, and 3 means other.
<i>seq_in_index</i>	smallint	NOT NULL.
<i>column_name</i>	varchar(32)	NOT NULL.
<i>collation</i>	char(1)	The value A means ascending; D means descending; and NULL means not applicable.
<i>cardinality</i>	int	Number of rows in the table or unique values in the index.
<i>pages</i>	int	Number of pages to store the index or table.

- The indexes in the results set appear in ascending order, ordered by the non-unique, type, index_name, and seq_in_index columns.
- The index type hashed accepts exact match or range searches, but searches involving pattern matching do not use the index.

Permissions

Any user can execute sp_statistics.

sp_stored_procedures

Description

Returns information about one or more stored procedures.

Syntax

sp_stored_procedures [*sp_name* [, *sp_owner* [, *sp_qualifier*]]]

Parameters

sp_name
 – is the name of the stored procedure. Use wildcard characters to request information about more than one stored procedure.

sp_owner

– is the owner of the stored procedure. Use wildcard characters to request information about procedures that are owned by more than one user.

sp_qualifier

– is the name of the database. This can be the current database or NULL.

Usage

- `sp_stored_procedures` returns information about stored procedures in the current database only.
- Table 25-15 shows the results set:

Table 25-15: Results set for `sp_stored_procedures`

Column	Datatype	Description
<code>procedure_qualifier</code>	<code>varchar(30)</code>	The name of the database.
<code>procedure_owner</code>	<code>varchar(30)</code>	
<code>procedure_name</code>	<code>varchar(41)</code>	NOT NULL.
<code>num_input_params</code>	<code>int</code>	NOT NULL. Always returns -1.
<code>num_output_params</code>	<code>int</code>	NOT NULL. The value ≥ 0 shows the number of parameters; -1 means the number of parameters is indeterminate.
<code>num_result_sets</code>	<code>int</code>	NOT NULL. Always returns -1.
<code>remarks</code>	<code>varchar(254)</code>	NULL.

- `sp_stored_procedures` can return the name of stored procedures for which the current user does not have execute permission. However, if the server attribute `accessible_sproc` is “Y” in the results set for `sp_server_info`, only stored procedures that are executable by the current user are returned.

Permissions

Any user can execute `sp_stored_procedures`.

sp_table_privileges

Description

Returns privilege information for all columns in a table or view.

Syntax

```
sp_table_privileges table_name [, table_owner
  [, table_qualifier]
```

Parameters

table_name

– is the name of the table. The use of wildcard characters in pattern matching is not supported.

table_owner

– is the name of the table owner. The use of wildcard characters in pattern matching is not supported. If you do not specify the table owner, sp_table_privileges looks for a table owned by the current user and then for a table owned by the Database Owner.

table_qualifier

– is the name of the database. This can be either the current database or NULL.

Usage

- Table 25-16 shows the results set:

Table 25-16: Results set for sp_table_privileges

Column	Datatype	Description
table_qualifier	varchar(32)	The name of the database. This field can be NULL.
table_owner	varchar(32)	
table_name	varchar(32)	NOT NULL.
grantor	varchar(32)	NOT NULL.
grantee	varchar(32)	NOT NULL.
privilege	varchar(32)	Identifies the table privilege. May be one of the following: <ul style="list-style-type: none"> • SELECT - The grantee is permitted to retrieve data for one or more columns of the table. • INSERT - The grantee is permitted to insert new rows containing data for one or more columns into the table. • UPDATE - The grantee is permitted to update the data in one or more columns of the table. • DELETE - The grantee is permitted to delete rows of data from the table. • REFERENCE - The grantee is permitted to refer to one or more columns of the table within a constraint.
is_grantable	varchar(3)	Indicates whether the grantee is permitted to grant the privilege to other users. The values are YES, NO, and NULL.

Permissions

Any user can execute sp_table_privileges.

sp_tables

Description	Returns a list of objects that can appear in a from clause.
Syntax	<code>sp_tables [table_name] [, table_owner] [, table_qualifier][, table_type]</code>
Parameters	<p><i>table_name</i> – is the name of the table. Use wildcard characters to request information about more than one table.</p> <p><i>table_owner</i> – is the table owner. Use wildcard characters to request information about more than one table.</p> <p><i>table_qualifier</i> – is the name of the database. Acceptable values are the name of the current database and NULL.</p> <p><i>table_type</i> – is a list of values, separated by commas, giving information about all tables of the table type(s) specified, including the following: " 'TABLE' , 'SYSTEM TABLE' , 'VIEW' "</p>

Note Enclose each table type with single quotation marks, and enclose the entire parameter with double quotation marks. Enter table types in uppercase.

Examples

```
sp_tables @table_type = "'TABLE' , 'VIEW' "
```

This procedure returns information about all tables in the current database of the type TABLE and VIEW and excludes information about system tables.

- Usage
- Adaptive Server does not necessarily check the read and write permissions on *table_name*. Access to the table is not guaranteed, even if you can display information about it.
 - The results set includes tables, views, and synonyms and aliases for gateways to DBMS products.
 - If the server attribute `accessible_tables` is “Y” in the results set for `sp_server_info`, only tables that are accessible by the current user are returned.
 - Table 25-17 describes the results set:

Table 25-17: Results set for sp_tables

Column	Datatype	Description
<i>table_qualifier</i>	varchar(30)	The database name. This field can be NULL.
<i>table_owner</i>	varchar(30)	
<i>table_name</i>	varchar(30)	NOT NULL. The table name.
<i>table_type</i>	varchar(32)	NOT NULL. One of the following: 'TABLE', 'VIEW', 'SYSTEM TABLE'.
<i>remarks</i>	varchar(254)	NULL

Permissions

Any user can execute sp_tables.

System Extended Stored Procedures

This chapter describes the system extended stored procedures (ESPs), which are supplied by Sybase. ESPs are created by installmaster at installation. They are located in the sybssystemprocs database and owned by the System Administrator. They can be run from any database.

Table 26-1 lists the system extended stored procedures discussed in this chapter.

Table 26-1: System extended stored procedures

Procedure	Description	Platform
xp_cmdshell	Executes a native operating system command on the host system running Adaptive Server.	All Supporting DLLs
xp_deletemail	Deletes a message from the Adaptive Server message inbox.	NT Only
xp_enumgroups	Displays groups for a specific Windows NT domain.	NT Only
xp_findnextmsg	Retrieves the message identifier of the next message in the Adaptive Server message inbox.	NT Only
xp_logevent	Provides for logging a user-defined event in the Windows NT Event Log.	NT Only
xp_readmail	Reads a message from the Adaptive Server message inbox.	NT Only
xp_sendmail	Sends a message to the specified recipients using the MAPI interface.	NT Only
xp_startmail	Starts an Adaptive Server mail session.	NT Only
xp_stopmail	Stops an Adaptive Server mail session.	NT Only

Permissions on system ESPs

Permissions are set in the sybssystemprocs database.

Users with the `sa_role` have default execution permissions on the system ESPs. These System Administrators can grant execution permissions to other users.

DLLs associated with system ESPs

You can get the names of the DLLs associated with the system ESPs by running `sp_helpextendedproc` in the `sysystemprocs` database.

Using system ESPs

The system ESPs follow the same calling conventions as the regular system procedures. The only additional requirement for system ESPs is that the Open Server application, XP Server, must be running. Adaptive Server starts XP Server the first time an ESP is invoked. XP Server continues to run until you shut down Adaptive Server.

xp_cmdshell

Description	Executes a native operating system command on the host system running Adaptive Server.
Syntax	<code>xp_cmdshell <i>command</i> [, no_output]</code>
Parameters	<i>command</i> – is the operating system command string; maximum length is 255 bytes. <code>no_output</code> – if specified, suppresses any output from the command.
Examples	Example 1 <code>xp_cmdshell 'copy C:\log A:\log.0102', no_output</code> Silently copies the file named <code>log</code> on the C drive to a file named <code>log.0102</code> on the A drive.

Example 2

```
xp_cmdshell 'date'
```

Executes the operating system's date command and returns the current date as a row of data.

Usage

- xp_cmdshell returns any output, including operating system errors, as rows of text in a single column.
- xp_cmdshell is run from the current directory of the XP Server.
- The width of the column of returned output is 80 characters. The output is not formatted.
- xp_cmdshell cannot perform commands that require interaction with the user, such as "login".
- The user context in which an operating system command is executed via xp_cmdshell is controlled by the value of the xp_cmdshell context configuration parameter. If this parameter is set to 1 (the default), xp_cmdshell restricts permission to users with System Administration privileges at the operating system level. If this parameter is set to 0, xp_cmdshell uses the security context of the operating system account under which Adaptive Server is running. Therefore, using xp_cmdshell with the xp_cmdshell context configuration parameter set to 0, any user can execute operating system commands using the permissions of the account running Adaptive Server. This account may have fewer restrictions than the user's own account.

For more information about the xp_cmdshell context, see the System Administration Guide

- Regardless of the value of xp_cmdshell context, if the user who is executing xp_cmdshell is not a System Administrator (does not have the sa_role), a System Administrator must have granted that user explicit permission to execute xp_cmdshell. For example, the following statement grants "joe" permission to execute xp_cmdshell:

```
grant execute on xp_cmdshell to joe
```

Permissions

By default, only a System Administrator can execute xp_cmdshell. A System Administrator can grant execute permission to other users.

xp_deletemail

Description	<i>Windows NT only</i> – Deletes a message from the Adaptive Server message inbox.
Syntax	xp_deletemail [<i>msg_id</i>]
Parameters	<i>msg_id</i> – is the message identifier of the mail message to be deleted.
Examples	Example 1 <pre>1> declare @cur_msg_id binary(255) 2> exec xp_deletemail @msg_id = @cur_msg_id</pre> <p>Deletes from the Adaptive Server message inbox the message with the message identifier specified in the <i>cur_msg_id</i> variable.</p> Example 2 <pre>xp_deletemail</pre> <p>Deletes the first message from the Adaptive Server message inbox.</p>
Usage	<ul style="list-style-type: none">• Obtain the <i>msg_id</i> using xp_findnextmsg.• If the <i>msg_id</i> parameter is not used, the message to be deleted defaults to the first message in the message inbox.
Permissions	By default, only a System Administrator can execute xp_deletemail. A System Administrator can grant this permission to other users.

xp_enumgroups

Description	<i>Windows NT only</i> – Displays groups for a specified Windows NT domain.
Syntax	xp_enumgroups [<i>domain_name</i>]
Parameters	<i>domain_name</i> – is the Windows NT domain for which you are listing user groups.
Examples	Example 1 <pre>xp_enumgroups</pre> <p>Lists all user groups on the Windows NT computer running XP Server.</p> Example 2 <pre>xp_enumgroups 'PCS'</pre>

	Lists all user groups in the PCS domain.
Usage	<ul style="list-style-type: none"> • <code>xp_enumgroups</code> displays all local user groups if no parameter is passed. • A <i>domain</i> is a named collection of computers that share a common user account database and security policy. • A return status of 0 indicates success; 1 indicates failure.
Permissions	By default, only a System Administrator can execute <code>xp_enumgroups</code> . A System Administrator can grant this permission to other users.

xp_findnextmsg

Description	<i>Windows NT only</i> – Retrieves the next message identifier from the Adaptive Server message inbox.
Syntax	<code>xp_findnextmsg @msg_id = @msg_id output [, type]</code> <code>[, unread_only = {true false}]</code>
Parameters	<p><i>msg_id</i></p> <ul style="list-style-type: none"> – on input, specifies the message identifier that immediately precedes the one you are trying to retrieve. Places the retrieved message identifier in the <i>msg_id</i> output parameter, which must be of type binary. <p><i>type</i></p> <ul style="list-style-type: none"> – is the input message type based on the MAPI mail definition. The only supported message type is CMC:IPM. A NULL value or no value defaults to CMC:IPM. <p><i>unread_only</i></p> <ul style="list-style-type: none"> – if this parameter is set to true, <code>xp_findnextmsg</code> considers only unread messages. If this parameter is set to false, <code>xp_findnextmsg</code> considers all messages, both read and unread, when retrieving the next message identifier. The default is true.

Examples

Example 1

```
xp_findnextmsg @msg_id = @out_msg_id output
```

Returns, in the `@out_msg_id` output variable, the message identifier of the next unread message after the message specified by the `@out_msg_id`.

Example 2

```
xp_findnextmsg @msg_id = @out_msg_id output, NULL,
```

```
@unread_only = false
```

Returns, in the `@out_msg_id` output variable, the message identifier of the next message after the message specified by the `@out_msg_id`. The message may be read or unread.

Usage

- When `xp_findnextmsg` can find no more messages in the inbox, it returns a status of 1.
- `xp_deletemail` and `xp_readmail` use the message identifier returned by `xp_findnextmsg`.

Permissions

By default, only a System Administrator can execute `xp_findnextmsg`. A System Administrator can grant this permission to other users.

xp_logevent

Description

Windows NT only – Provides for logging a user-defined event in the Windows NT Event Log from within Adaptive Server.

Syntax

```
xp_logevent error_number, message [, type]
```

Parameters

error_number

– is the user-assigned error number. It must be equal to or greater than 50000.

message

– is the text of the message that is displayed in the description field of the event viewer. The maximum length of the message is 255 bytes. Enclose the message in quotes.

type

– describes the urgency of the event. Values are informational, warning, and error. The default is informational. Enclose the value in quotes.

Examples

Example 1

```
xp_logevent 55555, 'Email message deleted.'
```

An informational event, number 55555, will be logged in the Windows NT Event Log. The text of the description in the event detail window is “Email message deleted”.

Example 2

```
xp_logevent 66666, 'DLL not found.', 'error'
```

An error event, number 66666, will be logged in the Windows NT Event Log. The text of the description in the event detail window is “DLL not found”.

Usage

- The following table describes the default event details for events generated with `xp_logevent`:

Detail	Value
User	N/A
Computer	Name of machine running XP Server
Event ID	12
Source	Name of Adaptive Server
Category	User

Permissions

Only a System Administrator can execute `xp_logevent`.

xp_readmail

Description

Windows NT only – Reads a message from the Adaptive Server message inbox.

