Release Bulletin
Adaptive Server® Enterprise Version 15.0.1 for Linux

Document ID: DC78565-01-1501-01

Last revised: August 25, 2006

1. Accessing current release bulletin information

A more recent version of this release bulletin may be available on the Web. To check for critical product or document information added after the release of the product CD, use the Sybase™ Technical Library Product Manuals Web site.

❖ **Accessing release bulletins at the Technical Library Product Manuals Web site**

1 Go to Product Manuals at http://www.sybase.com/support/manuals

2 Select Adaptive Server™ Enterprise and the appropriate language and click Go.

3 Select the appropriate version.

4 Select the Release Bulletins link.

5 Select the release bulletin for your platform. You can either download the PDF version or browse the document online.

2. Product summary

Enclosed is Sybase Adaptive Server Enterprise version 15.0.1. Server and client components are distributed on separate CDs.

For details on system requirements, including disk space and RAM, see the installation guide for your platform.

2.1 Installation kit

The installation kit includes:

- The Server CD – for contents list, see the installation guide for your platform.
- The PC-Client CD – contains software client components to be installed on Windows 2000, Windows XP Pro, and Windows 2003 computers.
- The Getting Started CD with the following documentation, specific to your platform:
  - Installation guide
  - Configuration guide
  - Release bulletin (this document)

2.2 Operating system requirements

Adaptive Server Enterprise version 15.0.1 has the following minimum operating system requirements:

- RHEL release 3 update 3:
  - kernel 2.4.2.1-20.EL
  - kernel-hugemem 2.4.21-20.EL (for Large Memory Support, LMS)
  - glibc-2.3.2.95-27
  - compat-libstdc++-7.3-2.96.128
- SLES 9 – patch level 2

- For details on different Linux distributions supported and operation system levels required, see the Red Hat Web site at http://certification.sybase.com/ucr/search.do.

Note The latest RH3 update that includes the required patch is RH3.0 Update 6 after the RH kernel 2.4.21-27.6.EL. The patch must be used, at the minimum, on a 64-bit kernel level. See "Running on RH3.0 64-bit distribution occasionally stops responding" on page 31for more detail.

Note Adaptive Server running on Linux on Power does not support the Job Scheduler, XA Interface, Directory Services, and Web Services. In addition, the following optional packages are currently not available on Linux on Power: High Availability (HA), Distributed Transaction Management (DTM), and Enhanced Full Text Search (EFTS).

3. Special installation instructions

Adaptive Server Enterprise 15.0.1 includes new versions of many supporting components. The most current version of other Sybase products (such as Enterprise Connect™ Data Access 12.6) contain earlier versions of those same components. Installing Adaptive Server version 15.0.1 into the same directory with existing products should not impact the existing products.

Sybase strongly recommends that you install Adaptive Server version 15.0.1 into its own directory. Where this is not practical, install Adaptive Server version 15.0.1 *last*.

See "Highlighted known installation issues" on page 29 for known installation problems.

Warning! Sybase recommends that you not install an older version of a Sybase product in a *$SYBASE* directory that already includes a newer Sybase product as this may not work. For example, attempting to install Replication Server 12.6 on top of Adaptive Server 15.0.1 may make either Adaptive Server Enterprise or Replication Server inoperable depending on the choices taken during install.

If the installation stops responding, see "ClearCase causes InstallShield to stop responding" on page 29 for more information.

3.1 Installing Adaptive Server interim releases on top of existing installations

When an Adaptive Server interim release is applied on top of an existing Adaptive Server high availability installation, it overwrites files that are modified and saved during the original Adaptive Server high availability installation.

The Adaptive Server interim release replaces:

- *ase_login_file* located in the *$SYBASE/ASE-12_5/SC-3_0/etc* directory
- *SY.ase* located in the *$SYBASE/ASE-12_5/SC-3_0/etc* directory
- *sybhauser* located in the *$SYBASE/ASE-12_5/install* directory

To avoid problems complete the following tasks before any interim release installation:

1 Save a copy of *ase_login_file*, *SY.ase*, and *sybhauser*.

2 Install the interim release.

3 Modify the new versions of the files installed by interim release using the saved copies of the files. Make any appropriate changes to the file contents.

4 Set the file ownership and permissions as existed prior to applying the interim release.

3.2 Using SySAM 2.0

Starting with Adaptive Server Enterprise version 15.0, the Sybase Software Asset Management System (SySAM) implementation has changed. With these changes, SySAM configuration is no longer optional, and requires some changes in the Adaptive Server Enterprise installation and configuration process. Review the updated SySAM implementation and plan your SySAM deployment before you install Adaptive Server Enterprise.

See the *Sybase Software Asset Management User's Guide* for complete details.

Warning! Adaptive Server Enterprise works for an initial period of 30 days without proper SySAM configuration. If your configuration does not identify a valid license within this 30-day grace period, Adaptive Server Enterprise shuts down.

3.3 Runtime libraries required for Linux on Power installation

Warning! You cannot install Linux on Power without the following procedure.

❖ **Installing Visual Age xlc compiler runtime libraries**

Adaptive Server is compiled and linked with the Visual Age compiler, xlc version 7.0. Before installing Adaptive Server, you must install the runtime libraries for the Visual Age compiler.

