

SYBASE®

New Features

Adaptive Server Enterprise

Version 12.5.3

DOCUMENT ID: DC00212-01-1253-01

LAST REVISED: November 2004

Copyright © 1989-2004 by Sybase, Inc. All rights reserved.

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

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

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

Sybase, the Sybase logo, AccelaTrade, 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, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-Translator, APT-Library, Backup Server, BizTracker, ClearConnect, Client-Library, Client Services, Convoy/DM, Copernicus, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DataWindow .NET, DB-Library, dbQueue, Developers Workbench, Direct Connect Anywhere, DirectConnect, Distribution Director, e-ADK, E-Anywhere, e-Biz Impact, e-Biz Integrator, E-Whatever, EC Gateway, ECMAP, ECRTP, eFulfillment Accelerator, 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, eProcurement Accelerator, EWA, Financial Fusion, Financial Fusion Server, Gateway Manager, GlobalFIX, iAnywhere, iAnywhere Application Alerts, iAnywhere Mobile Delivery, iAnywhere Mobile Document Viewer, iAnywhere Mobile Inspection, iAnywhere Mobile Marketing Channel, iAnywhere Mobile Pharma, iAnywhere Mobile Sales, iAnywhere Pylon, iAnywhere Pylon Application Server, iAnywhere Pylon Conduit, iAnywhere Pylon PIM Server, iAnywhere Pylon Pro, iAnywhere Solutions, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, Mail Anywhere Studio, MainframeConnect, Maintenance Express, Manage Anywhere Studio, M-Business Channel, M-Business Network, M-Business Server, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, My iAnywhere, My iAnywhere Media Channel, My iAnywhere Mobile Marketing, MySupport, Net-Gateway, Net-Library, New Era of Networks, ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Biz, Open Client, Open ClientConnect, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, Orchestration Studio, PB-Gen, PC APT Execute, PC DB-Net, PC Net Library, PocketBuilder, Pocket PowerBuilder, 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, Rapport, RepConnector, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Report-Execute, Report Workbench, Resource Manager, RW-DisplayLib, RW-Library, S-Designor, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILLS, 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, S.W.I.F.T. Message Format Libraries, 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, TotalFix, TradeForce, Transact-SQL, Translation Toolkit, UltraLite, UltraLite.NET, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, 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. 05/04

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., One Sybase Drive, Dublin, CA 94568.

Contents

About This Book	v
CHAPTER 1 Dumping and Loading Databases Across Platforms	1
Overview	1
Dump and load across platforms with the same endian architecture	2
Dump and load across platforms with different endian architecture.	2
Dumping a database	2
Loading a database.....	3
Restrictions.....	3
Performance Notes	4
CHAPTER 2 Importing Statistics For Proxy Tables	7
Limitations	7
CHAPTER 3 top n functionality	9
CHAPTER 4 New Functionality in Historical Server.....	11
Sending data to a database server	11
Setting up the receiving Adaptive Server	11
Starting Historical Server.....	11
Viewing the data.....	12
Data storage.....	12
CHAPTER 5 Secure Sockets Layer.....	15
Advanced Encryption Standard (AES) Algorithm.....	15
Setting SSL cipher suite preferences.....	16
Examples sp_ssladmin.....	20
Other considerations	21
@@ssl_ciphersuite	21
CHAPTER 6 Real Time Data Services Enhancements	23

CHAPTER 7	HP-UX.....	25
CHAPTER 8	Resource Governor	27
CHAPTER 9	Page Allocation for Partitioned DOL Tables	29
CHAPTER 10	User Connections	31
	Number of user connections	31
	User disconnections.....	31
CHAPTER 11	dtdValidation	33
CHAPTER 12	Migration Tool	35
CHAPTER 13	New Language Support.....	37
CHAPTER 14	Monitor Counters and sp_sysmon.....	39
	sp_sysmon noclear option	39
	Monitor counter concurrency.....	40
	New dbcc commands.....	41
CHAPTER 15	Changes to Stored Procedures, Functions, and Commands	43
	New stored procedure	43
	sp_post_xpload	43
	Changes to stored procedures	44
	sp_sysmon	44
	New function	44
	Changed commands	45
	alter table	45
	create table	45
	sp_help_resource_limit	46
	dbcc commands	46
Index		49

About This Book

Audience

System Administrators, Database Administrators, and System Security Officers may find this book useful as it covers the new features available in Adaptive Server® Enterprise version 12.5.3.

How to use this book

This book includes the following:

- Chapter 1, “Dumping and Loading Databases Across Platforms,” includes information about dumping and loading databases from one platform to another.
- Chapter 2, “Importing Statistics For Proxy Tables,” describes how to import statistics.
- Chapter 3, “top n functionality,” discusses the compatibility for the MicroSoft SQL, ASA, and ASIQ servers, in limiting the number of rows returned by a query.
- Chapter 4, “New Functionality in Historical Server,” describes how to download monitoring information from Historical Server to an Adaptive Server.
- Chapter 5, “Secure Sockets Layer,” describes the new cipher suites and administration options now available for SSL.
- Chapter 6, “Real Time Data Services Enhancements,” describes how to use `sp_configure` to set the number of native threads.
- Chapter 7, “HP-UX,” describes updates for HP-UX Itanium.
- Chapter 8, “Resource Governor,” describes changes for the resource governor, which prevents queries and transactions from monopolizing server resources.
- Chapter 9, “Page Allocation for Partitioned DOL Tables,” describes the new mechanism for space allocation.
- Chapter 10, “User Connections,” describes changes to identifying the reason for a user disconnect.
- Chapter 11, “dtdValidation,” describes the expansion of the `dtdValidation`.

-
- Chapter 12, “Migration Tool,” describes changes for sybmigrate across versions of Adaptive Server.
 - Chapter 13, “New Language Support,” discusses additional languages supported in Enhanced Full Text Search (EFTS).
 - Chapter 14, “Monitor Counters and sp_sysmon,” describes improvements to the monitoring in Adaptive Server.
 - Chapter 15, “Changes to Stored Procedures, Functions, and Commands,” has information about new and changed commands, stored procedures, and functions.