Syntax

```
xp_readmail [msg_id]
            [, recipients output]
            [, sender output]
            [, date_received output]
            [, subject output]
            [, cc output]
            [, message output]
            [, attachments output]
            [, suppress_attach = {true | false}]
            [, peek = {true | false}]
            [, unread = {true | false}]
            [, msg_length output]
            [, bytes_to_skip [output]]
            [, type [output]]
```

Parameters

msg_id

– specifies the message identifier of the message to be read by `xp_readmail`. If the *msg_id* parameter is not used, the message defaults to the first unread message in the message box, if `unread` is true, or to the first message in the message box, if `unread` is false.

recipients

– is a semicolon-separated list of the recipients of the message.

sender

– is the originator of the message.

date_received

– is the date the message was received.

subject

– is the subject header of the message.

cc

– is a list of the message's copied (cc'd) recipients (separated by semicolons).

message

– is the text of the message body. If the length of the message body, obtained from the *msg_length* output parameter, is greater than 255, use the *byte_to_skip* and *msg_length* parameters to read the message in 255-byte increments.

attachments

– is a list of the temporary paths of the attachments (separated by semicolons). *attachments* is ignored if *suppress_attach* is true.

suppress_attach

– if set to true, prevents the creation of temporary files for attachments. The default is true.

peek

– if set to false, flags the message as unread after it has been read. If set to true, flags the message as an unread message, even after it has been read. The default is false.

unread_only

– if set to true, *xp_readmail* considers only unread messages. If set to false, *xp_readmail* considers all messages, whether they are flagged as read or unread. The default is true.

msg_length

– is the total length of the message, in bytes. Used with the *bytes_to_skip* parameter, allows *xp_readmail* to read messages in 255-byte increments.

bytes_to_skip

– on input, if not 0, specifies the number of bytes to skip before reading the next 255 bytes of the message into the message output parameter. On output, contains the offset in the message (the previous value of *bytes_to_skip* plus the *msg_length* that is output with the call) from which to start reading the next 255-byte increment.

type

– is the message type based on the MAPI mail definition. The only supported message type is CMC:IPM. A NULL value or no value defaults to CMC:IPM.

Examples

Example 1

```
declare @msgid binary(255)
declare @originator varchar(20)
declare @mess varchar(255)
exec xp_findnextmsg @msg_id = @msgid output
exec xp_readmail @msg_id = @msgid,
@sender = @originator output,
@message = @mess output
```

`xp_readmail` reads the first unread message in the message inbox. It gets the message identifier for this message from the *@msgid* variable, where it has been stored by the `xp_findnextmsg` ESP. `xp_readmail` stores the sender's name in the *@originator* variable and the message body in the *@mess* variable.

Example 2

```
declare @msgid binary(255)
declare @mess varchar(255)
declare @msg_length char(255)
declare @len int
declare @skip int
exec xp_findnextmsg @msgid output
exec xp_readmail @msg_id = @msgid,
@message = @mess output,
@msg_length = @len output,
@bytes_to_skip = @skip output
print @mess
if (@len > 255)
begin
    while (@skip < @len)
    begin
        xp_readmail @msg_id = @msgid,
        @message = @mess output,
        @bytes_to_skip = @skip output
```

```
        print @mess
    end
end
```

Reads the first 255 bytes of the message for which the message identifier is output by xp_findnextmsg. If the total length of the message exceeds 255 bytes, reads the next 255 bytes and continues until there are no more bytes to read.

Usage	<ul style="list-style-type: none">• xp_readmail reads a message from the Adaptive Server message inbox.• To get the message identifier of the next message in the message inbox, use xp_findnextmsg.
Permissions	By default, only a System Administrator can execute xp_readmail. A System Administrator can grant this permission to other users.

xp_sendmail

Description *Windows NT only* – Sends a message to the specified recipients. The message is either text or the results of a Transact-SQL query.

Syntax

```
xp_sendmail recipient [, recipient] . . .
    [, subject]
    [, cc_recipient] . . .
    [, bcc_recipient] . . .
    [, {query | message}]
    [, attachname]
    [, attach_result = {true | false}]
    [, echo_error = {true | false}]
    [, include_file [, include_file] . . .]
    [, no_column_header = {true | false}]
    [, no_output = {true | false}]
    [, width]
    [, separator]
    [, dbuser]
    [, dbname]
    [, type]
    [, include_query = {true | false}]
```

Parameters *recipient*

- is the email address of the user who will receive the message. At least one recipient is required. Separate multiple recipients with semicolons.

subject

– is the optional message subject header. If not used, defaults to “Sybase SQL Server Message”.

cc_recipient

– is a list of the message’s copied (cc’d) recipients (separated by semicolons).

bcc_recipient

– is the list of the message’s blind- copied (bcc’d) recipients (separated by semicolons).

query

– is one or more Transact-SQL statements. The results are sent to the recipients of the message. If *query* is used, *message* cannot be used.

message

– is the text of the message being sent. If *message* is used, *query* cannot be used. For the complete list of options that are ignored when you use *message*, see the “Usage” section.

attachname

– is the name of the file containing the results of a query, which is included as an attachment to the message, when the *query* parameter is used. If *attachname* is used, *attach_result* must be set to true. If *attach_result* is true and *attachname* is not specified, the prefix of the attached file’s generated file name is “syb” followed by 5 random digits followed by the “.txt” extension, for example, *syb84840.txt*. This parameter is ignored if the *message* parameter is used.

attach_result

– if set to true, sends the results of a query as an attachment to the message. If set to false, sends the results directly in the message body. The default is false. This parameter is ignored if the *message* parameter is used.

echo_error

– if set to true, sends Adaptive Server messages, including the count of rows affected message, along with the query results. If set to false, does not send Adaptive Server messages. The default is true. This parameter is ignored if the *message* parameter is used.

include_file

– is a list of files to be included as attachments to the message, separated by semicolons. The files can be specified as file names, path names, or relative path names and can be either text or binary files.

no_column_header

– if set to true, column headers are sent with query results. If set to false, column headers are not sent. The default is false. This parameter is ignored if the *message* parameter is used.

no_output

– if set to true, no output is sent to the session that sent the mail. If set to false, the session sending the mail receives output. The default is false. This parameter is ignored if the *message* parameter is used.

width

– specifies, in characters, the width of the results sets when query results are sent in a message. *width* is the same as the /w option in isql. Result rows are broken by the newline character when the specified *width* is reached. The default is 80 characters. This parameter is ignored if the *message* parameter is used.

separator

– specifies the character to be used as a column separator when query results are sent in a message. *separator* is the same as the /s option in isql. The default is the tab character. This parameter is ignored if the *message* parameter is used.

dbuser

– specifies the database user name to be assumed for the user context for executing queries when the *query* parameter is used. The default is “guest.” This parameter is ignored if the *message* parameter is used.

dname

– specifies the database name to be assumed for the database context for executing queries when the *query* parameter is used. The default is “master.” This parameter is ignored if the *message* parameter is used.

type

– is the input message type based on the MAPI mail definition. The only supported message type is CMC:IPM. A NULL value or no value defaults to CMC:IPM.

include_query

– if set to true, the query or queries used in the *query* parameter are appended to the results set. If set to false, the query is not appended. The default is false. *include_query* is ignored if the *message* parameter is used.

Examples

Example 1

```
xp_sendmail @recipient = "sally;ramon",
```

```
@subject = "Adaptive Server Backup Status",
@message = "Adaptive Server Backup for SERVER2 is
complete.",
@copy_recipient="admin"
```

xp_sendmail sends a text message on the backup status of an Adaptive Server to “sally” and “ramon” with a copy to the “admin” group.

Example 2

```
xp_sendmail "peter",
@query = "select * from authors",
@attachname = "au_list.res",
@attach_result= true
```

Sends “peter” the results of a query on the *authors* table. The results are in an attachment to the message, which consists of a file named *au_lis.res*, which is in the directory from which the server is being executed.

Usage

- The following parameters are related to the results of queries sent in a message when the *query* parameter is used. They are ignored if the *message* parameter is used instead: *attachname*, *attach_result*, *echo_error*, *no_column_header*, *no_output*, *width*, *separator*, *dbuser*, *dname*, *include_query*.

Permissions

By default, only a System Administrator can execute xp_sendmail. A System Administrator can grant this permission to other users.

xp_startmail

Description

Windows NT only – Starts an Adaptive Server mail session.

Syntax

```
xp_startmail [mail_user ] [, mail_password]
```

Parameters

mail_user

– is a mail profile name used by Adaptive Server to log into the Windows NT mail system. If *mail_user* is not used, xp_startmail uses the mail user name that was used to set up Sybmail’s Adaptive Server account.

mail_password

– is the mail password used by Adaptive Server to log into the Windows NT mail system. If *mail_password* is not used, xp_startmail uses the mail password that was used to set up Sybmail’s Adaptive Server account.

Examples

Example 1

```
xp_startmail
```

Starts an Adaptive Server mail session using the mail user name and password for Sybmail's user account.

Example 2

```
xp_startmail "mailuser", "tre55uu"
```

Starts an Adaptive Server mail session with "mailuser" as the profile name and the password associated with that profile name.

Usage

- xp_startmail will not start an Adaptive Server mail session if one is already running.
- An Adaptive Server mail session must be started, either by an explicit call to xp_startmail or by configuring Adaptive Server to start an Adaptive Server mail session automatically at start-up, before any Sybmail-related system ESPs or the sp_processmail stored procedure can be executed. For information about initiating an Adaptive Server mail session automatically at start-up, see start mail session in the System Administration Guide.
- When the Windows NT automail session is not on, you must use the mail_user and mail_password parameters with xp_startmail.
- To see the default mail_user value from the fullname field for the "sybmail" user account, use the sp_displaylogin system procedure as follows:

```
sp_displaylogin sybmail
```

Permissions

By default, only a System Administrator can execute xp_startmail. A System Administrator can grant this permission to other users.

xp_stopmail

Description

Windows NT only – Stops an Adaptive Server mail session.

Syntax

```
xp_stopmail
```

Parameters

None

Examples

```
xp_stopmail
```

	Stops an Adaptive Server mail session.
Usage	<ul style="list-style-type: none">• Sybmail-related system ESPs and the <code>sp_processmail</code> stored procedure cannot be executed after an Adaptive Server mail session has been terminated with <code>xp_stopmail</code>.
Permissions	By default, only a System Administrator can execute <code>xp_stopmail</code> . A System Administrator can grant this permission to other users.

This chapter describes the dbcc stored procedures. These procedures access the tables only in the dbccdb database or in the alternate database, dbccalt. For details on setting up dbccdb or dbccalt, see the *System Administration Guide*. For information on the tables used in these databases, see Chapter 8, “dbccdb Tables,” in the *Reference Manual Volume 4 System Tables*.

Table 27-1 lists the dbcc stored procedures described in this chapter. For details on the dbcc system procedure sp_plan_dbccdb, see sp_plan_dbccdb. For more information on this system procedure and the dbcc stored procedures, see the *System Administration Guide*.

Table 27-1: dbcc stored procedures

Procedure Name	Description
sp_dbcc_alterws	Changes the size of the specified workspace to a specified value, and initializes the workspace.
sp_dbcc_configreport	Generates a report that describes the configuration information used by the dbcc checkstorage operation for the specified database.
sp_dbcc_createws	Creates a workspace of the specified type and size on the specified segment and database.
sp_dbcc_deletedb	Deletes from dbccdb all the information related to the specified target database.
sp_dbcc_deletehistory	Deletes the results of dbcc checkstorage operations performed on the target database before the specified date and time.
sp_dbcc_differentialreport	Generates a report that highlights the changes in I/O statistics and faults that took place between two dbcc operations
sp_dbcc_evaluatedb	Recomputes configuration information for the target database and compares it to the current configuration information.
sp_dbcc_faultreport	Generates a report covering fault statistics for the dbcc checkstorage operations performed for the specified object in the target database on the specified date.
sp_dbcc_fullreport	Runs sp_dbcc_summaryreport, sp_dbcc_configreport, sp_dbcc_statisticsreport, and sp_dbcc_faultreport.
sp_dbcc_runcheck	Runs dbcc checkstorage on the specified database, then runs sp_dbcc_summaryreport or a report you specify.
sp_dbcc_statisticsreport	Generates an allocation statistics report on the specified object in the target database.
sp_dbcc_summaryreport	Generates a summary report on the specified database.

Procedure Name	Description
sp_dbcc_updateconfig	Updates the dbcc_config table in dbccdb with the configuration information of the target database.

Specifying the Object Name and Date

Several dbcc stored procedures use parameters for the object name and date. This section provides important information on specifying the object name and date.

Specifying the Object Name

The object name specifies only the name of the table or index for which to generate a report. When you specify an object name, you must also specify a database name (*dbname*). You cannot specify an owner for the object. If the specified object name is not unique in the target database, the system procedure generates a report on all objects with the specified name.

Specifying the Date

Use the following syntax to specify the date and time (optional):

mm/dd/yy[:hh:mm:ss]

A 24-hour clock is assumed.

When you specify the date, the system procedures interpret it as follows:

- If both the date and the time are specified, the dbcc operation that completed at the specified date and time is selected for the report.
- If the specified date is the current date, and no time is specified, the time is automatically set to the current time. The dbcc operation that completed within the previous 24 hours with a finish time closest to the current time is selected for the report.
- If the specified date is not the current date, and no time is specified, the time is automatically set to “23:59:59”. The dbcc checkstorage operation that completed with a finish date and time closest to the specified date and system-supplied time is selected for the report.

For example, suppose the most recent `dbcc checkstorage` operation completed on March 4, 1997 at 10:20:45.

If you specify the date as “03/04/97”, the system procedure interprets the date as 03/04/97:23:59:59. This date and time are compared to the actual finish date and time of 03/04/97:10:20:45.

If you specify the date as “03/04/97:10:00:00”, the operation that completes at 10:20:45 is not selected for the report because only the operations that complete on or before the specified time meet the criteria.

If you specify the date as “03/06/97”, no report is generated because the most recent operation completed more than 24 hours earlier.

sp_dbcc_alterws

Description	Changes the size of the specified workspace to a specified value, and initializes the workspace.
Syntax	<code>sp_dbcc_alterws <i>dbname</i>, <i>wsname</i>, "<i>wssize</i>[K M]"</code>
Parameters	<p><i>dbname</i></p> <ul style="list-style-type: none">– is the name of the database in which the workspace resides. Specify either <code>dbccdb</code> and <code>dbccalt</code>. <p><i>wsname</i></p> <ul style="list-style-type: none">– specifies the name of the workspace to alter. <p><i>wssize</i></p> <ul style="list-style-type: none">– is the new size of the workspace, specified by K (kilobytes) or M (megabytes). If you do not specify K or M, <i>wssize</i> specifies the number of pages. Page size is platform-dependent. The minimum size for a workspace is 24 pages.
Examples	<pre>sp_dbcc_alterws dbccdb, scan_ws_000001, "30M"</pre> <p>Workspace <code>scan_ws_000001</code> has been altered successfully to size 30MB.</p> <p>Changes the size of the <code>scan_ws_000001</code> workspace on <code>dbccdb</code> to 30MB.</p>
Usage	<ul style="list-style-type: none">• <code>sp_dbcc_alterws</code> changes the size of the specified workspace to the specified value and initializes the workspace.• To achieve maximum performance, make sure you have configured a buffer pool of at least 16K before you alter a workspace.

- Use `sp_plan_dbccdb` to determine size estimates before altering the workspace.
- The workspace must exist before it can be altered. For information on creating workspaces, see `sp_dbcc_createws`.
- To delete a workspace, in `dbccdb` issue:


```
drop table workspace_name
```
- For more information on the scan and text workspaces, and the `dbccalt` database, see the System Administration Guide.

Permissions Only a System Administrator or the Database Owner can run `sp_dbcc_alterws`.

See also *Commands* – `dbcc`
dbcc stored procedures – `sp_dbcc_createws`, `sp_dbcc_evaluatedb`
System procedures – `sp_plan_dbccdb`, `sp_helpdb`

sp_dbcc_configreport

Description Generates a report that describes the configuration information used by the `dbcc checkstorage` operation for the specified database.

Syntax `sp_dbcc_configreport [dbname]`

Parameters *dbname*
 – specifies the name of the database. If *dbname* is not specified, the report contains information on all databases in `dbccdb..dbcc_operation_log`.

Examples `sp_dbcc_configreport`

Reporting configuration information of database `sybssystemprocs`.

Parameter Name	Value	Size
database name	sybssystemprocs	51200K
dbcc named cache	default data cache	1024K
text workspace	textws_001 (id = 544004969)	128K
scan workspace	scanws_001 (id = 512004855)	1024K
max worker processes	1	
operation sequence number	2	

	Generates a report on the configuration information related to dbcc for the sybssystemprocs database. The “Value” column lists the object name, where applicable, and the size.
Usage	<ul style="list-style-type: none">• <code>sp_dbcc_configreport</code> generates a report that describes the configuration information used by dbcc operations for the specified database. This information is stored in the <code>dbcc_config</code> table.• To evaluate the most current configuration parameters, run <code>sp_dbcc_updateconfig</code> before running <code>sp_dbcc_configreport</code>.• To change the configuration values for a workspace, use <code>sp_dbcc_alterws</code>.
Permissions	Any user can run <code>sp_dbcc_configreport</code> .
See also	<i>Commands</i> – dbcc <i>dbcc stored procedures</i> – <code>sp_dbcc_alterws</code> , <code>sp_dbcc_fullreport</code> , <code>sp_dbcc_statisticsreport</code> , <code>sp_dbcc_summaryreport</code> , <code>sp_dbcc_updateconfig</code>

sp_dbcc_createws

Description	Creates a workspace of the specified type and size on the specified segment and database.
Syntax	<code>sp_dbcc_createws <i>dbname</i>, <i>segname</i>, [<i>wsname</i>], <i>wstype</i>, "wssize[K M]"</code>
Parameters	<p><i>dbname</i></p> <ul style="list-style-type: none">– is the name of the database in which the workspace is to be created. Values are <code>dbccdb</code> and <code>dbccalt</code>. <p><i>segname</i></p> <ul style="list-style-type: none">– is the name of the segment for the workspace. <p><i>wsname</i></p> <ul style="list-style-type: none">– is the name of the workspace. If the value is null, <code>sp_dbcc_createws</code> generates the name <code>scan_ws_nnnnnn</code> for the scan workspace and <code>text_ws_nnnnnn</code> for the text workspace, where <code>nnnnnn</code> is a unique 6-digit number. <p><i>wstype</i></p> <ul style="list-style-type: none">– specifies the type of workspace to be create. Values are <code>scan</code> and <code>text</code>.

wssize

– is the workspace size, specified with K (kilobytes) or M (megabytes).
If you do not specify K or M, *wssize* specifies the number of pages. The minimum size for a workspace is 24 pages.

Examples

Example 1

```
sp_dbcc_createws dbccdb, scanseg, scan_ws_pubs2,  
scan, "10M"
```

Creates a 10MB scan workspace named scan_ws_pubs2 on the scanseg segment in dbccdb.

Example 2

```
sp_dbcc_createws dbccdb, textseg, text, "14M"
```

Creates a 14MB scan workspace named text_ws_000001 on the textseg segment in dbccdb.

Usage

- sp_dbcc_createws creates a workspace with the specified name and size and initializes it.
- Before you create a workspace, create the segment with sp_addsegment.
- Before you create a workspace, make sure you have configured a buffer pool of at least 16K, to achieve maximum performance.
- Use sp_plan_dbccdb to determine size estimates.
- After creating a workspace, run sp_dbcc_updateconfig to record the new configuration information in dbcc_config.
- Each workspace must have a unique name.
- To delete a workspace, in dbccdb issue:

```
drop table workspace_name
```

- For more information on the scan and text workspaces, see the System Administration Guide.
- For information on the dbccalt database, see the System Administration Guide.

Permissions

Only a System Administrator or the Database Owner can run sp_dbcc_createws.

See also *Commands* – `dbcc`
dbcc stored procedures – `sp_dbcc_alterws`, `sp_dbcc_evaluatedb`
System procedures – `sp_addsegment`, `sp_plan_dbccdb`, `sp_helpsegment`

sp_dbcc_deletedb

Description Deletes from `dbccdb` all the information related to the specified target database.

Syntax `sp_dbcc_deletedb [dbname]`

Parameters *dbname*
 – specifies the name of the target database for which you want the configuration information deleted. If you do not specify a value for *dbname*, Adaptive Server deletes data from all databases in `dbccdb..dbcc_config`. If the target database is `dbccdb`, and `dbccalt` exists, Adaptive Server deletes the data from `dbccalt`.

Examples

```
sp_dbcc_deletedb "engdb"
```

 All information for database `engdb` has been deleted from `dbccdb`.
 Deletes all information for the database named `engdb` from `dbccdb`.

Usage

- `sp_dbcc_deletedb` deletes from `dbccdb` all the information related to the specified target database, including configuration information and the results of previous `dbcc` checkstorage operations.
- If the deleted database is `dbccdb`, and the `dbccalt` database exists, `sp_dbcc_deletedb` deletes the configuration information and results of `dbccdb` from `dbccalt`.
- To remove the results of `dbcc` checkstorage operations created before a specific date, use `sp_dbcc_deletehistory`.
- For information about the `dbccalt` database, see the System Administration Guide.

Permissions Only a System Administrator or the Database Owner can run `sp_dbcc_deletedb`.

See also *Commands* – `dbcc`
dbcc stored procedures – `sp_dbcc_deletehistory`, `sp_dbcc_evaluatedb`
System procedures – `sp_plan_dbccdb`

sp_dbcc_deletehistory

Description	Deletes the results of dbcc checkstorage operations performed on the target database before the specified date and time.
Syntax	sp_dbcc_deletehistory [<i>cutoffdate</i> [, <i>dbname</i>]]
Parameters	<p><i>cutoffdate</i></p> <p>– deletes all entries made on or before this date. This parameter is of type datetime. If a date is not specified, only the results of the last operation are retained. For more information, see “Specifying the Date” on page 458.</p> <p><i>dbname</i></p> <p>– specifies the name of the database for which the data must be deleted. If not specified, sp_dbcc_deletehistory deletes the history information for all databases in dbccdb..dbcc_config.</p>
Examples	<pre>sp_dbcc_deletehistory "03/04/1997", "pubs2"</pre> <p>Deletes results of all operations performed on the database pubs2 on or before March 4, 1997.</p>
Usage	<ul style="list-style-type: none">• sp_dbcc_deletehistory deletes the results of dbcc checkstorage operations performed on the target database before the specified date and time.• If the target database is dbccdb, and the dbccalt database exists, sp_dbcc_deletehistory deletes historical data for dbccdb from dbccalt.• The value specified for <i>cutoffdate</i> is compared to the finish time of each dbcc operation.• To see the dates when dbcc checkstorage was run so that you can choose the value for <i>cutoffdate</i>, run sp_dbcc_summaryreport.• For information on the dbccalt database, see the System Administration Guide.
Permissions	Only a System Administrator or the Database Owner can run sp_dbcc_deletehistory on a specific database. Only a System Administrator can run sp_dbcc_deletehistory without specifying a database name.
See also	<p><i>Commands</i> – dbcc</p> <p><i>dbcc stored procedures</i> – sp_dbcc_deletedb, sp_dbcc_evaluatedb</p> <p><i>System procedures</i> – sp_plan_dbccdb</p>

sp_dbcc_differentialreport

Description	Generates a report that highlights the changes in I/O statistics and faults that took place between two dbcc operations.
Syntax	<code>sp_dbcc_differentialreport [dbname [, objectname]], [db_op] [, "date1" [, "date2"]]</code>
Parameters	<p><i>dbname</i></p> <ul style="list-style-type: none"> – specifies the name of the database. If you do not specify a <i>dbname</i>, the report contains information on all databases in <code>dbccdb..dbcc_operation_log</code>. <p><i>objectname</i></p> <ul style="list-style-type: none"> – specifies the name of the table or index for which you want the report generated. If <i>object_name</i> is not specified, statistics on all objects in the target database are reported. <p><i>db_op</i></p> <ul style="list-style-type: none"> – specifies the source of the data to be used for the report. The only value is <code>checkstorage</code>. The report is generated on the data specified by <i>db_op</i> on <i>date1</i> and <i>date2</i> for the specified object in the target database. If dates are not specified, the last two operations of the type <i>db_op</i> are compared. <p><i>date1</i></p> <ul style="list-style-type: none"> – specifies the first date of a dbcc checkstorage operation to be compared. <p><i>date2</i></p> <ul style="list-style-type: none"> – specifies the last date of a dbcc checkstorage operation to be compared.
Examples	<pre>sp_dbcc_differentialreport master, sysprocedures, checkstorage, "05/01/97", "05/04/97"</pre> <p>Generates a report that shows the changes in I/O statistics and faults that occurred in the <code>sysprocedures</code> table between May 1, 1997 and May 4, 1997</p>
Usage	<ul style="list-style-type: none"> • <code>sp_dbcc_differentialreport</code> generates a report that highlights the changes in I/O statistics and faults that occurred between two dbcc operations. It compares counter values reported from two instances of dbcc checkstorage. Only the values that have been changed are reported.

- If only one date is specified, the results of the dbcc checkstorage operation selected by the specified date are compared to the results of the dbcc checkstorage operation immediately preceding the selected operation.
- If no dates are specified, the results of last two dbcc checkstorage operations are compared.
- If sp_dbcc_differentialreport returns a number for *object_name*, it means the object was dropped after the dbcc checkstorage operation completed.
- If no changes occurred between the specified operations, sp_dbcc_differentialreport does not generate a report.

Permissions Any user can run sp_dbcc_differentialreport.

See also *Commands* – dbcc

dbcc stored procedures – sp_dbcc_fullreport, sp_dbcc_statisticsreport, sp_dbcc_summaryreport, sp_dbcc_updateconfig

sp_dbcc_evaluatedb

Description Recomputes configuration information for the target database and compares it to the current configuration information.

Syntax sp_dbcc_evaluatedb [*dbname*]

Parameters *dbname*
– specifies the name of the target database. If *dbname* is not specified, sp_dbcc_evaluatedb compares all databases listed in the dbcc_config table.

Examples sp_dbcc_evaluatedb

Recommended values for workspace size, cache size and worker process count are:

```
Database name : sybsemprocs
current scan workspace size : 400K      suggested scan workspace size : 272K
current text workspace size : 208K     suggested text workspace size : 208K
current cache size : 1024K            suggested cache size : 640K
current process count : 1              suggested process count : 1
```

	Recomputes configuration information for the current database, sybssystemprocs, and suggests new values for some parameters.
Usage	<ul style="list-style-type: none">• <code>sp_dbcc_evaluatedb</code> recomputes configuration information for the target database and compares the data to the current configuration information. It uses counter values recorded for the target database in the <code>dbcc_counters</code> table.• The cache size is the size of the 16K buffer pool in the cache. For a 2K buffer pool, the minimum size of this cache must be the recommended value, plus 512.• When the size and data distribution pattern of the target database changes, run <code>sp_dbcc_evaluatedb</code> to optimize the configuration information.• To gather configuration information for the target database the first time, use <code>sp_plan_dbccdb</code>.• To make sure you are evaluating the most current configuration parameters, run <code>sp_dbcc_updateconfig</code> before running <code>sp_dbcc_evaluatedb</code>.
Permissions	Only System Administrator or the Database Owner can run <code>sp_dbcc_evaluatedb</code> . Only a System Administrator can run <code>sp_dbcc_evaluatedb</code> without specifying a database name.
See also	<i>Commands</i> – <code>dbcc</code> <i>dbcc stored procedures</i> – <code>sp_dbcc_updateconfig</code> <i>System procedures</i> – <code>sp_plan_dbccdb</code>

sp_dbcc_faultreport

Description	Generates a report covering fault statistics for the <code>dbcc checkstorage</code> operations performed for the specified object in the target database on the specified date.
Syntax	<code>sp_dbcc_faultreport [report_type [, dbname [, objectname [, date]]]]</code>
Parameters	<i>report_type</i> – specifies the type of fault report. Valid values are short and long. The default is short.

dbname

– specifies the name of the target database; for example, master..sysdatabases. If *dbname* is not specified, the report contains information on all databases in dbccdb..dbcc_operation_log.

object_name

– specifies the name of the table or index for which you want the report generated. If *object_name* is not specified, statistics on all objects in the target database are reported.

date

– specifies exact date and time that the dbcc checkstorage operation finished. You can find this value in dbcc_operation_log.finish. You can create the value by combining the date from start time and the hours and minutes from end time in the sp_dbcc_summaryreport output. If you do not specify *date*, Adaptive Server uses the date of the most recent operation.

Examples

Example 1

```
sp_dbcc_faultreport "short"
```

Database Name : sybssystemprocs

Table Name	Index	Type Code	Description	Page Number
sysprocedures	0	100031	page not allocated	5702
sysprocedures	1	100031	page not allocated	14151
syslogs	0	100022	chain start error	24315
syslogs	0	100031	page not allocated	24315

Generates a short report of the faults found in tables in the sybssystemprocs database. The report includes the table name, the index number in which the fault occurred, the type code of the fault, a brief description of the fault, and the page number on which the fault occurred.

```
sp_dbcc_faultreport "long"
```

Generating 'Fault Report' for object sysprocedures in database sybssystemprocs.

```
Type Code: 100031; Soft fault, possibly spurious
Page reached by the chain is not allocated.
page id: 14151
page header:
0x0000374700003788000037460000005000648B803EF0001000103FE0080000F
Header for 14151, next 14216, previous 14150, id = 5:1
time stamp = 0x0001000648B8, next row = 1007, level = 0
```

```
free offset = 1022, minlen = 15, status = 128(0x0080)
```

```
.  
.
.
```

Generates a long report of the faults found in tables in the `sysystemprocs` database. This example shows the first part of the output of a long report. The complete report repeats the information for each object in the **target database** in which `dbcc checkstorage` found a fault. The data following the long string of numbers shown under the "page header" field ("Header for 14151, next 14216, previous 14150 ...") describes the components of the "page header" string.

Usage

- `sp_dbcc_faultreport` generates a report that shows all faults for the specified object in the target database.
- If `sp_dbcc_faultreport` returns a number for *object_name*, it means the object was dropped after the `dbcc checkstorage` operation completed.
- For information on the fault ID, see the `type_code` column described in the System Administration Guide.
- For information on the fault status, see the System Administration Guide.

Permissions

Any user can run `sp_dbcc_faultreport`.