Related documents

The Sybase Adaptive Server Enterprise documentation set comprises these documents:

- 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.
- *What’s New in Adaptive Server Enterprise?* – describes the new features in Adaptive Server version 12.5.1, the system changes added to support those features, and the changes that may affect your existing applications.
- *ASE Replicator User’s Guide* – describes how to use the ASE Replicator feature of Adaptive Server to implement basic replication from a primary server to one or more remote Adaptive Servers.
- *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.
- *Configuring Adaptive Server Enterprise* for your platform – provides instructions for performing specific configuration tasks for Adaptive Server.
- *EJB Server User’s Guide* – explains how to use EJB Server to deploy and execute Enterprise JavaBeans in Adaptive Server.

- *Error Messages and Troubleshooting Guide* – explains how to resolve frequently occurring error messages and describes solutions to system problems frequently encountered by users.
- *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.
- *Glossary* – defines technical terms used in the Adaptive Server documentation.
- *Historical Server User's Guide* – describes how to use Historical Server to obtain performance information for SQL Server[®] and Adaptive Server.
- *Java in Adaptive Server Enterprise* – describes how to install and use Java classes as data types, functions, and stored procedures in the Adaptive Server database.
- *Job Scheduler User's Guide* – provides instructions on how to install and configure, and create and schedule jobs on a local or remote Adaptive Server using the command line or a graphical user interface (GUI).
- *Monitor Client Library Programmer's Guide* – describes how to write Monitor Client Library applications that access Adaptive Server performance data.
- *Monitor Server User's Guide* – describes how to use Monitor Server to obtain performance statistics from SQL Server and Adaptive Server.
- *Performance and Tuning Guide* – is a series of four books that explains how to tune Adaptive Server for maximum performance:
 - *Basics* – the basics for understanding and investigating performance questions in Adaptive Server.
 - *Locking* – describes how the various locking schemas can be used for improving performance in Adaptive Server.
 - *Optimizer and Abstract Plans* – describes how the optimizer processes queries and how abstract plans can be used to change some of the optimizer plans.
 - *Monitoring and Analyzing* – explains how statistics are obtained and used for monitoring and optimizing performance.
- *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.

-
- *Reference Manual* – is a series of four books that contains the following detailed Transact-SQL[®] information:
 - *Building Blocks* – Transact-SQL datatypes, functions, global variables, expressions, identifiers and wildcards, and reserved words.
 - *Commands* – Transact-SQL commands.
 - *Procedures* – Transact-SQL system procedures, catalog stored procedures, system extended stored procedures, and dbcc stored procedures.
 - *Tables* – Transact-SQL system tables and dbcc tables.
 - *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.
 - *System Tables Diagram* – illustrates system tables and their entity relationships in a poster format. Available only in print version.
 - *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.
 - *Using Adaptive Server Distributed Transaction Management Features* – explains how to configure, use, and troubleshoot Adaptive Server DTM features in distributed transaction processing environments.
 - *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.
 - *Utility Guide* – documents the Adaptive Server utility programs, such as isql and bcp, which are executed at the operating system level.
 - *Web Services User's Guide* – explains how to configure, use, and troubleshoot Web Services for Adaptive Server.
 - *XA Interface Integration Guide for CICS, Encina, and TUXEDO* – provides instructions for using the Sybase DTM XA interface with X/Open XA transaction managers.

- *XML Services in Adaptive Server Enterprise* – describes the Sybase native XML processor and the Sybase Java-based XML support, introduces XML in the database, and documents the query and mapping functions that comprise XML Services.

Other sources of information

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

- The Getting Started CD contains release bulletins and installation guides in PDF format, and may also contain other documents or updated information not included on the Technical Library CD. It is included with your software. To read or print documents on the Getting Started CD, you need Adobe Acrobat Reader (downloadable at no charge from the Adobe Web site, using a link provided on the CD).
- The Technical Library CD contains product manuals and is included with your software. The DynaText reader (included on the Technical Library CD) 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.

- The 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 EBFs/Maintenance, Technical Documents, Case Management, Solved Cases, newsgroups, and the Sybase Developer Network.

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

Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

❖ Finding the latest information on product certifications

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Select Products from the navigation bar on the left.
- 3 Select a product name from the product list and click Go.
- 4 Select the Certification Report filter, specify a time frame, and click Go.
- 5 Click a Certification Report title to display the report.

Sybase EBFs and software maintenance

❖ **Creating a personalized view of the Sybase Web site (including support pages)**

Set up a MySybase profile. MySybase is a free service that allows you to create a personalized view of Sybase Web pages.

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click MySybase and create a MySybase profile.

❖ **Finding the latest information on EBFs and software maintenance**

- 1 Point your Web browser to the Sybase Support Page at <http://www.sybase.com/support>.
- 2 Select EBFs/Maintenance. If prompted, enter your MySybase user name and password.
- 3 Select a product.
- 4 Specify a time frame and click Go. A list of EBF/Maintenance releases is displayed.

Padlock icons indicate that you do not have download authorization for certain EBF/Maintenance releases because you are not registered as a Technical Support Contact. If you have not registered, but have valid information provided by your Sybase representative or through your support contract, click Edit Roles to add the “Technical Support Contact” role to your MySybase profile.

- 5 Click the Info icon to display the EBF/Maintenance report, or click the product description to download the software.

Conventions

In the regular text of this document, the names of files and directories appear in *italics*, for example:

- In Windows NT: *%SYBASE%\bin*
- In UNIX platforms: *\$SYBASE*

Note Substitute your Sybase installation drive and directory for *\$SYBASE* in UNIX, and *%SYBASE%* in Windows NT.

Table 1 details the typographic (font and syntax) conventions as used in this document.