See also

Commands – dbcc

dbcc stored procedures – `sp_dbcc_fullreport`, `sp_dbcc_statisticsreport`, `sp_dbcc_summaryreport`, `sp_dbcc_updateconfig`

sp_dbcc_fullreport

Description

Runs `sp_dbcc_summaryreport`, `sp_dbcc_configreport`, `sp_dbcc_statisticsreport`, and `sp_dbcc_faultreport` short for *database..object_name* on or before the specified *date*.

Syntax

```
sp_dbcc_fullreport [dbname [, objectname [, date]]]
```

Parameters

dbname

- specifies the name of the database. If you do not specify *dbname*, the report contains information on all databases in `dbccdb..dbcc_operation_log`.

object_name

– specifies the name of the table or index for which you want the report generated. If you do not specify *object_name*, statistics on all objects in the target database are reported.

date

– specifies the date on which the dbcc checkstorage operation was performed. If you do not specify a *date*, the date of the last operation is used.

Examples

```
sp_dbcc_fullreport master, sysprocedures
```

Runs `sp_dbcc_summaryreport`, `sp_dbcc_configreport`, `sp_dbcc_statisticsreport`, and `sp_dbcc_faultreport` short for the most recent dbcc checkstorage operation run on the `sysprocedures` table in the master database.

Usage

- `sp_dbcc_fullreport` runs `sp_dbcc_summaryreport`, `sp_dbcc_configreport`, `sp_dbcc_statisticsreport`, and `sp_dbcc_faultreport` short for the specified database object on or before the specified date.

Permissions

Any user can run `sp_dbcc_fullreport`.

See also

Commands – dbcc

dbcc stored procedures – `sp_dbcc_statisticsreport`, `sp_dbcc_summaryreport`, `sp_dbcc_updateconfig`

sp_dbcc_runcheck

Description

Runs dbcc checkstorage on the specified database, then runs `sp_dbcc_summaryreport` or a report you specify.

Syntax

```
sp_dbcc_runcheck dbname [, user_proc]
```

Parameters

dbname

– specifies the name of the database on which the check is to be performed.

user_proc

– specifies the name of the dbcc stored procedure or a user-created stored procedure that is to be run instead of `sp_dbcc_summaryreport`.

Examples

Example 1

```
sp_dbcc_runcheck "engdb"
```

Checks the database engdb and generates a summary report on the information found.

Example 2

```
sp_dbcc_runcheck "pubs2", sp_dbcc_fullreport
```

Checks the database pubs2 and generates a full report.

Usage

- `sp_dbcc_runcheck` runs `dbcc checkstorage` on the specified database.
- After the `dbcc checkstorage` operation is complete, `sp_dbcc_runcheck` runs `sp_dbcc_summaryreport` to generate a summary report. If you specify one of the other report-generating `dbcc` stored procedures for *dbcc_report*, `sp_dbcc_runcheck` runs that procedure instead of `sp_dbcc_summaryreport`. For a brief description and examples of all the report-generating stored procedures provided with `dbccdb`, see the System Administration Guide.
- You can write your own report-generating stored procedure and specify its name for `user_proc`. The stored procedure must be self-contained. `sp_dbcc_runcheck` cannot pass any parameters to Adaptive Server.

Permissions

Only a System Administrator or the Database Owner can run `sp_dbcc_runcheck`.

See also

Commands – `dbcc`

dbcc stored procedures – `sp_dbcc_summaryreport`

sp_dbcc_statisticsreport

Description

Generates an allocation statistics report on the specified object in the target database.

Syntax

```
sp_dbcc_statisticsreport [dbname [, objectname  
[, date]]]
```

Parameters

dbname

- specifies the **target database**. If *dbname* is not specified, the report contains information on all databases in `dbccdb..dbcc_operation_log`.

objectname

– specifies the name of the table or index for which you want the report generated. If you do not specify *objectname*, Adaptive Server reports statistics on all objects in the target database.

date

– specifies the date on which the dbcc checkstorage operation was performed. If you do not specify *date*, Adaptive Server uses the date of the most recent operation.

Examples

```
sp_dbcc_statisticsreport 'sybsystemprocs',
'sysobjects'
```

Statistics Report on object sysobjects in database sybsystemprocs

Parameter Name	Index Id	Value
count	0	241.0
max size	0	99.0
max count	0	22.0
bytes data	0	19180.0
bytes used	0	22113.0
count	1	14.0
max size	1	9.0
max level	1	0.0
max count	1	14.0
bytes data	1	56.0
bytes used	1	158.0
count	2	245.0
max level	2	1.0
max size	2	39.0
max count	2	71.0
bytes data	2	4377.0
bytes used	2	6995.0

Parameter Name	Index Id	Partition	Value	Dev_name
page gaps	0	1	13.0	master
pages used	0	1	15.0	master
extents used	0	1	3.0	master
overflow pages	0	1	0.0	master
pages overhead	0	1	1.0	master
pages reserved	0	1	7.0	master
page extent gaps	0	1	11.0	master
ws buffer crosses	0	1	2.0	master
page extent crosses	0	1	11.0	master
pages used	1	1	2.0	master

extents used	1	1	1.0	master
overflow pages	1	1	0.0	master
pages overhead	1	1	1.0	master
pages reserved	1	1	6.0	master
page extent gaps	1	1	0.0	master
ws buffer crosses	1	1	0.0	master
page extent crosses	1	1	0.0	master
page gaps	2	1	4.0	master
pages used	2	1	6.0	master
extents used	2	1	1.0	master
overflow pages	2	1	0.0	master
pages overhead	2	1	1.0	master
pages reserved	2	1	2.0	master
page extent gaps	2	1	0.0	master
ws buffer crosses	2	1	0.0	master
page extent crosses	2	1	0.0	master

Generates a statistics report on the sysobjects table in the master database.

Usage

- `sp_dbcc_statisticsreport` generates an allocation statistics report on the specified object in the **target database**. It uses data from the `dbcc_counters` table, which stores information about page utilization and error statistics for every object in the target database.
- If `sp_dbcc_statisticsreport` returns a number for *object_name*, it means the object was dropped after the dbcc checkstorage operation completed.
- `sp_dbcc_statisticsreport` reports values recorded in the `dbcc_counters` table for the datatypes 5000–5019. See the System Administration Guide.

For bytes data, bytes used, and overflow pages, `sp_dbcc_statisticsreport` reports the sum of the values reported for all partitions and devices.

For count, max count, max size and max level, `sp_dbcc_statisticsreport` reports the largest of the values reported for all partitions and devices.

`sp_dbcc_statisticsreport` reports information for each device and partition used by objects in the **target database** for the following rows:

- extents used
- io errors
- page gaps

- page extent crosses
- page extent gaps
- page format errors
- pages reserved
- pages overhead
- pages misallocated
- pages not allocated
- pages not referenced
- pages used

The page gaps, page extent crosses, and page extent gaps indicate how the data pages for the objects are distributed on the database devices. Large values indicate less effectiveness in using larger buffer sizes and in data prefetch.

- If multiple dbcc checkstorage operations were run on a target database on the same day, sp_dbcc_statisticsreport generates a report based on the results of the last dbcc checkstorage operation that finished before the specified time.

Permissions

Any user can run sp_dbcc_statisticsreport.

See also

Commands – dbcc

dbcc stored procedures – sp_dbcc_fullreport, sp_dbcc_summaryreport, sp_dbcc_updateconfig

sp_dbcc_summaryreport

Description

Generates a summary report on the specified database.

Syntax

sp_dbcc_summaryreport [*dbname* [, *date*] [, *opname*]]

Parameters

dbname

– specifies the name of the database for which you want the report generated. If you do not specify *dbname*, sp_dbcc_summaryreport generates reports on all databases in dbccdb..dbcc_operation_log for which the date is on or before the date and time specified by the *date* option.

date

– specifies the date on which dbcc checkstorage was performed. If you do not specify a date, sp_dbcc_summaryreport uses the date of last dbcc checkstorage operation performed on the **target database**. This parameter is of the datatype datetime. If both the date and the time are specified for *date*, summary results of all the operations performed on or before the specified time are reported. If no date is specified, all operations are reported.

opname

– specifies the operation. *opname* may be either checkstorage, which is the default, or checkverify, or both. If *opname* is not specified, reports are generated for all operations.

Examples

Example 1

```
sp_dbcc_summaryreport
```

```
DBCC Operation : checkstorage
```

Database Name	Start time	End Time	Operation ID
Hard Faults	Soft Faults	Text Columns	Abort Count
User Name			
sybssystemprocs	05/11/1999 14:53:11	14:53:32:163	1
0	0	0	0
sa			
sybssystemprocs	05/11/1999 14:55:06	14:55:29:200	2
0	0	0	0
sa			
sybssystemprocs	05/11/1999 14:56:10	14:56:27:750	3
0	0	0	0
sa			

```
DBCC Operation : checkverify
```

Database Name	Start time	End Time	Operation ID
Hard Faults	Soft Faults	User Name	
sybssystemprocs	05/11/1999 14:55:29	14:55:29:310	2
0	0	sa	

Generates a summary report on the master database, providing information on all dbcc checkstorage and dbcc checkverify operations performed.

Example 2

```
sp_dbcc_summaryreport "testdb"
```

DBCC Operation : checkstorage

Database Name	Start time	End Time	Operation ID
Hard Faults	Soft Faults	Text Columns	Abort Count
User Name			
testdb	05/11/1999 14:55:29	14:55:49:903	1
0	0	0	0
sa			
testdb	05/11/1999 14:55:50	14:56:9:546	2
0	0	0	0
sa			
testdb	05/11/1999 14:56:28	14:56:40:666	3
0	0	0	0
sa			

Generates a summary report on the user database testdb, providing information on all dbcc checkstorage operations performed. dbcc checkstorage was the only operation run on this database, so no dbcc checkverify information appears on the report.

Example 3

```
sp_dbcc_summaryreport null, "checkverify"
```

DBCC Operation : checkverify

Database Name	Start time	End Time	Operation ID
Hard Faults	Soft Faults	User Name	
sybsystemprocs	05/11/1999 14:55:29	14:55:29:310	2
0	0	sa	

Generates a summary report on the master database, providing information on all dbcc checkverify operations performed. Because dbcc checkverify was the specified operation, no dbcc checkstorage information appears on the report.

Example 4

```
sp_dbcc_summaryreport sybssystemprocs,
"checkstorage"
```

DBCC Operation : checkstorage

Database Name	Start time	End Time	Operation ID
Hard Faults	Soft Faults	Text Columns	Abort Count
User Name			
-----	-----	-----	-----
-----	-----	-----	-----
sybssystemprocs	05/11/1999 14:53:11	14:53:32:163	1
0	0	0	0
sa			
sybssystemprocs	05/11/1999 14:55:06	14:55:29:200	2
0	0	0	0
sa			
sybssystemprocs	05/11/1999 14:56:10	14:56:27:750	3
0	0	0	0
sa			

Generates a summary report on the sybssystemprocs database, providing information on all dbcc checkstorage operations performed. Because dbcc checkstorage was the specified operation, no dbcc checkverify information appears on the report.

Usage

- sp_dbcc_summaryreport generates a summary report of checkstorage or checkverify operations, or both, on the specified database.
- The report indicates the name of the database that was checked, the start and end time of the dbcc checkstorage run and the number of soft and hard faults found.
- The “Operation ID” column contains a number that identifies the results of each dbcc checkstorage operation on a given database at a specific time. The number provided in the report comes from the opid column of the dbcc_operation_log table. For more information, see the System Administration Guide.
- The “Text Columns” column shows the number of non-null text columns found by dbcc checkstorage during the run.
- The “Abort Count” column shows the number of tables that contained errors, which caused dbcc checkstorage to abort the check on the table. For details on the errors, run sp_dbcc_faultreport.

Permissions

Any user can run sp_dbcc_summaryreport.

See also

Commands – dbcc

dbcc stored procedures – sp_dbcc_fullreport, sp_dbcc_statisticsreport, sp_dbcc_updateconfig

sp_dbcc_updateconfig

Description

Updates the dbcc_config table in dbccdb with the configuration information of the target database.

Syntax

```
sp_dbcc_updateconfig dbname, type, "str1" [, "str2"]
```

Parameters

dbname

– is the name of the target database for which configuration information is being updated.

type

– specifies the type name from the dbcc_types table. Table 27-2 on page 480 shows the valid values for *type*.

str1

– specifies the first configuration value for the specified *type* to be updated in the dbcc_config table. Table 27-2 on page 480 describes the expected value of *str1* for the specified *type*.

str2

– specifies the second configuration value for the specified *type* that you want to update in the dbcc_config table. Table 27-2 on page 480 describes the expected value of *str2* for the specified *type*.

Examples

Example 1

```
sp_dbcc_updateconfig pubs2, "max worker processes",  
"4"
```

Updates dbcc_config with the maximum number of worker processes for dbcc checkstorage to use when checking the pubs2 database. The new maximum number of worker processes is 4.

Example 2

```
sp_dbcc_updateconfig pubs2, "dbcc named cache",  
pubs2_cache, "10K"
```

Updates dbcc_config with the size of the dbcc named cache “pubs2_cache”. The new size is 10K.

Example 3

```
sp_dbcc_updateconfig pubs2, "scan workspace",
scan_pubs2
```

Updates dbcc_config with the new name of the scan workspace for the pubs2 database. The new name is scan_pubs2. This update is made after using sp_dbcc_alterws to change the name of the scan workspace.

Example 4

```
sp_dbcc_updateconfig pubs2, "text workspace",
text_pubs2
```

Updates dbcc_config with the new name of the text workspace for the pubs2 database. The new name is text_pubs2. This update is made after using sp_dbcc_alterws to change the name of the text workspace.

Example 5

```
sp_dbcc_updateconfig pubs2, "OAM count threshold", 5
```

Updates dbcc_config with the OAM count threshold value for the pubs2 database. The new value is 5.

Example 6

```
sp_dbcc_updateconfig pubs2, "IO error abort", 3
```

Updates dbcc_config with the I/O error abort value for the pubs2 database. The new value is 3.

Example 7

```
sp_dbcc_updateconfig pubs2, "linkage error abort", 8
```

Updates dbcc_config with the linkage error abort value for the pubs2 database. The new value is 8.

Usage

- sp_dbcc_updateconfig updates the dbcc_config table for the **target database**.
- If the name of the target database is dbccdb, and the database dbccalt exists, sp_dbcc_updateconfig updates the dbcc_config table in dbccalt.
- If the target database name is not found in dbcc_config, sp_dbcc_updateconfig adds it and sets the operation sequence number to 0 before updating other configuration information.
- If the expected value for the specified *type* is a number, sp_dbcc_updateconfig converts the values you provide for *str1* and *str2* to numbers.

- Table 27-2 shows the valid type names to use for *type* and the expected value for *str1* or *str2*.