Table 1: Font and syntax conventions for this document

Element	Example
Command names, command option names, database names, datatypes, utility names, utility flags, and other keywords are in Helvetica.	dsedit
Variables, or words that stand for values that you fill in, are in <i>italics</i> .	select <i>column_name</i> from <i>table_name</i> where <i>search_conditions</i>
<i>Parentheses</i> must be typed as part of the command.	compute row_aggregate (column_name)
<i>Curly braces</i> indicate that at least one of the enclosed options is required by the command (see comma).	{cheese, sauce} Note Do not type the curly braces.
<i>Brackets</i> mean that choosing one or more of the enclosed options is optional.	[anchovies, pineapple, bell_peppers] Note Do not type the brackets.
The <i>vertical bar</i> means you may select only one of the options shown.	{cash check credit} Note Do not type the curly braces.
The <i>comma</i> means you may choose as many of the options shown as you like; be sure to separate multiple choices in a command with commas.	[extra_cheese, avocados, sour_cream] Note Do not type the brackets.
An <i>ellipsis</i> (...) means that you can <i>repeat</i> the unit that the ellipsis follows as many times as you like.	buy <i>thing</i> = price [cash check credit] [, <i>thing</i> = price [cash check credit]]... <ul style="list-style-type: none"> • You must buy at least one <i>thing</i> (item) and give its price. • You may choose a method of payment: one of the options enclosed in square brackets. • You may choose also to buy additional items: as many of them as you like. For each item you buy, provide its name, its price, and (optionally) a method of payment.
Syntax statements, which display the utility's syntax including all its options, appear as shown here, either in a serif font for flags and options (-v), or italics for user-supplied values (<i>username</i>).	charset [-P <i>password</i>] [-S <i>server</i>] [-I <i>interface</i>] <i>sort_order</i> <i>charset</i>
Examples that illustrate computer output appear in Courier, as shown:	pub_id pub_name city state ----- 0736 New Age Books Boston MA 0877 Binnet & Hardley Washington DC (2 rows affected)

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.

Dumping and Loading Databases Across Platforms

Overview

Adaptive Server Enterprise version 12.5.2 supported the dump and load of databases across platforms with the same endian architecture.

With Adaptive Server version 12.5.3, you can now dump and load databases across platforms with different endian architecture. This means you can perform dump database and load database from either a big endian platform to a little endian platform, or from a little endian platform to a big endian platform.

In a big-endian system, the most significant byte of storage, such as integer or long, has the lower address. The reverse is true for a little-endian system.

There is no syntax change with dump or load database in version 12.5.3. Adaptive Server automatically detects the architecture type of the originating system of the database dump file during a load database, then performs the necessary conversions. Loads from older versions, such as 11.9 and 12.0, are also supported. The dump and load can be from 32-bit to 64-bit platforms, and vice versa.

Platforms supported:

Big-endian	Sun Solaris	IBM AIX	Silicon Graphics	HP-UX	MAC OS X
Little-endian	Linux IA	Windows	HP True 64	Sun Solaris x86	

Dump and load across platforms with the same endian architecture

When you perform dump database and load database across platforms with the same endian architecture, user and system data do not require conversions. There are no limitations on operations with the dump and load of a database.

Stored procedures and other compiled objects are recompiled from the SQL text in *syscomments* at the first execution after the load database for certain combination platforms.

Dump and load across platforms with different endian architecture

Adaptive Server allows a dump and load database between big endian and little endian architectures and vice versa.

Dumping a database

Before you run dump database, for a cross platform dump and load, use the following procedures to move the database to a transactional quiescent status:

- 1 Verify the database runs cleanly by executing `dbcc checkdb` and `dbcc checkalloc`.
- 2 To prevent concurrent updates from open transactions by other processes during dump database, use `sp_dboption` to place the database in a single-user mode.
- 3 Flush statistics to `systabstats` using `sp_flushstats`.
- 4 Wait for 10 to 30 seconds, depending on the database size and activity.
- 5 Run checkpoint against the database to flush updated pages.
- 6 Run dump database.

Loading a database

Once you load the database, Adaptive Server automatically identifies the endian type on the dump file and performs all necessary conversions during the load database and online database.

Note After Adaptive Server converts the index rows, the order of index rows may be incorrect. Adaptive Server marks following indexes on user tables as suspect indexes during online database.

- Non-clustered index on APL table.
- Clustered index on DOL table.
- Non-clustered index on DOL table.

See “sp_post_xpload” on page 43 for information about checking and rebuilding indexes on user tables.

Restrictions

- dump transaction and load transaction is not allowed across platforms.
- dump database and load database to or from a remote backupserver are not supported across platforms.
- You cannot load a password-protected dump file across platforms.
- If you perform dump database and load database for a parsed XML object, you must parse the text again after the load database is completed.
- You cannot perform dump database and load database across platforms on Adaptive Servers versions earlier than 11.9.
- Adaptive Server cannot translate embedded data structures stored as binary, varbinary, or image columns.
- load database is not allowed on the master database across platforms.
- Stored procedures and other compiled objects are recompiled from the SQL text in syscomments at the first execution after the load database.

If you do not have permission to recompile from text, then the person who does has to recompile from text using `dbcc upgrade_object` to upgrade objects.

Note If you migrate login records in *syslogins* system table in the master database from Solaris to Linux, you can use `bcp` with character format. The login password from the Solaris platform is compatible on Linux without a trace flag from this release. For all other combinations and platforms, login records need to be recreated because the passwords are not compatible.

Performance Notes

Due to the design of indexes within a datasever that provides an optimum search path, index rows are ordered for fast access to the table's data row. Index rows which contain row identifiers (RIDs), are treated as binary to achieve a fast access to the user table.

Within the same architecture platform, the order of index rows remains valid and search order for a selection criteria takes its normal path. However, when index rows are translated across different architectures, this invalidates the order by which optimization was done. This results in an invalid index on user tables when the cross platform dump and load feature is performed.

A database dump from a different architecture, such as big endian to little endian, is loaded, certain indexes are marked as suspect:

- Non-clustered index on APL table.
- Clustered index on DOL table.
- Non-clustered index on DOL table.

To fix indexes on the target system, after load from a different architecture dump, you could use one of two methods:

- 1 Drop and recreate all of the indexes.
- 2 Use `sp_post_xpload`, see page 43.

Since the data point and information varies from usage on indexes, the schema, user data, number of indexes, index key length, and number of index rows, in general, it requires planning to recreate indexes on large tables as it can be a lengthy process. `sp_post_xpload` validates indexes, drops invalid indexes, and recreates dropped indexes, in a single command on databases.

Since `sp_post_xpload` performs many operations it can take longer than drop and recreate indexes. Sybase recommends that you use the drop and recreate indexes on those databases larger than 10G.

Importing Statistics For Proxy Tables

In Adaptive Server version 12.5.3, when you perform update statistics on a remote server proxy table, if the relevant table and index statistics are available, the table catalogs are imported to the local systabstats and sysstatistics.

By default, update statistics for proxy tables always attempts to import the required statistics data. But when the statistics data is unavailable or incomplete on the remote table, Component Integration Services (CIS) reverts to the prior mechanism of gathering statistic data.

You can also force CIS to revert to the prior mechanism of gathering statistic data by turning on Traceflag 11229. This allows you to obtain all data from the database, then calculating the statistics.

Limitations

Key limitations:

- The proxy table must be mapped to another Adaptive Server version 11.9 or later.
- Excludes proxy tables mapped to RPCs, external files, and system directories.
- If the remote servers are not Adaptive Server Enterprise version 11.9 or later, or of another server class, CIS continues to obtain statistics data using prior mechanisms.

top n functionality

top n functionality implements compatibility with MicroSoft SQL server, Adaptive Server Anywhere, and ASIQ.

Use the top n clause to limit the number of rows in the result set to the number of rows specified by the integer. The integer can be any unsigned 32-bit value, in the range 0 through $2^{32}-1$ (4GB-1 or 4,294,967,295). Zero indicates “no” rows.

Adaptive Server Enterprise, version 12.5.3 supports the top n clause in outer query select statements, but not in the select list of a subquery. This differs from the MicroSoft SQL server. Any attempt to use the top n clause in a subquery yields a syntax error.

Use top n with select...into statements to limit the number of rows inserted in the target table. This is different from set rowcount, which is ignored during a select...into.

select statement

This is the syntax for the top n clause a select statement:

```
SELECT
all_distinct_clause
TOP unsigned_integer
select_list
into_clause
from_where_clause
group_by_clause
having_clause
```

Example:

```
select top 5 col1 from t1
```

union statements

The syntax to use the top n clause in a select statement with a union is:

```
select top 2 c1 from t1
union all
select top 3 c2 from t2
```

This returns five rows, assuming t1 has at least 2 rows and t2 has at least 3 rows. The top limit applies to the individual selects that form a union, not to the union as a whole.

update statement In an update statement insert the top n clause immediately after the keyword:

```
UPDATE
TOP unsigned_integer
object_identifier
SET
set_clause_list
from_where_clause
for_clause
abstract_plan_clause
```

delete statement In a delete statement insert the top n clause immediately after the keyword:

```
DELETE
TOP unsigned_integer
result_table
from_where_clause
for_clause
abstract_plan_clause
```

- Usage
- When used with delete, update, or in a view, you cannot specify ordering. If there is an implied order on the table from a clustered index, that order applies, otherwise, the results are unpredictable as they can be in any order.
 - When used with cursors, top n limits the overall size of the result set. If you specify set cursor rowcount, top n limits the results of a single fetch.
 - When a view definition contains select top n and a query with a where clause uses it, the results may be inconsistent.

New Functionality in Historical Server

Adaptive Server version 12.5.3 allows you to send monitoring data from Historical Server to a database on a specific Adaptive Server.

Sending data to a database server

Setting up the receiving Adaptive Server

Once you select an Adaptive Server:

- 1 Create a database for the Historical Server to use for storing monitoring data. The default for the name of the database is *hs_monitoring*. If you choose to have a different name for the database, create that named database, then change the database name in the `hs_directload.sql` script.
- 2 Run the install script `hs_directload.sql` on the database. The install script creates two catalog tables, sessions and views, and the stored procedure `sp_hs_dboutput`.

Starting Historical Server

You must specify:

- The destination for the monitoring data from Historical Server to a specific Adaptive Server and database, rather than a flat file.
- The user name and password for the output if they are not the same as those on Historical Server. If the user name and the password are not specified, they default to the ones specified in the `-U` and `-P` parameters on the command line.

The target Adaptive Server and database for the historical data must be available when Historical Server is started.

The command line syntax is:

```
histserver -U<user name>  
- P<password> -D<output dir>  
-l<log file> -I<interfaces file>  
[-d<delimiter>] [-O<ASE name>] [-o<DATABASE name>] [-f ]  
-u<outputASE user name> - p<outputASEpassword>
```

- -O *ASE name* – the name of the target Adaptive Server.
- -o *DATABASE name* – the database name to which the monitoring data is sent, if it is not *hs_monitoring*, identify the database name on this option.
- -u *outputASE user name* – login name for the connection to the Adaptive Server with the output.
- -p *outputASEpassword* – password for the login name for the output Adaptive Server.
- -f – must be used when the -O option is specified if you want Historical Server to send data to files in the output directory as well as the database.

You must have access to Historical Server and update permissions to the target database.

Viewing the data

Data storage

The data from Historical Server is in the same format as defined when the view was created on the server. To store the data, Historical Server creates two system tables:

- Sessions table provides a record of every recording session that used the output database.
- Views table lists the views that were used by each recording session.

The table structures are similar to the structure of the output data files, where:

- First column is the monitoring session ID.

- Second column is the monitored server name.
- Third column is a date/timestamp.
- Subsequent columns per data item specified in the view definition.

This structure, for the date and timestamp, and the data items, is identical to the structure defined in the DDL scripts Historical Server provides when a bulk copy is executed on Historical Server data files into another Adaptive Server.

Secure Sockets Layer

Adaptive Server Enterprise version 12.5.3 has added the following changes for Secure Sockets Layer (SSL):

- New cipher suites using the Advanced Encryption Standard (AES) algorithm
- New options to `sp_ssladmin` to set preferences for cipher suites accepted by Adaptive Server
- New global variable `@@ssl_ciphersuite` to tell the client which cipher suite was chosen by the SSL handshake

These enhancements improve the ability of the System Security Officer to manage SSL, and for client applications to determine the encryption algorithms used on their connection.

Advanced Encryption Standard (AES) Algorithm

Two new cipher suites are available in Adaptive Server version 12.5.3. They are:

- `TLS_RSA_WITH_AES_256_CBC_SHA`
- `TLS_RSA_WITH_AES_128_CBC_SHA`

These cipher suites use the Advanced Encryption Standard (AES), which is now available for Secure Sockets Layer (SSL) in Adaptive Server version 12.5.3. AES is a FIPS-197 - approved standard for symmetric encryption. The AES algorithm provides the strongest encryption available for Adaptive Server.