Table 27-2: Type names and expected values

type Name	Value expected for <i>str1</i> or <i>str2</i>
dbcc named cache	The name of the cache, specified by <i>str1</i> , and the new size (in kilobytes or megabytes) or the number of 2K pages, specified by <i>str2</i> .
IO error abort	The new error count, specified by <i>str1</i> . The value must be a number greater than 0. <i>str2</i> is not used with this type.
linkage error abort	The new linkage error count value specified in <i>str1</i> . The value must be a number greater than 0. <i>str2</i> is not used with this type.
max worker processes	The new number of worker processes, specified by <i>str1</i> . The value must be a number greater than 0. <i>str2</i> is not used with this type.
OAM count threshold	The new threshold count, specified by <i>str1</i> . The value must be a number greater than 0. <i>str2</i> is not used with this type.
scan workspace	The new name for the scan workspace, specified by <i>str1</i> . <i>str2</i> is not used with this type.
text workspace	The new name of the text workspace, specified by <i>str1</i> . <i>str2</i> is not used with this type.

- For more information on the *type* names and values, see the System Administration Guide.

Permissions

Only a System Administrator or the Database Owner can run `sp_dbcc_updateconfig`.

See also

Commands – dbcc
dbcc stored procedures – sp_dbcc_evaluatedb
System procedures – sp_plan_dbccdb

Index

Symbols

- @ (at sign)
 - procedure parameters and 3
- ::= (BNF notation)
 - in SQL statements xv
- , (comma)
 - in SQL statements xv
 - in user-defined datatypes 59
- { } (curly braces)
 - in SQL statements xv
- .. (dots) in database object names 34
- () (parentheses)
 - in SQL statements xv
 - in user-defined datatypes 59
- “ ” (quotation marks)
 - enclosing parameter values 3, 418
 - enclosing reserved words 107
 - single, and **quoted_identifier** 113
- [] (square brackets)
 - in SQL statements xv

Numerics

- 0 return status 1, 418
- 7-bit ASCII characters, checking with **sp_checknames** 100
- 7-bit terminal, **sp_helpsort** output 282
- 8-bit terminal, **sp_helpsort** output 283

A

- abort tran on log full** database option 146
- abstract plan groups
 - adding 37
 - dropping 188
 - exporting 212
 - importing 289

- renaming 359
- abstract plans
 - information about 268
 - viewing with **sp_help_qplan** 268
- accounting, chargeback
 - sp_clearstats** 121
 - sp_reportstats** 359–360
- accounts. *See* logins
- actions
 - modifying for resource limits 311
 - resource limit information on 272
 - specifying for resource limits 42
- adding
 - abstract plan groups 37
 - aliases 17–37
 - date strings 27–30
 - dump devices 62–64
 - engine groups 22
 - engines to a group 22
 - execution classes 22
 - foreign keys 224–225
 - group to a database 26–27
 - limits 40
 - logins to Server 30–32
 - messages to *sysusermessages* 32–34
 - named time ranges 55
 - remote logins 38–40
 - resource limits 40
 - segments 47–48
 - servers 48–51
 - thresholds 51–55
 - time ranges 55
 - user-defined datatypes 58–62
 - users to a database 64–65
 - users to a group 64–65, 99–100
- ad hoc** auditing option 68
- aliases, language
 - assigning 368–369
 - defining 27–30
- aliases, server 49

Index

- aliases, user
 - See also* logins; users
 - assigning 17–18
 - assigning different names compared to 64
 - database ownership transfer and 98
 - dropping 175, 201
 - help on 286
 - sysalternates* table 17, 175
- all** auditing option 68
- allow nulls by default** database option 146
- alter** auditing option 68
- alter database** command
 - sp_dbremap** and 153
- alternate identity. *See* aliases, user
- alternate languages. *See* languages, alternate
- ANYENGINE* engine group 22
- applications
 - applying resource limits to 41
 - dropping resource limits from 192
 - modifying resource limits for 310
 - resource limit information on 271
- ASCII characters
 - checking for with **sp_checknames** 100
- asynchronous prefetch
 - configuring limits 340
- at sign (@)
 - procedure parameters and 3
- @@connections** global variable
 - sp_monitor** and 320
- @@cpu_busy** global variable
 - sp_monitor** and 319
- @@idle** global variable
 - sp_monitor** and 320
- @@io_busy** global variable
 - sp_monitor** and 319
- @@ncharsize** global variable
 - sp_addtype** and 61
- @@pack_received** global variable
 - sp_monitor** and 320
- @@pack_sent** global variable
 - sp_monitor** and 320
- @@packet_errors** global variable
 - sp_monitor** and 320
- @@thresh_hysteresis** global variable
 - threshold placement and 52
- @@total_errors** global variable
 - sp_monitor** and 320
- @@total_read** global variable
 - sp_monitor** and 320
- @@total_write** global variable
 - sp_monitor** and 320
- attributes
 - execution classes 23
 - server (**sp_server_info**) 429
 - sp_addobjectdef** and 37
- audit trail
 - adding comments 18
- auditing
 - adding an audit table 20
 - options, displaying 164
- auditing options
 - adhoc** 68
 - all** 68
 - alter** 68
 - bcp** 68
 - bind** 68
 - cmdtext** 68
 - create** 68
 - dbaccess** 68
 - dbcc** 68
 - delete** 68
 - disk** 68
 - drop** 68
 - dump** 68
 - errors** 68
 - exec_procedure** 68
 - exec_trigger** 68
 - func_dbaccess** 68
 - func_obj_access** 68
 - grant** 68
 - insert** 68
 - load** 68
 - login** 68
 - logout** 68
 - reference** 68
 - revoke** 68
 - rpc** 69
 - security** 69
 - select** 69
 - setting 68
 - setuser** 69
 - table_access** 69

- truncate** 69
 - unbind** 69
 - update** 69
 - view_access** 69
 - authority. *See* permissions
 - authorizations. *See* permissions
 - auto identity** database option 146
- ## B
- Backup Server
 - See also* *Utility Guide*
 - amount dumped, specifying 203
 - information about 281
 - multiple 50
 - volume handling messages 410–413
 - Backus Naur Form (BNF) notation xiv, xv
 - basic display level for configuration parameters 168
 - bcp** (bulk copy utility)
 - select into/bulkcopy/pllsort** and 149
 - bcp** auditing option 68
 - binary sort order of character sets 283
 - bind** auditing option 68
 - binding
 - data caches 75–78
 - defaults 78–80
 - objects to data caches 75–78
 - rules 84–85
 - unbinding and 403–405, 406
 - user messages to constraints 83
 - blanks
 - catalog stored procedure parameter values 418
 - in system procedure parameter values 3
 - blocking process
 - sp_lock** report on 216, 295
 - sp_who** report on 415
 - BNF notation in SQL statements xiv, xv
 - brackets. *See* square brackets []
 - bulk copying. *See* **bcp** (bulk copy utility)
- ## C
- caches, data
 - binding objects to 75
 - configuring 87–95
 - dropping 94
 - information about 90, 239
 - logonly** type 94
 - memory pools 336–341
 - overhead 94, 239
 - recovery and 90
 - status 92
 - unbinding all objects from 405
 - unbinding objects from 403
 - case sensitivity
 - in SQL xvi
 - catalog stored procedures 417–439
 - list of 417
 - return status 418
 - syntax 418–419
 - chained transaction mode
 - sp_procxmode** and 347
 - changing
 - database options 143–150
 - database owners 98–99
 - dbccdb* workspace size 459
 - language alias 368
 - memory pools within data caches 336
 - names of abstract plan groups 359
 - object names 355–356
 - passwords for login accounts 331–332
 - resource limits 310
 - thresholds 314–318
 - time ranges 312
 - user's group 99–100
 - character sets
 - changing names of 110, 112
 - checking with **sp_checknames** 100
 - checking with **sp_checkreswords** 106
 - multibyte 283
 - sp_helpsort** display of 282
 - chargeback accounting
 - sp_clearstats** procedure 121–122
 - sp_reportstats** procedure 359–360
 - check constraints
 - binding user messages to 83
 - displaying source text of 283
 - renaming 355–356
 - checking passwords. *See* passwords; **sp_remotoption** system procedure

Index

- checkpoint** command
 - setting database options and 146
- clearing accounting statistics 121–122
- clustered indexes
 - indid* not equal to one 256
- cmd* returned by **sp_who** 416
- cmdtext** auditing option 68
- codes
 - datatype 425
 - ODBC datatype 420
- collating sequence. *See* sort order
- collisions
 - hash key 268
- column data. *See* datatypes
- column name
 - changing 109, 355–356
 - checking with **sp_checknames** 101
- column pairs. *See* joins; keys
- columns
 - common key 125–126
 - datatypes 422
 - defaults for 78–80
 - dependencies, finding 109
 - foreign keys 224–225, 426
 - joins and 262
 - permissions on 420
 - primary key 341
 - returned by **sp_who** 415
 - rules 84–85
 - unbinding defaults from 405–406
 - unbinding rules with **sp_unbindrule** 409–410
- comma (,)
 - in SQL statements xv
 - in user-defined datatypes 59
- commands
 - display syntax of 389–390
- comments
 - adding to audit trail 18
- common keys
 - See also* foreign keys; joins; primary keys
 - defining 125–126
 - dropping 181
 - join candidates and 262
 - reporting 262–264
- companion servers
 - configuring 127–129
- comparing plan groups 122
- comparing plans 122, 124
- compiled objects
 - checking for source text of 114
 - displaying source text of 283
 - hiding source text of 287
- compiling
 - sp_recompile** and 349
- comprehensive display level for configuration
 - parameters 168
- concurrency optimization 117
- concurrency_opt_threshold** option, **sp_chgattribute** 117
- configuration parameters
 - changing 130–134
 - display levels 168
 - help information on 240
- constraints
 - binding user messages to 83
 - displaying source text of 283
 - information about 237, 244
 - renaming 355–356
 - unbinding messages with **sp_unbindmsg** 408
- contention, lock
 - monitoring with **sp_object_stats** 324–327
- controller, device
 - sp_helpdevice** and number 252
- conventions
 - See also* syntax
 - Transact-SQL syntax xiv
 - used in the Reference Manual xiv
- copying
 - plan groups 135
 - plans 135, 136
- corrupt databases
 - listing 291
 - recovery fault isolation mode 376
- corrupt pages
 - bringing online 222–224
 - isolating on recovery 376–378, 378–379
 - listing 292
- CPU usage
 - monitoring 319
- @*cpu_busy* global variable
 - sp_monitor** and 319
- create** auditing option 68

create database command
 log on option compared to **sp_logdevice** 299

create index command
 sp_extendsegment and 214

create table command
 sp_extendsegment and 214

creating
 abstract plan groups 37
 datatypes 58–62
 dbccdb workspaces 461
 execution classes 22
 extended stored procedures 23–24
 limits 40
 named time ranges 55
 resource limits 40
 thresholds 51–55
 time ranges 55
 user aliases 17–18
 user groups 26
 user-defined audit records 68

curly braces ({}) in SQL statements xv

current database
 information from **sp_helpdb** 250
 space used by 385–387

current locks, **sp_lock** system procedure 293

current usage statistics 359–360

cursors
 information about 138

custom audit records 68

custom datatypes. *See* user-defined datatypes

D

data caches
 binding objects to 75
 configuring 87–95
 dropping 94
 information about 90, 239
 logonly type 94
 memory pools 336–341
 overhead 94, 239
 recovery and 90
 status 92
 unbinding all objects from 405
 unbinding objects from 403

data dependency. *See* dependencies, database object

database design
 dropping keys 182
 logical relationships in 126, 224

database devices
 defaulton or **defaultoff** status 163–164
 dropping 176–177
 dropping segments from 196–197
 dsynch setting of 162
 listing of 251
 sp_helpdevice system procedure 251
 status 163

database files. *See* files

database object owners
 sp_depends system procedure and 156

database objects
 binding defaults to 78–80
 binding rules to 84
 binding to caches 75
 dependencies of 156–161
 finding 161, 236
 listings of 231
 permissions on 273
 remapping 350
 renaming 355–356
 sp_tables list of 439
 space used by 385–387

database options 146–150
 See also individual option names
 listing 143–150
 showing settings 145, 248

database owners
 See also database object owners; permissions
 adding users 64
 changing 98
 dbo use only database option 147
 information about 285–286
 transferring ownership 98

database recovery order
 sp_dbrecovery_order system procedure 150–152
 system databases and 151

databases
 See also database objects
 adding groups 26
 adding users 64
 binding to data caches 75, 77

- changing user's default 307
- checking with **sp_checknames** 101
- dropping row lock promotion thresholds for 195
- dropping segments from 196–197
- dropping users from 200
- help on 248
- information on storage space used 250, 385
- listing suspect 291
- listing suspect pages in 292
- listing with **sp_databases** 424
- listing with **sp_helpdb** 248
- lock promotion thresholds for 370
- options 143–150
- ownership 98
- renaming 356–358
- running out of space in 393
- setting row lock promotion thresholds for 374
- storage information 385
- thresholds 393
- unbinding from data caches 403
- datatype precedence. *See* precedence
- datatypes
 - codes 420, 425
 - defaults and 78–80
 - dropping user-defined 200
 - hierarchy 61
 - ODBC 420
 - physical 59
 - sp_datatype_info** information on 425
 - sp_help** information on 231–237
 - unbinding defaults from 405–406
 - unbinding rules with **sp_unbindrule** 409–410
- datatypes, custom. *See* user-defined datatypes
- date parts
 - order of 28
- day-long time ranges 56
- days
 - alternate language 28
 - in time ranges 56
- dbaccess** auditing option 68
- dbcc** (Database Consistency Checker)
 - scripts and **sp_checkreswords** 108
 - space allocation and 333
- dbcc** auditing option 68
- dbccdb* database
 - changing workspace size in 459
 - creating workspaces in 461
 - deleting **dbcc checkstorage** history from 464
 - deleting target database information from 463
 - reporting allocation statistics from 471
 - reporting comprehensive information from 469
 - reporting configuration information from 460, 467, 469
 - reporting fault information from 465, 467
 - reporting full details from 469
 - reporting I/O statistics from 465
 - stored procedures for use with 457
- DB-Library programs
 - changing identifier names and 108
- dbo use only** database option
 - setting with **sp_dboption** 146
- ddl in tran** database option 147
- default database
 - See also sysdevices* table
 - assigning with **sp_addlogin** 30
 - changing user's 307
- default database devices
 - setting status with **sp_diskdefault** 163
 - sp_helpdevice** and 251
- default language id** configuration parameter 30
- default* segment
 - dropping 197
 - mapping 48
- default settings
 - changing login 32, 307
 - configuration parameters 133
 - language 30
- defaulton | defaultoff** option, **sp_diskdefault** 163
- defaults
 - binding 78–80
 - checking name with **sp_checkreswords** 105
 - displaying source text of 283
 - remapping 350
 - renaming 108, 355–356
 - system tables and 79
 - unbinding 405–406
- defncopy** utility command 107
- delete** auditing option 68
- deleting
 - See also* dropping
 - dbcc checkstorage** history from *dbccdb* 464
 - files 176