The AES algorithm and dependent cipher suites used in SSL are available by default in Adaptive Server 12.5.3 and Open Client 12.5.1. No action is required by users or administrators to use this algorithm unless algorithms are currently restricted client applications. If you have clients that currently restrict algorithms, you may want to consider using AES algorithm instead.

Setting SSL cipher suite preferences

In Adaptive Server version 12.5.3, `sp_ssladmin` has two new command options, `lsciphers` and `setciphers`. With these new options, the set of cipher suites that Adaptive Server uses can be restricted, giving control to the System Security Officer over the kinds of encryption algorithms that may be used by client connections to the server or outbound connections from Adaptive Server. The default behavior for use of SSL cipher suites in Adaptive Server is the same as in earlier versions; it uses an internally defined set of preferences for cipher suites.

To display the values for any set cipher suite preferences, enter:

```
sp_ssladmin lsciphers
```

To set a specific cipher suite preference, enter:

```
sp_ssladmin setciphers, {"FIPS" | "Strong" | "Weak" |  
"All" | quoted_list_of_ciphersuites }
```

where:

- “FIPS” – is the set of encryptions, hash, and key exchange algorithms that are FIPS-compliant. The algorithms included in this list are AES, 3DES, DES, and SHA1.
- “Strong” – is the set of encryption algorithms using keys longer than 64 bits.
- “Weak” – is the set of encryption algorithms from the set of all supported cipher suites that are not included in the strong set.
- “All” – is the set of default cipher suites.
- `quoted_list_of_ciphersuites` – specifies a set of cipher suites as a comma-separated list, ordered by preference. Use quotes (“”) to mark the beginning and end of the list. The quoted list can include any of the predefined sets as well as individual cipher suite names. Unknown cipher suite names cause an error to be reported, and no changes are made to preferences.

The detailed contents of the predefined sets are in Table 5-1 on page 18.

`sp_ssladmin setciphers` sets cipher suite preferences to the given ordered list. This restricts the available SSL cipher suites to the specified set of “FIPS”, “Strong”, “Weak”, “All”, or a quoted list of cipher suites. This takes effect on the next listener started, and requires that you restart Adaptive Server to ensure that all listeners use the new settings.

You can display any cipher suite preferences that have been set using `sp_ssladmin lsciphers`. If no preferences have been set, `sp_ssladmin lsciphers` returns 0 rows to indicate no preferences are set and Adaptive Server uses its default (internal) preferences.

Table 5-1: Predefined cipher suite sets in Adaptive Server version 12.5.3

Set name	Cipher suite names included in the set
FIPS	TLS_RSA_WITH_AES_256_CBC_SHA TLS_RSA_WITH_AES_128_CBC_SHA TLS_RSA_WITH_3DES_EDE_CBC_SHA TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA TLS_RSA_WITH_DES_CBC_SHA TLS_DHE_DSS_WITH_DES_CBC_SHA TLS_DHE_RSA_WITH_DES_CBC_SHA TLS_RSA_EXPORT1024_WITH_DES_CBC_SHA TLS_DHE_DSS_EXPORT1024_WITH_DES_CBC_SHA
Strong	TLS_RSA_WITH_AES_256_CBC_SHA TLS_RSA_WITH_AES_128_CBC_SHA TLS_RSA_WITH_3DES_EDE_CBC_SHA TLS_RSA_WITH_RC4_128_SHA TLS_RSA_WITH_RC4_128_MD5 TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA TLS_DHE_DSS_WITH_RC4_128_SHATLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA
Weak	TLS_RSA_WITH_DES_CBC_SHA TLS_DHE_DSS_WITH_DES_CBC_SHA TLS_DHE_RSA_WITH_DES_CBC_SHA TLS_RSA_EXPORT1024_WITH_DES_CBC_SHA TLS_RSA_EXPORT1024_WITH_RC4_56_SHA TLS_DHE_DSS_EXPORT1024_WITH_RC4_56_SHA TLS_DHE_DSS_EXPORT1024_WITH_DES_CBC_SHA TLS_RSA_EXPORT_WITH_RC4_40_MD5 TLS_RSA_EXPORT_WITH_DES40_CBC_SHA TLS_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA TLS_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA

Set name	Cipher suite names included in the set
All	TLS_RSA_WITH_AES_256_CBC_SHA TLS_RSA_WITH_AES_128_CBC_SHA TLS_RSA_WITH_3DES_EDE_CBC_SHA TLS_RSA_WITH_RC4_128_SHA TLS_RSA_WITH_RC4_128_MD5 TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA TLS_DHE_DSS_WITH_RC4_128_SHA TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA TLS_RSA_WITH_DES_CBC_SHA TLS_DHE_DSS_WITH_DES_CBC_SHA TLS_DHE_RSA_WITH_DES_CBC_SHA TLS_RSA_EXPORT1024_WITH_DES_CBC_SHA TLS_RSA_EXPORT1024_WITH_RC4_56_SHA TLS_DHE_DSS_EXPORT1024_WITH_RC4_56_SHA TLS_DHE_DSS_EXPORT1024_WITH_DES_CBC_SHA TLS_RSA_EXPORT_WITH_RC4_40_MD5 TLS_RSA_EXPORT_WITH_DES40_CBC_SHA TLS_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA TLS_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA

Warning! Sybase recommends that you do not use any cipher suites that are not included in predefined sets because they pose potential security vulnerabilities.

Table 5-2: Cipher suites to avoid

Reason to avoid	Cipher suites
Cipher suites with “anon” for certificate exchange algorithm; the server is not authenticated by its certificate.	TLS_DH_anon_WITH_3DES_EDE_CBC_SHA TLS_DH_anon_WITH_RC4_128_MD5 TLS_DH_anon_WITH_DES_CBC_SHA TLS_DH_anon_EXPORT_WITH_DES40_CBC_SHA TLS_DH_anon_EXPORT_WITH_RC4_40_MD5
Cipher suites with “NULL” for the symmetric key encryption algorithm do not encrypt data for transmission on the network.	TLS_RSA_WITH_NULL_SHA TLS_RSA_WITH_NULL_MD5

Examples `sp_ssladmin`

On initial startup, before any cipher suite preferences have been set, no preferences are shown by `sp_ssladmin lscipher`.

```
1> sp_ssladmin lscipher
2> go
```

Output:

```
 Cipher Suite Name   Preference
-----
(0 rows affected)
(return status = 0)
```

The following example specifies the set of cipher suites that use FIPS algorithms.

```
1> sp_ssladmin setcipher, 'FIPS'
2> go
```

A preference of 0 (zero) `sp_ssladmin` output indicates a cipher suite is not used by Adaptive Server. The other, non-zero numbers, indicate the preference order that Adaptive Server uses the algorithm during the SSL handshake. The client side of the SSL handshake chooses one of these cipher suites that matches its list of accepted cipher suites.

This example uses a quoted list of cipher suites to set preferences in Adaptive Server:

```
1> sp_ssladmin setcipher,
'TLS_RSA_WITH_AES_128_CBC_SHA,
TLS_RSA_WITH_AES_256_CBC_SHA'
```

```
2> go
```

Other considerations

When you upgrade to Adaptive Server version 12.5.3, the cipher suite preferences are the server defaults, and `sp_ssladmin` option `iscipher` displays no preferences. The server uses its default preferences, those defined by "All". The System Security Officer should consider the security policies employed at his or her site and the available SSL cipher suites to decide whether to restrict cipher suites and which cipher suites are appropriate for the security policies.

If you downgrade from Adaptive Server version 12.5.3 to an earlier version, any cipher suite preferences are ignored, and the Adaptive Server defaults for that version are used.

If you have set SSL cipher suite preferences and want to remove all preferences from the server and use default preferences, delete the preferences from their storage location in system catalogs using the following commands:

```
1> sp_configure 'allow updates to system tables', 1
2> go

1> delete from master..sysattributes where class=24
2> go

1> sp_configure 'allow updates to system tables', 0
2> go
```

These commands can be executed only by the System Security Officer or System Administrator.

@@ssl_ciphersuite

The Transact-SQL® global variable `@@ssl_ciphersuite` has been added to allow users to know which cipher suite was chosen by the SSL handshake and verify that an SSL or a non-SSL connection was established.

For example, an `isql` connection using SSL protocol displays the cipher suite chosen for it.

```
1> select @@ssl_ciphersuite
2> go
```

Output:

```
-----  
TLS_RSA_WITH_AES_128_CBC_SHA  
  
(1 row affected)
```

Real Time Data Services Enhancements

In Adaptive Server version 12.5.3, use `sp_configure` to set the number for native threads and the wait time for messaging.

max native threads per engine

Summary information	
Default value	50
Maximum values	1000
Status	Dynamic
Display level	Intermediate
Required role	System Administrator

Use to define the maximum number of native threads the server spawns per engine. When the limit for the native threads is reached, Adaptive Server sessions that require a native thread, sleep until another session releases a native thread.

rtm thread idle wait period

Summary information	
Default value	50 seconds
Maximum value	4026531839 seconds
Status	Dynamic
Display level	Intermediate
Required role	System Administrator

Use to define the time a native thread used by Adaptive Server waits when it has no work to do. When the time set for a native thread is reached, the thread automatically fades out.



This chapter contains information on HP-UX for Adaptive Server Enterprise version 12.5.3. This platform includes the following for the release of Adaptive Server:

- Feedback Optimization (FBO) – On the Itanium platform, Adaptive Server Enterprise supports the FBO server on HP-11.23 Itanium, which improves performance.
- High Availability – On the Itanium platform, the High Availability option is available on HP-UX 11.23 with HP ServiceGuard A.11.15. For more information, see *Using Sybase Failover in A High Availability System* for more information.
- For HP-UX, XML services are enabled for the HP-UX 32-bit and 64-bit platforms.



Resource Governor

Adaptive Server Enterprise provides resource limits to help System Administrators prevent queries and transactions from monopolizing server resources. Resource limits, however, are not fully specified until they are bound to a time range.

A resource limit is a set of parameters specified by a System Administrator to prevent an individual login or application from:

- Exceeding a given estimated or actual I/O costs.
- Returning excessive rows on a per query basis.
- Exceeding a given elapsed time on a query batch or transaction basis.
- Utilizing excessive tempdb space per session.

In Adaptive Server Enterprise version 12.5.3, when the System Administrator modifies a resource limit, all users logged in the session see the change, including the System Administrator.



Page Allocation for Partitioned DOL Tables

In prior versions of Adaptive Server Enterprise there was a problem where inserts into a partitioned data-only locked table with a placement index may have caused wasted space. The space wastage is more pronounced in tables with a higher number of partitions and also for larger page sizes.

Adaptive Server Enterprise version 12.5.3 avoids wasting extra space by filling up existing allocated extents in the target allocation page even though these extents are assigned to other partitions. The net effect is that we only allocate new extents when there are no free extents in the target allocation page.

To turn off the new space allocation mechanism, activate command line trace flag 646 (-T646).



User enhancements done for Adaptive Server Enterprise version 12.5.3 include an updated error message and a correction on reserved sockets.

Number of user connections

When Sybase introduced Enterprise Java Beans (EJB) it required some of Adaptive Server's internal socket structures. This feature simply reserved one third of available sockets to provide for EJB, regardless of whether or not EJB was configured. Once reserved, these sockets were unavailable for normal Adaptive Server work.

With Adaptive Server version 12.5.3, no sockets are automatically reserved for EJB, it must be done manually.