- plans 175, 189
 - target database information from *dbccdb* 463
 - delimited identifiers
 - testing 107
 - using 106, 113–114
 - denying access to a user 297
 - dependencies, database object
 - changing names of 107
 - recompilation and 356
 - sp_depends** system procedure 156–161
 - detail** option, **sp_helpconstraint** 245
 - device fragments
 - sp_helpdb** report on 248
 - devices
 - See also sysdevices* table
 - changing names of 110, 112
 - dsync setting for 162
 - information on log 265
 - direct updates
 - to system tables 110
 - disk** auditing option 68
 - disk devices
 - adding 62–64
 - disk mirroring
 - sp_who** report on 416
 - disk** option, **sp_addumpdevice** 62
 - display
 - character sets 282
 - database options 143–150
 - source text of compiled objects 283
 - syntax of modules 389
 - distributed Transaction Management (DTM) 395
 - dots (..) for omitted name elements 34
 - drop** auditing option 68
 - drop logins** option, **sp_dropserver** 197
 - dropmessages** option, **sp_droplanguage** 183
 - dropping
 - See also deleting*
 - abstract plan groups 188
 - aliased user 175
 - database devices 176–177
 - groups 180–181
 - lock promotion thresholds 180
 - plans 175, 189
 - procedures 178
 - remote logins 191–192, 198
 - remote servers 197–198
 - resource limits 192
 - row lock promotion thresholds 195
 - segment from a database 196–197
 - time ranges 199
 - user from a database 200–201
 - user from a group 100
 - user-defined datatype 200
 - user-defined messages 184–185
 - workspaces 462
 - dsync setting 162
 - dump** auditing option 68
 - dump devices
 - See also database devices; log device*
 - adding 62–64
 - dropping 176–177
 - listing 251
 - permission and ownership problems 63
 - dump transaction** command
 - sp_logdevice** and 299
 - dump, database
 - interrupted 153
 - dumping databases 203
 - dumpvolume** option
 - dump database** 411
 - Dynamic Link Libraries (DLLs)
 - unloading 225–226
- ## E
- 8-bit terminal, **sp_helpsort** output 283
 - encryption
 - compiled object source text 287
 - reversing 288
 - ending days of named time ranges 55
 - ending times of named time ranges 56
 - enforcing resource limits 42
 - errors
 - number of 320
 - errors** auditing option 68
 - exclusive locks 216, 295
 - exclusive row locks 296
 - exec_procedure** auditing option 68
 - exec_trigger** auditing option 68
 - execution

Index

- operating system commands 442
 - exp_row_size** option
 - sp_chgattribute** 117
 - sp_help** report on 236
 - expand_down** parameter
 - sp_activeroles** 15
 - sp_displayroles** 172
 - sp-displayroles** 172
 - exporting plan groups 212
 - extended stored procedures
 - creating 23–24
 - displaying 252
 - dropping 178
 - extending
 - segments 213
- ## F
- failures, media
 - trunc log on chkpt** database option and 149
 - family of worker processes
 - fid* reported by **sp_lock** 295
 - sp_familylock** report on *fid* 215
 - fault isolation
 - index level 221, 291
 - fid* (family ID) number 215
 - sp_lock** report 295
 - file names
 - configuration file 131
 - DLL 225
 - files
 - See also* tables; transaction log
 - deleting 176
 - inaccessible after **sp_dropdevice** 176
 - interfaces, and server names 49
 - localization 112
 - fillfactor** option
 - sp_chgattribute** 116
 - finding
 - cache bindings 87, 239
 - character sets 282
 - configuration parameters 240
 - constraints 244
 - database objects 236
 - database options 143
 - database settings 248
 - datatypes 231
 - devices 251
 - languages 264
 - object dependencies 156, 161
 - object information 231
 - partition information 237
 - permissions 273
 - reserved words 101
 - resource limits 270
 - segments 279
 - server names 281
 - thresholds 285
 - users in a database 285
 - first page
 - log device 265
 - partition, displaying with **sp_helppartition** 237
 - foreign keys
 - dropping 181
 - inserting 224–225
 - sp_fkeys** information on 426
 - sp_helpkey** and 262
 - format strings
 - in user-defined error messages 34
 - formats
 - times in named time ranges 56
 - formulas
 - max_rows_per_page** of nonclustered indexes 119
 - fragments, device space
 - sp_placeobject** and 333
 - from** keyword
 - sp_tables** list of objects appearing in clause 439
 - full name
 - changing with **sp_modifylogin** 307
 - specifying with **sp_addlogin** 32
 - func_dbaccess** auditing option 68
 - func_obj_access** auditing option 68
 - future space allocation. *See* space allocation;
 - sp_placeobject** system procedure
 - futureonly** option
 - sp_bindefault** 79
 - sp_bindrule** 84, 85
 - sp_unbindefault** 406
 - sp_unbindrule** 409

G

getting messages. *See* **sp_getmessage** system procedure

global variables

See also individual variable names

sp_monitor report on 319

grant auditing option 68

grant option

sp_helprotect 273

sp_role 362

groups

See also “public” group

changing 99–100

dropping 180–181

information about 254

sp_addgroup 26–27

sp_adduser procedure 64

Windows NT domain 444

guest users

sybssystemprocs database 2

H

hash-key collisions 268

help

sp_syntax display 389

sp_sysmon display 390

help reports

See also information (server); system procedures

constraints 244

database devices 251

database object 231

databases 248

datatypes 231

dump devices 251

extended stored procedures 252

groups 254

indexes 255

joins 261

keys 262

language, alternate 264

logins 270

permissions 273

remote servers 281

resource limits 270

segments 279

source text for compiled objects 283

system procedures 231–286

tables 231

thresholds 285

users 285–286

hierarchy

See also precedence

data cache bindings 76

lock promotion thresholds 370, 375

roles, displaying with **sp_activeroles** 15

user-defined datatypes 61

hierarchy of roles. *See* role hierarchies

high availability

configuring Adaptive Server for 127

holdlock keyword

select 294

I

I/O

concurrency_opt_threshold and 117

configuring size 336

limiting 41

log size 302

usage statistics 359

identifiers

delimited 106

quoted 106

renaming 107

reserved words and 101–114

set quoted_identifier on 106, 113–114

sp_checkreswords and 106

identities

alternate 17

IDENTITY columns

automatic 146, 150

database options using 148

nonunique indexes 148

identity in nonunique index database option

setting with **sp_dboption** 148

identity keyword

sp_addtype and 59

identity of user. *See* aliases; logins; users

identity_gap option

Index

- sp_chgattribute** 117
 - @@*idle* global variable
 - sp_monitor** and 320
 - IDs, time range 57
 - IDs, user
 - See also* logins
 - image* datatype
 - size of 386
 - importing abstract plan groups 289
 - index pages
 - locks on 296
 - indexes
 - binding to data caches 75
 - checking name with **sp_checknames** 101
 - checking name with **sp_checkreswords** 105
 - estimating space and time requirements 209
 - IDENTITY columns in nonunique 148
 - information about 255
 - order of, reported by **sp_helpindex** 257
 - renaming 108, 355–356
 - sp_placeobject** space allocation for 332–333
 - sp_statistics** information on 435
 - space used by 386
 - suspect 290–291
 - unbinding from data caches 403
 - information (server)
 - current locks 293
 - indexes 255
 - suspect indexes 290–291
 - information about, reporting
 - cache bindings 77
 - current locks 293
 - data caches 90
 - database devices 251
 - database objects 231
 - database owners 285
 - databases 248
 - datatypes 231
 - dump devices 251
 - extended stored procedures 252
 - first page of log 265
 - groups 254, 285
 - indexes 255
 - join columns 261
 - keys 262
 - languages 264
 - locks 293, 324
 - logins 413
 - performance 390
 - permissions 273
 - remote server logins 270
 - remote servers 281
 - resource limits 270
 - segments 279
 - server processes 413
 - server users 169, 413
 - source text for compiled objects 283
 - space usage 385
 - statistics, monitor 318
 - suspect indexes 290
 - thresholds 285
 - transaction log device 265
 - users, database 285
 - input packets, number of 320
 - insert** auditing option 68
 - intent table locks 216, 295
 - interfaces file
 - changing server names in 112
 - sp_addserver** and 49
 - intermediate display level for configuration parameters 168
 - @@*io_busy* global variable
 - sp_monitor** and 319
 - isolation levels
 - catalog stored procedures 418
 - identity in nonunique index** database option and 148
 - system procedures 1
- ## J
- Java items
 - sp_helpjava** system procedure 259
 - joins
 - information about 261
 - sp_commonkey** 125
- ## K
- keys, table

See also common keys; indexes
 dropping 182
 information about 262
syskeys table 126, 224, 341
 keywords
 as identifiers 101
kill command
 sp_who and 415

L

language defaults 30
 adding 27–30
 changing user's 32
 languages, alternate
 alias for 368
 changing names of 110, 112
 checking with **sp_checkreswords** 106
 date formats in 27
 dropping 182–183
 dropping messages in 184
 information on 264
 installing 27
 official name 368
 syslanguages table 265
 system messages and 226
 user-defined messages 33
 without Language Modules 27
 last-chance thresholds 53, 316
LASTONLINE engine group 22
 limit types 41
 elapsed time 41
 I/O cost 41
 modifying values 311
 number of rows returned 41
 specifying values 41
 limited days
 modifying for time ranges 313
 resource limit information on 271
 specifying for time ranges 55
 limited times
 modifying for time ranges 313
 resource limit information on 271
 specifying for time ranges 56
 linking users. *See* alias, user

listing
 database options 143
 devices 251
 lists
 catalog stored procedures 417
 dbcc stored procedures 457
 system procedures 1–13
load auditing option 68
 local alias, language 368
local option, **sp_addserver** 49
 local servers 49
 See also remote servers; servers
 localization
 changing language names and files 112
lock | unlock option, **sp_locklogin** 297
 lock promotion thresholds 369
 dropping row with **sp_droprowlockpromote** 195
 setting row with **sp_setrowlockpromote** 373
 sp_help report on 237
 sp_helpdb report on database setting 248
 locking
 cache binding and 77
 cache unbinding and 404
 control over 293–296
 logins 297
 monitoring contention 324
 locking scheme
 sp_help report on 236
 locks
 displaying information about 293, 324
 exclusive page 216
 exclusive table 216
 exclusive table and page 295
 “FAM DUR” status 217
 intent table 216, 295
 page 216, 295
 reported by **sp_lock** 293
 row 296
 shared page 216, 295
 shared table 216, 295
 sp_familylock system procedure 215–217
 sp_lock system procedure 293–296
 types of 216, 295
 log device
 information 265
log on option

Index

- create database**, and **sp_logdevice** 299
 - log segment
 - sp_helplog** report on 265
 - sp_helpthreshold** report on 285
 - logging
 - messages 67
 - user-defined events 446
 - user-defined messages 33
 - logical device name 62, 163
 - syslogs* table 298
 - login** auditing option 68
 - logins
 - See also* remote logins; users
 - accounting statistics 121, 360
 - adding to Servers 30–32
 - alias 17, 175
 - applying resource limits to 40
 - changing current database owner 98
 - dropping 183, 197
 - dropping resource limits from 192
 - information about 270
 - information on 169
 - locking 297–298
 - modifying accounts 307–310
 - modifying resource limits for 310
 - number of 320
 - options for remote 351
 - password change 331–332
 - “probe” 360
 - remote 191–192, 197
 - resource limit information on 271
 - sysremotelogins* table 38–40, 191, 198
 - unlocking 297–298
 - logout** auditing option 68
 - logsegment* log storage
 - dropping 197
 - lower and higher roles. *See* role hierarchies
- M**
- machine ticks 319
- mail messages, server
 - deleting 444
 - processing 342–344
 - reading 447
 - sending 450
 - starting session 453
 - stopping session 454
- mapping
 - databases 153
 - remote users 38
- master* database
 - checking with **sp_checkreswords** 105
 - sp_dboption** and 145
 - system procedure tables 4
 - thresholds and 53, 316
- max_rows_per_page** option
 - changing with **sp_relimit** 116
 - sp_chgattribute** 116
- memory
 - freeing from XP Server 225–226
 - mapping 153
 - used by configuration parameters 241
- memory pools
 - configuring 336
 - configuring asynchronous prefetch limits 340
 - configuring wash percentage 340
 - defaults 88
 - minimum size of 339
 - sp_logiosize** and 303
 - transaction logs and 339
- message **output** parameter, **sp_getmessage** 226
- messages
 - adding user-defined 32–34
 - dropping system with **sp_droplanguage** 183
 - dropping user-defined 184–185
 - language setting for 184, 226
 - logging 67
 - number for 32, 67, 184, 226
 - sp_getmessage** procedure 226–227
 - system procedure 4
 - sysusermessages* table 32–34
 - unbinding with **sp_unbindmsg** 408
- mirroring. *See* disk mirroring
- model* database
 - changing database options 145
- modifying
 - configuration parameter display level 168
 - configuration parameters 130
 - login accounts 307
 - named time ranges 312

- resource limits 310
- thresholds 314
- modifying abstract plans 372
- modules, display syntax of 389
- monitoring
 - lock contention 324
 - space remaining 51, 53, 316
 - system activity 318
- month values
 - alternate language 28
- moving
 - indexes 332
 - tables 332
 - transaction logs 298
 - user to new group 99
- MRU replacement strategy
 - disabling 95
- multibyte character sets
 - sort order 283
 - sp_helpsort** output 283
- mutual authentication** option, **sp_serveroption** 365

N

- named time ranges
 - adding 55
 - “at all times” 57, 199
 - changing active time ranges 57
 - creating 55
 - dropping 199
 - entire day 56
 - IDs for 57
 - modifying 312
 - overlapping 57
- names
 - alias 17, 175, 201
 - assigning different, compared to aliases 64
 - changing database object 355–356
 - changing identifier 107
 - checking with **sp_checknames** 101
 - checking with **sp_checkreswords** 101
 - DLL file 225
 - remote user 191
 - server 49
 - server attribute 429

- user’s full 30
- naming
 - groups 27
 - time ranges 55
 - user-defined datatypes 61
- @@*ncharsize* global variable
 - sp_addtype** and 61
- nesting
 - cursors 139
- net password encryption** option
 - sp_serveroption** 365
- no chkpt on recovery** database option
 - setting with **sp_dboption** 148
- no free space acctg** database option
 - setting with **sp_dboption** 148
- not null values
 - sp_addtype** and 59
 - for user-defined data 59
- null values
 - sp_addtype** and 59
 - for user-defined datatypes 59
- number (quantity of)
 - databases reported by **sp_countmetadada** 137
 - groups per user 100
 - indexes 137
 - messages per constraint 83
 - open objects 137
- numbers
 - See also* IDs, user
 - datatype code 420
 - device 252
 - global variable unit 319
 - message 32, 67, 184, 226
 - ODBC datatype code 420
 - spid* (server process ID) 413
 - weekday names and 28

O

- object names, database
 - checking with **sp_checknames** 101
 - checking with **sp_checkreswords** 105
- object owners. *See* database object owners
- objects. *See* database objects; databases

Index

ODBC. *See* Open Database Connectivity (ODBC) API datatypes
official language name 29, 368
See also aliases; languages, alternate
Open Client applications
 connection security with 40
Open Database Connectivity (ODBC) API datatypes 420
operating system commands 442
optdiag utility
 flushing in-memory statistics 219
optimization
 queries (**sp_recompile**) 349
options
 See also configuration parameters
 database 143–150
 remote logins 351–352
 remote servers 365–368
order
 See also indexes; precedence; sort order
 of date parts 28
output
 packets, number of 320
output option
 sp_getmessage 226
overhead
 data caches 94
overlapping time ranges 57
owners. *See* database object owners; database owners
ownership
 dump devices and 63

P

@@pack_received global variable
 sp_monitor and 320
@@pack_sent global variable
 sp_monitor and 320
@@packet_errors global variable
 sp_monitor and 320
page locks
 types of 216, 295
pages, data
 computing number of, with **sp_spaceused** 386
 locks held on 216, 295
pair of columns. *See* common keys; joins

parameters, procedure
 ways to supply 3, 418
parentheses ()
 in SQL statements xv
 in user-defined datatypes 59
passthrough mode
 sp_autoconnect system procedure 73
 sp_passthru system procedure 329
 sp_remotesql system procedure 353
passwords
 date of last change 171
 encryption over network 367
 setting with **sp_addlogin** 30
 sp_password 331–332
 sp_remotoption and 351
 sp_serveroption and 367
 trusted logins or verifying 351
path name
 dump device 62
pattern matching
 catalog stored procedure parameters 419
PC DB-Library. *See* DB-Library programs
performance
 concurrency optimization 117
 information about 390
permissions
 displaying user's 169
 dump devices and 63
 granting 274
 information on 273
 new database owner 99
 new database user 309
 revoking 274
 sp_column_privileges information on 420
 system procedures 1
physical datatypes 59
physical device name 62
placeholders
 error message percent sign (%) 34
plan groups
 adding 37
 comparing 122
 copying 135
 copying to a table 212
 creating 37
 dropping 188

- dropping all plans in 175
- exporting 212
- information about 266
- reports 266
- plans
 - changing 372
 - comparing 122, 124
 - copying 135, 136
 - deleting 175
 - dropping 175, 189
 - finding 217
 - modifying 372
 - searching for 217
 - sp_showplan** output 382
- pools, memory
 - configuring 336
 - defaults 88
- precedence
 - binding defaults to columns and datatypes 80
 - resource limits 44
 - rule binding 85
- precision, datatype
 - sp_help** report on 234
 - user-defined datatypes 59
- prefetch
 - disabling 95
 - enabling 95
- primary keys
 - sp_dropkey** procedure 181
 - sp_foreignkey** and 224
 - sp_helpkey** and 262
 - sp_primarykey** definition of 341
- priority
 - sp_setpsex** 372
- “probe” login account 360
- probe process, two-phase commit 360
- procedures. *See* stored procedures; system procedures
- process logical name. *See* logical device name
- processes (server tasks)
 - checking locks held 293
 - checking locks on 215–217, 293–296
 - ID number 413
 - sp_showplan** display of 382–383
 - sp_who** report on 413–416
- promotion, lock 369
- protection system

- groups 27
- locking logins 297
- “public” group
 - See also* groups
 - information report 255
 - sp_addgroup** and 27
 - sp_adduser** and 65
 - sp_changegroup** and 100
 - sp_helpgroup** report on 255
- punctuation
 - enclosing in quotation marks 3, 418
 - in user-defined datatypes 59

Q

- queries
 - compilation and optimization 349
 - sp_tables** and 439
- query plans
 - recompiling with **sp_recompile** 349
- query processing
 - limiting with **sp_add_resource_limit** 41
 - modes 344–345
- quotation marks (“ ”)
 - enclosing parameter values 3, 418
 - enclosing reserved words 107
 - single, and **quoted_identifier** 113
- quoted identifiers
 - testing 107
 - using 106, 113–114

R

- range
 - specifying for resource limits 41
- range locks 296
- read only** database option
 - setting with **sp_dboption** 148
 - setting with **sp_setsuspect_granularity** 376
- readonly** option, **sp_serveroption** 365
- recompilation
 - stored procedures 349
- records, audit 18
- recovery

Index

- data caches and 90
- displaying mode 376
- forcing suspect pages online with **sp_forceonline_db** 219
- forcing suspect pages online with **sp_forceonline_page** 222
- listing offline pages 292
- listing suspect databases 291
- setting mode 376
- setting threshold 378
- recovery fault isolation 221, 291
- reference** auditing option 68
- reference information
 - catalog stored procedures 417
 - dbcc** stored procedures 457
 - system extended stored procedures 441
 - system procedures 1–4
- referencing, object. *See* dependencies, database object
- referential integrity constraints
 - binding user messages to 83
 - renaming 355–356
- regulations
 - for finding objects 161, 236
- reindex** option, **dbcc**
 - after **sp_indsuspect** 290
- remapping database objects 350
- remote logins
 - See also* logins; users
 - dropping 191–192
 - information on 270
 - sp_remotoption** for 351–352
 - sysremotelogins* table 38–40
 - trusted or untrusted mode 351
- remote procedure calls
 - sp_password** 332
- remote servers
 - See also* servers
 - changing names of 110, 112
 - dropping logins 191
 - information on 281
 - information on logins of 270
 - names of 49
 - passwords on 332
 - sp_remotoption** and 351–352
- remote users. *See* remote logins
- removing. *See* dropping; deleting
- renaming 355–356
 - See also* **sp_rename** system procedure
 - a database 356–358
 - warnings about 356, 358
- replacing user-defined messages 33
- reporting from *dbccdb* database
 - allocation statistics 471
 - comprehensive information 469
 - configuration information 460, 467, 469
 - fault information 465, 467
 - full details 469
 - I/O statistics 465
- reports
 - plan groups 266
 - sp_who** 413–416
- reserved words
 - catalog stored procedures and 418
 - as identifiers 101–114
 - system procedures and 3
- reservepagegap** option
 - sp_chgattribute** 116
 - sp_help** report on 237
- resource limits
 - creating 40
 - dropping 192
 - information about 270
 - modifying 310
 - types of 41
- retrieving
 - error message text 226
- return status
 - catalog stored procedures 418
 - sp_checkreswords** 105
 - system procedures 1
- reversing encryption of source text 288
- revoke** auditing option 68
- revoke** option, **sp_role** 362
- role hierarchies, displaying
 - using **sp_activeroles** 15
 - using **sp_displayroles** 171
- roles
 - displaying with **sp_activeroles** 15
- row lock promotion thresholds
 - dropping with **sp_droprowlockpromote** 195
 - setting with **sp_setrowlockpromote** 373
 - sp_helppdb** report on database setting 248

row locks 296

rows, table

- computing number of, with **sp_spaceused** 386
- limiting how many returned 41

rpc auditing option 69

rpc security model A option, **sp_serveroption** 365

rpc security model B option, **sp_serveroption** 365

rules

- binding 84–85
- changing names of 108
- checking name with **sp_checkreswords** 105
- displaying source text of 283
- naming user-created 84
- remapping 350
- renaming 355–356
- system tables and 85
- unbinding 409–410

S

scale, datatype

- in user-defined datatypes 59

scope of resource limits

- changes to active time ranges and 58
- information on 271
- specifying 43

security auditing option 69

security mechanism option, **sp_serveroption** 365

segments

- See also* database devices; log segment; space allocation
- adding 47–48
- changing names of 110, 112
- checking names with **sp_checkreswords** 105
- dropping 196–197
- extending 48, 213
- information about 279
- mapping 48
- monitoring remaining space 51–55, 314–318
- sp_helpthreshold** report on 285

select auditing option 69

server aliases 49

server information options. *See* information (server)

server process ID number. *See* processes (server tasks)

servers

- See also* processes (server tasks); remote servers
- adding 48–51
- attribute names 429
- dropping 197–198
- information on remote logins 270
- local 49
- monitoring activity of 318
- names of 49
- options, changing with **sp_serveroption** 365–368
- remote 281
- setting row lock promotion thresholds for 373
- sp_server_info** information on 429
- upgrading and **sp_checknames** 101
- upgrading and **sp_checkreswords** 105

set command

- sp_setlangalias** and **language** option 368

setting

- auditing options 68
- setuser** auditing option 69
- 7-bit terminal, **sp_helpsort** output 282
- shared locks 216, 295
- shared row locks 296
- single quotes. *See* quotation marks
- single user** database option
- setting with **sp_dboption** 149
- single-user mode 149
- sp_renamedb** and 357

size

- image* datatype 386
- log device 299
- text* storage 386

size of auto identity column configuration parameter

- 146, 150

sort order

- changing, and **sp_indsuspect** system procedure 290
- information about 282

source text

- checking for existence of 114
- displaying 283
- encryption, reversing 288
- hiding 287

sp_activeroles system procedure 15

sp_add_qpgroup system procedure 37

sp_add_resource_limit system procedure 40–45

sp_add_time_range system procedure 55–58

- sp_addalias** system procedure 17–18
- sp_addauditrecord** system procedure 18–19
- sp_addauditable** system procedure 20
- sp_addengine** system procedure 21
- sp_addexceclass** system procedure 22
- sp_addextendedproc** system procedure 23–24
- sp_addexternlogin** system procedure 24–26
- sp_addgroup** system procedure 26–27
- sp_addlanguage** system procedure 27–30
- sp_addlogin** system procedure 30–32
- sp_addmessage** system procedure 32–34
- sp_addobjectdef** system procedure 34–37
- sp_addremotelogin** system procedure 38–40
- sp_addsegment** system procedure 47–48
 - in mixed data and log databases 48
- sp_addserver** system procedure 48–51
- sp_addthreshold** system procedure 51–55
- sp_addtype** system procedure 58–62
- sp_addumpdevice** system procedure 62–64
- sp_adduser** system procedure 64–65
- sp_altermessage** system procedure 67
- sp_audit** system procedure 68–72
- sp_auditdisplay** system procedure 164–168
- sp_autoconnect** system procedure 72–74
- sp_bindcache** system procedure 75–78
- sp_bindefault** system procedure 78–80
 - create default and 79
- sp_bindexceclass** system procedure 80
- sp_bindmsg** system procedure 83
- sp_bindrule** system procedure 84–85
- sp_cacheconfig** system procedure 87–95
- sp_cachestrategy** system procedure 95–98
- sp_changedbowner** system procedure 98–99
- sp_changegroup** system procedure 99–100
 - sp_dropgroup** and 181
- sp_checknames** system procedure 100–101
- sp_checkreswords** system procedure 101–114
 - return status 105
- sp_checksourc** system procedure 114
- sp_chgattribute** system procedure 116–119
- sp_clearpsex** system procedure 120
- sp_clearstats** system procedure 121–122
- sp_cmp_all_qplans** system procedure 122
- sp_cmp_qplans** system procedure 124
- sp_column_privileges** catalog stored procedure 420–422
- sp_columns** catalog stored procedure 422–424
 - datatype code numbers 420
 - and **sp_datatype_info** 426
- sp_commonkey** system procedure 125–126
- sp_companion** system procedure 127–130
- sp_configure** system procedure 130–134
 - setting display levels for 168
- sp_copy_all_qplans** system procedure 135
- sp_copy_qplan** system procedure 136
- sp_countmetadata** system procedure 136
- sp_cursorinfo** system procedure 138–141
- sp_databases** catalog stored procedure 424
- sp_datatype_info** catalog stored procedure 425–426
- sp_dbcc_alterws** stored procedure 459–460
- sp_dbcc_configreport** stored procedure 460–461
- sp_dbcc_createws** stored procedure 461–463
- sp_dbcc_deletedb** stored procedure 463
- sp_dbcc_deletehistory** stored procedure 464
- sp_dbcc_differentialreport** stored procedure 465–466
- sp_dbcc_evaluatedb** stored procedure 466–467
- sp_dbcc_faultreport** stored procedure 467–469
- sp_dbcc_fullreport** stored procedure 469–470
- sp_dbcc_plandb** system procedure 333–335
- sp_dbcc_runcheck** stored procedure 470–471
- sp_dbcc_statisticsreport** stored procedure 471–474
- sp_dbcc_summaryreport** stored procedure 474–478
- sp_dbcc_updateconfig** stored procedure 478–480
- sp_dboption** system procedure 143–150
- sp_dbremap** system procedure 152–153
- sp_defaultloc** system procedure 153–155
- sp_depends** system procedure 156–161
- sp_deviceattr** system procedure 162–163
- sp_diskdefault** system procedure 163–164
- sp_displaylevel** system procedure 168–169
- sp_displaylogin** system procedure 169–171
- sp_displayroles** system procedure 171
- sp_drop_all_qplans** system procedure 175
- sp_drop_qpgroup** system procedure 188
- sp_drop_qplan** system procedure 189
- sp_drop_resource_limit** system procedure 192–195
- sp_drop_time_range** system procedure 199–200
- sp_dropalias** system procedure 175
- sp_dropdevice** system procedure 176–177
- sp_dropengine** system procedure 177
- sp_dropexceclass** system procedure 177
- sp_dropextendedproc** system procedure 178

- sp_dropexternlogin** system procedure (Component Integration Services only) 179
- sp_dropglockpromote** system procedure 180
- sp_dropgroup** system procedure 180–181
 - See also **sp_changegroup**
- sp_dropkey** system procedure 181–182
- sp_droplanguage** system procedure 182–183
- sp_droplogin** system procedure 183–184
- sp_dropmessage** system procedure 184–185
- sp_dropobjectdef** system procedure (Component Integration Services only) 187–188
- sp_dropremotelogin** system procedure 191–192
- sp_drowlockpromote** system procedure 195
- sp_dropsegment** system procedure 196–197
 - sp_placeobject** and 197
- sp_dropserver** system procedure 197–198
- sp_droptreshold** system procedure 198–199
- sp_droptype** system procedure 200
- sp_dropuser** system procedure 200–201
- sp_dumpoptimize** system procedure 203–207
- sp_estspace** system procedure 209–212
- sp_export_qpgroup** system procedure 212
- sp_extendsegment** system procedure 213–214
- sp_familylock** system procedure 215–217
- sp_find_qplan** system procedure 217–219
- sp_fkeys** catalog stored procedure 426–428
- sp_flushstats** system procedure 219
- sp_forceonline_db** system procedure 219–220
- sp_forceonline_object** system procedure 221–222
- sp_forceonline_page** system procedure 222–224
- sp_foreignkey** system procedure 224–225
- sp_freelI** system procedure 225–226
- sp_getmessage** system procedure 226–227
- sp_grantlogin** system procedure (Windows NT only) 227
- sp_ha_admin** system procedure 229–230
 - installing with *installhasvss* 230
- sp_help** system procedure 231–237
- sp_help_qpgroup** system procedure 266–268
- sp_help_qplan** system procedure 268–269
- sp_help_resource_limit** system procedure 270–273
- sp_helppartition** system procedure 237
- sp_helpcache** system procedure 239–240
- sp_helpconfig** system procedure 240–244
- sp_helpconstraint** system procedure 244–248
- sp_helppdb** system procedure 248–250
- sp_helpdevice** system procedure 251–252
- sp_helpextendedproc** system procedure 252–253
- sp_helpexternlogin** system procedure (Component Integration Services only) 253
- sp_helpgroup** system procedure 254–255
- sp_helpindex** system procedure 255–257
- sp_helpjava** system procedure 259–261
- sp_helpjoins** system procedure 261–262
- sp_helpkey** system procedure 262–264
- sp_hellanguage** system procedure 264–265
- sp_helplog** system procedure 265
- sp_helpobjectdef** system procedure (Component Integration Services only) 265
- sp_helpremotelogin** system procedure 270
- sp_helpprotect** system procedure 273–277
- sp_helpsegment** system procedure 279–281
- sp_helpserver** system procedure 281–282
- sp_helpsort** system procedure 282–283
- sp_helptext** system procedure 283–284
- sp_helpthreshold** system procedure 285
- sp_helpuser** system procedure 285–286
- sp_hidetext** system procedure 287
- sp_import_qpgroup** system procedure 289–290
- sp_indsuspect** system procedure 290–291
- sp_listsuspect_db** system procedure 291
- sp_listsuspect_object** system procedure 291–292
- sp_listsuspect_page** system procedure 292–293
- sp_lock** system procedure 293–296
- sp_locklogin** system procedure 297–298
- sp_logdevice** system procedure 298–300
 - log on** extension to **create database** and 299
- sp_loginconfig** system procedure (Windows NT only) 300–301
- sp_logininfo** system procedure (Windows NT only) 301–302
- sp_logiosize** system procedure 302
- sp_modify_resource_limit** system procedure 310–312
- sp_modify_time_range** system procedure 312–314
- sp_modifylogin** system procedure 307–310
- sp_modifythreshold** system procedure 314–318
- sp_monitor** system procedure 318–320
- sp_monitorconfig** system procedure 320–324
- sp_object_stats** system procedure 324–327
- sp_passthru** system procedure 329–330
- sp_password** system procedure 331–332