User disconnections

An enhancement to error 1608 now displays the host name and login name when a user is disconnected from Adaptive Server. The System Administrator can now analyze and confirm the cause of the disconnect. The error message appears as follows:

```
00:00000:00016:2004/11/08 16:00:58.52 kernel  Cannot send, host process
disconnected: TONYI-XP 1496 suid: 1
00:00000:00016:2004/11/08 16:00:58.52 kernel  Cannot send, host process
disconnected: TONYI-XP 1496 suid: 1
00:00000:00016:2004/11/08 16:00:59.08 server  Error: 1608, Severity: 18, State:
4
00:00000:00016:2004/11/08 16:00:59.08 server  A client process exited
abnormally, or a network error was encountered. Unless other errors occurred,
continue processing normally.
00:00000:00017:2004/11/08 16:01:13.98 server  DBCC TRACEON 7844, SPID 17
00:00000:00017:2004/11/08 16:01:21.03 kernel  Cannot send, host process
```

User disconnections

```
disconnected: TONYI-XP 512 suid: 1
00:00000:00017:2004/11/08 16:01:23.10 server Error: 1608, Severity: 18, State:
4
00:00000:00017:2004/11/08 16:01:23.10 server A client process exited
abnormally, or a network error was encountered. Unless other errors occurred,
continue processing normally.
00:00000:00017:2004/11/08 16:01:23.10 kernel extended error information:
hostname: TONYI-XP login: sa
```

For Adaptive Server Enterprise version 12.5.3, the `dtdValidate` option expanded to include:

- `dtdValidate='no'`
No validation is performed whether the document has an embedded DTD or references an external DTD. If the document contains an embedded DTD, `dtdValidate` verifies that it is correct but does not validate.
- `dtdValidate='yes'`
Validates only if there is an embedded DTD or a reference to an external DTD.
- `dtdValidate='strict'`
The document must contain an embedded or reference to an external DTD, and validates against the DTD.



Migration Tool

The `sybmigrate` tool allows you to migrate across versions of Adaptive Server Enterprise and supports source servers from 12.0 through 12.5.3.

When migrating from a source server with Adaptive Server Enterprise versions 12.0 through 12.5.0.1, you must include the size of and location of the database on the target server.



New Language Support

The Enhanced Full-Text Search Specialty Data Store (EFTS) now includes language support for:

- Traditional Chinese on the Windows and Solaris platforms
- Arabic, Hebrew, Thai, and Russian on the Linux platform



Monitor Counters and *sp_sysmon*

In Adaptive Server version 12.5.3, there are a number of enhancements to improve some of the most commonly used monitoring tools.

With the introduction of the *noclear* parameter for *sp_sysmon* and the monitor counter concurrency improvements, you can now run multiple concurrent sessions of the *sp_sysmon* and other monitoring applications, such as Monitor Server, Historical Server or other *sp_sysmon* sessions. By using the *noclear* parameter, you can ensure that *sp_sysmon* does not alter the values of monitor counters used by other applications.

sp_sysmon noclear option

noclear is a new parameter to *sp_sysmon*. When the *noclear* option is included in the *sp_sysmon*, *sp_sysmon* does not zero out the monitor counters. As a result, data collected by other applications using the monitor counters at the same time are not be affected by the *sp_sysmon* report.

When used with the *noclear* parameter, *sp_sysmon* utilizes a temporary table. However, if space in the user's default temporary database is limited, you may need to increase the size of the database to run *sp_sysmon* with *noclear*.

The report contents generated from *sp_sysmon* with or without the *noclear* should be the same. However, since the *noclear* parameter requires storing some data in a temporary table, the *sp_sysmon* report may show a small amount of additional server activity when run with the *noclear* parameter compared to execution without it.

A field called “Sample Mode” has been added to the sp_sysmon report header to indicate whether the report was generated using “No Clear” or “Reset Counters” mode. The report header indicates the sample start and end times with the labels “Sampling Started at” and “Sampling Ended at” when run with the noclear parameter. It displays “Statistics Cleared at” and “Statistics Sampled at” when run without the noclear parameter.

When noclear is used with a specific sample interval, for any parameter in the sp_sysmon, the command runs in noclear mode. If noclear is not specified, sp_sysmon clears the counters.

Syntax `sp_sysmon interval [, noclear,[,section [, applmon]]]`

Example Report usage without clearing the counters:

```
sp_sysmon "00:01:00", kernel, noclear
sp_sysmon "00:01:00", noclear
```

Note You can use the noclear parameter only when you specify a sample interval in sp_sysmon. If you specify begin_sample or end_sample you cannot use noclear.

Monitor counter concurrency

Adaptive Server version 12.5.3 tracks the number of applications using the monitor counters. Adaptive Server does not disable or stop data collection by the monitor counters as long as one or more applications are known to be using them. This allows applications such as sp_sysmon and the Monitor Server to operate concurrently.

This change does not require any modifications on applications using monitor counters. Adaptive Server tracks the number of different applications or user connections where the monitor counters have been enabled or disabled. The monitor counters are disabled only when all connections have disabled them.

The monitor counter usage count is not automatically decremented when an application that enabled the monitor counters logs out of Adaptive Server. This means that, if an application enables the monitor counters and does not disable them before logging out, the usage count continues to reflect the same number of users as it did prior to the application logging out. This can be corrected by using new dbcc commands described in the following section.

New *dbcc* commands

The System Administrator can now manually modify monitor counter usage count. When an application enables monitor counters, then fails to disable them before logging off Adaptive Server, the System Administrator can use these commands to terminate monitor counter data collection:

```
dbcc monitor (increment, <group name>)
dbcc monitor (decrement, <group name>)
dbcc monitor (reset, <group name>)
```

Where *<group name>* can be one of the following:

- 'all'
- *spinlock_s*
- *appl*

increment and *decrement* increase and decrease usage counts for the monitor counters in the specified group by 1. *reset* sets the usage count for the monitor counters in the specified group to zero. This turns off collection of monitoring data for this group.

You can determine the usage count for all, which comprises most of the monitor counters, by selecting the *@@monitors_active* global variable.

The usage counts for the *spinlock_s* and *appl* groups are reported by the *dbcc resource* command.

Changes to Stored Procedures, Functions, and Commands

This chapter describes new and changed stored procedures, functions, and commands for Adaptive Server version 12.5.3.

New stored procedure

sp_post_xpload

Description	checks and rebuilds indexes after a cross-platform database load where the endian types are different.
Syntax	<code>sp_post_xpload</code>
Example	Once the database is loaded from another platform, rebuild its indexes by executing: <code>sp_post_xpload</code>
Options	None.
Permissions	Can be only executed by a System Administrator.
Comments	<ul style="list-style-type: none"> • The following indexes are rebuilt on all user tables in the database: <ul style="list-style-type: none"> • Nonclustered index on an APL table • Clustered index on a DOL table • Nonclustered index on a DOL table • Indexes on system tables are not processed with this stored procedure. System table indexes are rebuilt when online database is executed. • You can also rebuild indexes using drop index and create index. • Only run this stored procedure when the database is loaded across platforms with different endians.

- Where the index status is suspect, reset the index by executing the `sp_xpload`, `drop index`, or `create index`.
- Stored procedures are recompiled from the SQL text in syscomments at the first execution after the load database. Use `dbcc upgrade_object` to upgrade objects if you do not have permission recompile from text.

Changes to stored procedures

sp_sysmon

When the `noclear` option is included in the `sp_sysmon`, `sp_sysmon` does not zero out the monitor counters.

Syntax `sp_sysmon interval [, noclear[,section [, applmon]]]`

Example Report usage without clearing the counters:

```
sp_sysmon "00:01:00", kernel, noclear
sp_sysmon "00:01:00", noclear
```

Note You can use the `noclear` parameter only when you specify a sample interval in `sp_sysmon`. If you specify `begin_sample` or `end_sample` you cannot use `noclear`.

See Chapter 14, “Monitor Counters and `sp_sysmon`” for more information.

New function

Function `getutcdate()`

Description Returns a date and time whose value is in Universal Coordinated Time (UTC), also known as Greenwich Mean Time (GMT). `getutcdate` is calculated each time a row is inserted or selected.

Syntax `insert t1 (c1, c2, c3) select c1, getutcdate(), getdate() from t2`

Changed commands

Identity columns are used to store sequential numbers, such as invoice numbers or employee numbers, that are generated automatically by Adaptive Server. The value of the identity column uniquely identifies each row in a table.

identity indicates that the column has the identity property. Each table in a database can have one identity column with a datatype of integer, small integer, tiny integer or numeric with a scale of 0. Identity columns are not updatable and do not allow nulls.

alter table

Sets the identity column to int, smallint, tinyint, or numeric.

Syntax

```
alter table [[database.][owner].]table_name
    { add column_name datatype
      [default {constant_expression | user | null}]
      (id { int | smallint | tinyint | numeric(n) }
        {identity | null | not null} ...
    }
...

```

Parameters id – the identity column is set as integer or numeric.

Example

```
alter table
    id numeric

```

Permissions The System Administrator, Database Administrator and table owner use this clause when altering a table.

create table

Sets the identity column to int, smallint, tinyint, or numeric.

Syntax

```
create table [database .][owner .].]table_name (column_name datatype
    [default {constant_expression | user | null}]
    (id { int | smallint | tinyint | numeric(n) }
      {{{identity | null | not null}}}
      [off row | [ in row [ (size_in_bytes) ] ]
    .....
    )

```

Parameters id – the identity column is set as integer or numeric

```
create table.....  
id{int)
```

Permissions The System Administrator, Database Administrator and Table Owner use this clause when creating a table.

sp_help_resource_limit

Description When the parameter @verbose is used, the output is displayed in the verbose mode.

Syntax sp_help_resource_limit [name [, appname[, limittime[, limitday [, scope[, action [, verbose]]]]]]]]]

Parameters @verbose – value is 1 or 0 (zero).

Example To list all limits in verbose mode:

```
sp_help_resource_limit null,null,null,null,null,null,1
```

Another way to list all limits in verbose mode:

```
sp_help_resource_limit @verbose=1
```

Usage To list all resource limits in verbose mode.

dbcc commands

New commands are now available to allow the System Administrator to manually modify the monitor counter usage count. Under normal conditions it would not be necessary for the System Administrator to use these commands. However, in cases where an application enables the monitor counters, then fails to disable them before logging off Adaptive Server, the System Administrator needs to use these commands to terminate monitor counter data collection.

The new DBCC commands are:

```
dbcc monitor (increment, <group name>)  
dbcc monitor (decrement, <group name>)  
dbcc monitor (reset, <group name>)
```

Where <group name> can be one of the following:

- 'all'
- spinlock_s

- appl

The increment and decrement commands increase and decrease the usage counts for the monitor counters in the specified group by 1. The reset command sets the usage count for the monitor counters in the specified group to zero. This will turn off collection of monitoring data for this group.

The usage count for the all group, which comprises most of the monitor counters, can be determined by selecting the @@*monitors_active* global variable.

The usage counts for the spinlock_s and appl groups are reported by the dbcc resource command.

Index

A

- allocation
 - space 29
- alter table**
 - identity column 45

C

- cipher suites
 - Secure Sockets Layers 15
- configuration parameter
 - max native threads per engine** 23
 - rtm thread idle wait period** 23
- conventions, syntax x, xi
- counters
 - sp_sysmon** 39
- create table**
 - identity column 45

D

- disconnection
 - user 31
- dtdValidate 33
- dump and load database
 - across platforms 1–4
- dump database**
 - across platforms 2

E

- EJB server
 - reserved sockets 31

F

- fbo. *See* HP-UX Itanium
- feedback optimization
 - HP-UX Itanium 25
- font conventions x, xi
- functions
 - top-N 9

G

- getutcdate** to obtain the GMT 44

H

- Historical Server 11
- HP-UX Itanium
 - feedback optimization 25
 - XML 25

I

- identity column
 - alter table** 45
 - create table** 45
- import statistics
 - proxy tables 7

L

- languages 37
- load**
 - rebuild indexes 43
- load database**
 - across platforms 3

M

max native threads per engine configuration parameter 23
migration 35

N

native threads
sp_configure 23

P

proxy tables
import statistics 7
update statistics 7

R

rebuild indexes
sp_post_xoload 43
reserved sockets
EJB server 31
resource governor 27
rtm thread idle wait period configuration parameter 23

S

Secure Sockets Layers
cipher suites 15
sp_configure
native threads 23
sp_post_xoload
rebuild indexes 43
sp_sysmon 44
sp_sysmoncounters 39
space
allocation 29
SSL *See* Secure Sockets Layers
sybmigrate 35
syntax
conventions x

T

top-N functionality 9

U

update statistics
proxy tables 7
user
disconnection 31

X

XML
HP-UX Itanium 25