- sp_pkeys** catalog stored procedure 428
- sp_placeobject** system procedure 332–333
- sp_poolconfig** system procedure 336–341
- sp_primarykey** system procedure 341–342
 - sp_foreignkey** and 224
- sp_processmail** system procedure 342–344
- sp_procqmode** system procedure 344–345
- sp_procxmode** system procedure 346–347
- sp_recompile** system procedure 349
- sp_remap** system procedure 350
- sp_remotefoption** system procedure 351–352
- sp_remotesql** system procedure 353–355
- sp_rename** system procedure 355–356
- sp_rename_qpgroup** system procedure 359
- sp_renamedb** system procedure 109, 356–358
- sp_reportstats** system procedure 359–360
- sp_revokeloglein** system procedure (Windows NT only) 361
- sp_role** system procedure 361–362
- sp_sendmsg** system procedure 363–364
- sp_server_info** catalog stored procedure 429–431
 - sp_tables** and 439
- sp_serveroption** system procedure 365–368
- sp_set_qplan system procedure** 372
- sp_setlangalias** system procedure 368–369
- sp_setpglockpromote** system procedure 369–371
- sp_setpsexex** system procedure 371
- sp_setrowlockpromote** system procedure 373
- sp_setsuspect_granularity** system procedure 376–378
- sp_setsuspect_threshold** system procedure 378–379
- sp_showcontrolinfo** system procedure 379
- sp_showexexclass** system procedure 381
- sp_showplan** system procedure 382
- sp_showpsexex** system procedure 383
- sp_spaceused** system procedure 385–387
- sp_special_columns** catalog stored procedure 432
- sp_sproc_columns** catalog stored procedure 433
 - datatype code numbers 420
- sp_ssladmin** system procedure 387–389
- sp_statistics** catalog stored procedure 435
- sp_stored_procedures** catalog stored procedure 436
 - sp_server_info** information 431
- sp_syntax** system procedure 389–390
- sp_sysmon** system procedure 390–393
- sp_table_privileges** catalog stored procedure 437–438
- sp_tables** catalog stored procedure 439
 - sp_server_info** information 431
- sp_thresholdaction** system procedure 393–395
 - threshold procedure 53, 316
- sp_transactions** system procedure 395–402
- sp_unbindcache** system procedure 403–405
- sp_unbindcache_all** system procedure 405
- sp_unbinddefault** system procedure 405–406
- sp_unbindexexclass** system procedure 407
- sp_unbindmsg** system procedure 408
- sp_unbindrule** system procedure 409–410
- sp_volchanged** system procedure 410–413
- sp_who** system procedure 413–416
 - columns returned 415
- space
 - See also* size; space allocation
 - estimating table and index size 209–212
 - monitoring remaining with **sp_modifythreshold** 314–318
 - sp_spaceused** procedure 385–387
 - unused 387
- space allocation
 - See also* database devices; segments
 - future 332–333
 - log device 299
 - sp_placeobject** procedure 332–333
- space management properties
 - changing with **sp_chgattribute** 116
- spid* number
 - sp_who** output 415
- spt_committab* table 4
- spt_datatype_info* table 419
- spt_datatype_info_ext* table 419
- spt_monitor* table 4
- spt_server_info* table 419
- spt_values* table 4
- SQL standards
 - SQL pattern matching 419
 - user-defined datatypes and 59
- square brackets []
 - in SQL statements xv
- starting days of named time ranges 55
- starting times of named time ranges 56
- statistics
 - flushing to *systabstats* 219
 - returned by global variables 318
 - sp_clearstats** procedure 121

- sp_monitor** 318
- sp_reportstats** 359–360
- status
 - database device 163
- stored procedures
 - See also* database objects; system procedures
 - cache binding and 77, 404
 - catalog 417–439
 - changing transaction modes with **sp_procxmode** 346–347
 - for *dbccdb* database 457
 - displaying query processing modes with **sp_procqmode** 344–345
 - object dependencies and 156–161
 - remapping 350
 - renamed database and 358
 - renaming 355–356
 - sp_checkreswords** and 106
 - sp_recompile** and 349
 - sp_sproc_columns** information on 433
 - sp_stored_procedures** information on 436
- suspect databases, listing 291
- suspect indexes
 - forcing online 221, 291
- suspect pages
 - bringing online 219–220, 222–224
 - isolating on recovery 376–378, 378–379
 - listing 292
- sybdiagdb* database 243
- sybsyntax* database 390
- sybsystemprocs* database
 - permissions and 2
- symbols
 - in SQL statements xiv, xv
- syntax
 - catalog stored procedures 418–419
 - checking for reserved words 105
 - display procedure (**sp_syntax**) 389–390
- syntax conventions, Transact-SQL xiv
- sysalternates* table
 - aliases 17
 - sp_dropalias** and 175
 - sysusers* table and 17
- syscomments* table
 - source text in 284
- sysconstraints* table
 - sp_bindmsg** and 83
- sysdatabases* table 424
- sysdevices* table 163, 251
- syskeys* table
 - sp_dropkey** and 182
 - sp_foreignkey** and 224
 - sp_primarykey** and 341
- syslanguages* table 265
 - sp_droplanguage** and 182
- syslkstats* table 326
- syslogs* table 298
 - put on a separate device 299
- sysmessages* table
 - error message text 226
- sysremotelogins* table 38–40, 198
 - sp_droptremotelogin** and 191
- sysresourcelimits* table
 - applicable limits for a login session 44
 - sp_help_resource_limit** and 273
- sys.servers* table
 - sp_addserver** and 49
 - sp_helpserver** and 282
- sys.sessions*
 - removing old entries 230
- sys.tabstats* table
 - flushing statistics to 219
- system extended stored procedures 441–455
- system procedure tables 4
 - catalog stored procedures and 419
- system procedures
 - catalog stored 417–439
 - changing names of 108
 - displaying source text of 283
 - displaying syntax of 389–390
 - extended stored procedures 441–455
 - help reports 231–286
 - list of 1–13
 - permissions 1
 - return status 1
 - using 1
- system procedures results. *See* information (server)
- system roles
 - displaying with **sp_active roles** 15
- system* segment
 - dropping 197
 - mapping 48

Index

- system tables
 - binding to caches 77
 - defaults and 79
 - direct updates dangerous to 110
 - rules and 85
 - space allocation 333
 - updating 1
 - sysimeranges* table
 - ID number storage in 57
 - range name storage in 41
 - systypes* table 200
 - sysusermessages* table
 - error message text 226
 - sp_dropmessage** and 184
 - sysusers* table
 - sysalternates* table and 17
- ## T
- table_access** auditing option 69
 - tables
 - binding to data caches 75
 - changing names of 108
 - checking name with **sp_checkreswords** 105
 - column information 422
 - column permission information from
 - sp_column_privileges** 420–422
 - common key between 125–126
 - dropping keys between 181
 - dropping row lock promotion thresholds for 195
 - estimating space for 209
 - joined common key 125–126
 - lock promotion thresholds for 370
 - locks held on 216, 295
 - locks, types of 216, 295
 - object dependencies and 156–161
 - primary keys on 341
 - renaming 355–356
 - setting row lock promotion thresholds for 374
 - sp_placeobject** space allocation for 332–333
 - sp_recompile** and 349
 - sp_table_privileges** information on 437–438
 - sp_tables** 439
 - space used by 386
 - with suspect indexes 290
 - system procedure 4, 419
 - unbinding from data caches 403
 - tape dump devices
 - adding 62–64
 - tape** option, **sp_addumpdevice** 62
 - tempdb* database
 - auto_identity** database option and 146
 - unique auto_identity index** database option and 150
 - temporary names. *See* alias, user
 - temporary tables
 - sp_help** and 236
 - system procedure 4
 - terminals
 - 7-bit, **sp_helpsort** output example 282
 - 8-bit, **sp_helpsort** output example 283
 - text
 - copying with **defncopy** 107
 - user-defined message 33
 - text* datatype
 - size of storage 386
 - @@thresh_hysteresis** global variable
 - threshold placement and 52
 - threshold procedures 53
 - creating 393
 - executing 54, 317
 - parameters passed to 53, 317
 - thresholds
 - adding 51–55
 - changing 314–318
 - crossing 52
 - disabling 55, 199, 318
 - hysteresis value 52, 316
 - information about 285
 - last-chance 53, 199, 316
 - maximum number 53, 316
 - optimization for reducing I/O 117
 - removing 198–199
 - row lock promotion 374
 - space between 53
 - time interval
 - estimating index creation 209
 - limiting 41
 - since **sp_monitor** last run 319
 - time ranges
 - adding 55

- “at all times” 57, 199
- changing active time ranges 57
- creating 55
- dropping 199
- entire day 56
- IDs for 57
- modifying 312
- overlapping 57
- timeouts** option, **sp_serveroption** 365
- @total_errors** global variable
 - sp_monitor** and 320
- @total_read** global variable
 - sp_monitor** and 320
- @total_write** global variable
 - sp_monitor** and 320
- transaction logs
 - data caches and 339
 - log I/O size and 339
 - on a separate device 298–300
 - thresholds and 199
- transactions
 - modes 346–347
- Transact-SQL
 - reserved words 105
- translation
 - of user-defined messages 34
- triggers
 - changing names of 108
 - checking name with **sp_checkreswords** 105
 - displaying source text of 283
 - object dependencies and 156–161
 - remapping 350
 - renamed database and 358
 - renaming 355–356
 - sp_recompile** and 349
- true | false** clauses
 - sp_dboption** 143
 - sp_remotoption** 351
- true** option, **sp_changedbowner** 98
- trunc log on chkpt** database option 149
- truncate** auditing option 69
- trusted mode
 - remote logins and 40
- trusted** option, **sp_remotoption** 351
- two-phase commit
 - probe process 360

U

- UDP messaging 363
- unbind** auditing option 69
- unbinding
 - data caches 403–405
 - defaults 405–406
 - objects from caches 403–405
- unencrypting source text 288
- unique auto_identity index** database option 150
- unlocking login accounts 297
- unmapping a segment from a database 196–197
- unused space
 - sp_spaceused** reporting of 387
- update** auditing option 69
- update row locks 296
- us_english language 29
- usage statistics 359
- use message confidentiality** server option 365
- use message integrity** server option 365
- user context for operating system commands
 - (**xp_cmdshell**) 443
- User Datagram Protocol messaging 363
- user IDs
 - changing with **sp_import_qpgroup** 289
 - displaying 171
 - dropping with **sp_droplogin** and 184
- user names
 - See also* database object owners; logins
 - changing 110
 - checking with **sp_checkreswords** 105
- user permissions. *See* database owners; permissions
- user-created objects. *See* database objects
- user-defined audit records 68
- user-defined datatypes
 - binding defaults to 78–80
 - binding rules to 84
 - changing names of 109
 - checking name with **sp_checkreswords** 105
 - creating 58–62
 - dropping 200
 - hierarchy 61
 - naming 61
 - unbinding defaults from 405–406
 - unbinding rules with **sp_unbindrule** 409–410
- user-defined event logging (**xp_logevent**) 446
- user-defined messages 32–34

- unbinding with **sp_unbindmsg** 408
- user-defined procedures
 - creating ESPs with **sp_addextendedproc** 23
- user-defined roles
 - displaying with **sp_activeoles** 15
- users
 - accounting statistics 121, 360
 - adding 30–32, 64–65
 - change group for 99–100
 - changing names of 112, 307–310
 - dropping aliased 175
 - dropping from databases 200–201
 - dropping from servers 183–184
 - dropping remote 198
 - information on 169, 285
 - password change for accounts 331–332
 - permissions of 273
 - remote 270
 - sp_who** report on 413–416
 - system procedure permissions and 2
 - sysusers* table 17
- utility commands
 - See also Utility Programs* manual
 - display syntax 389–390

V

- values
 - displaying with **sp_server_info** 429
- view_access** auditing option 69
- views
 - checking name with **sp_checkreswords** 105
 - columns 422
 - common key between 125–126
 - displaying source text of 283
 - dropping keys between 181
 - object dependencies and 156–161
 - primary keys on 341
 - remapping 350
 - renamed database and 358
 - renaming 108, 355–356
- virtual page numbers 252
- volume handling 410

W

- wash area
 - configuring 340
 - defaults 340
- wash** keyword, **sp_poolconfig** 336
- weekday date value
 - first 28
 - names and numbers 28
- wildcard characters
 - SQL standards pattern matching (\$ and _) 419
- workspaces
 - dropping 462

X

- XP Server 442
 - freeing memory from 225–226
- xp_cmdshell context** configuration parameter 443
- xp_cmdshell** system extended stored procedure 442
- xp_deletemail** system extended stored procedure 444
 - sp_processmail** and 343
- xp_enumgroups** system extended stored procedure 444
- xp_findnextmsg** system extended stored procedure 445
 - sp_processmail** and 343
- xp_logevent** system extended stored procedure 446
- xp_readmail** system extended stored procedure 447
 - sp_processmail** and 343
- xp_sendmail** system extended stored procedure 450
 - sp_processmail** and 343
- xp_startmail** system extended stored procedure 453
- xp_stopmail** system extended stored procedure 454