1 Confirm that the runtime libraries are installed by entering:

```
rpm -q vacpp.rte-7.0.0
```

If this command returns `vacpp.rte-7.0.0`, the Visual Age runtime packages are installed and you can skip the following steps.

If the command returns the message that `vacpp.rte-7.0.0` is not installed, you must install them.

2 Go to the IBM Web site,
at http://www-1.ibm.com/support/docview.wss?uid=swg24007906

Follow the instructions to download and install the runtime libraries on your operating system. Select the package that is appropriate for your Linux distribution. For details on different Linux distributions supported and operating system level required, go to the Sybase certification web site at http://certification.sybase.com/ucr/search.do.

3 After installing the xlc runtime libraries, verify that LD_LIBRARY_PATH is set to include the xlc runtime library path.

3.4 Monitoring and Diagnostic (MDA) tables

The *installmontables* script for installing Monitoring and Diagnostic tables (MDA) has been modified to now use the materialized parameter when defining the MDA proxy tables. This change results in improved performance for many queries using the MDA tables, making it no longer necessary to define a “loopback” remote server in your sysservers table in order to use the MDA tables.

Sybase recommends you reinstall the MDA tables using the *installmontables* script contained in this release.

Note No changes are needed to the *installmontables* script unless you plan to use the MDA tables to monitor remote servers.

❖ **Monitoring remote servers using MDA tables**

Make these changes to use the MDA tables to monitor remote servers:

1 Update the *installmontables* script by changing all occurrences of "materialized" to "external procedure."

2 Modify the at clause in the proxy table definitions in the *installmontables* script to add the name of the remote server followed by the string "..." before the MDA RPC name.

3 You can identify MDA RPC names by the "$" (dollar sign) preceding the name. The remote server name must be the same as the remote server name as it appears in the *interfaces* file on the server where the MDA tables are installed.

For example, for the monState table definition, if your remote server name appearing in the *interfaces* file is SALESPROD, change:

```
at "$monState"
```

to:

```
at "SALESPROD...$monState"
```

If the remote server has not already been registered with the server on which you are installing the MDA tables, use the sp_addserver stored procedure to do so.

When installing the MDA tables in order to reference a remote server as described above, be sure to use the version of the *installmontables* script that came with the version of the Adaptive Server Enterprise installed at the remote server.

Use of the materialized directive has changed the way some error messages from the MDA tables are sent to the client application. If all configuration parameters required by an MDA table used in a query are not enabled, then the client application receives an error message.

4. Special upgrade instructions

This section describes special upgrade instructions for Adaptive Server.

4.1 Upgrading compiled objects

Sybase recommends that you read "Upgrading compiled objects with dbcc upgrade_object" in Chapter 6, "Troubleshooting," in the *Adaptive Server Installation Guide for Linux*.

4.2 Disabling High Availability during upgrade

When High Availability (HA) is enabled, upgrading the master database in one server (for example, PRIMARY) attempts to apply the same changes to the other server (for example, SECONDARY).

The result is that you cannot upgrade syspartitionkeys and sysslices because those tables do not yet exist. Consequently, the first server that is upgraded cannot upgrade the master database.

The upgrade on the second server succeeds, although you see several warning messages about the privileges not being revoked (message 4622, state 1). When you re-run the upgrade process on the first server, it then succeeds, although it also produces the same messages.

To avoid this problem:

1 Disable HA in the *old* secondary server:

```
sp_companion primary-server-name, 'drop'
sp_configure 'enable HA', 0
```

2 Disable HA in the *old* primary server:

```
sp_configure 'enable HA', 0
```

3 Restart the servers.

4 Perform the upgrade to re-run installmaster.

5 In the *new* primary and secondary server, enter:

```
sp_configure 'enable HA', 1
<reboot>
```

6 Re-run installhasvss on both servers.

7 In the *new* secondary server, enter:

```
sp_companion primary-server-name,
    configure[, with_proxydb]
```

8 Restart both servers.

4.3 Upgrading High Availability and cluster support

Adaptive Server version 15.0.1 supports the following cluster platforms for High Availability:

- HP-UX – MCSG 11.15
- IBM AIX – HACMP 5.2
- Sun Solaris – VCS4.0, SunCluster 3.0/3.1
- Linux – VCS4.0
- Win2000 – Cluster Manager 5.0
- Win2003 – Cluster Manager 5.2

Adaptive Server does not support SunCluster2.2 and VCS1.3/2.0 HA agents on Sun Solaris. If you are using these agents, upgrade the respective cluster versions to configure Adaptive Server 15.0.1 for High Availability on Sun Solaris.

4.3.1 Upgrading the cluster subsystem

You can upgrade the cluster subsystem in two ways:

- Perform a major upgrade, which involves cluster down time, then shutting down and restarting all servers. In this case, you must:
  a Run sp_companion suspend, as described in *Using Sybase Failover in a High Availability System*.
  b Take the resource groups of primary and secondary companions on both nodes offline. Make sure this shuts down both the companion server and corresponding resource groups, and that they are not automatically brought online until the cluster system upgrade is complete.
  c Upgrade the cluster subsystem following the instructions from the cluster system vendor.

You may find options to migrate the current resource groups to the new cluster version. If such an option is not available (or if the resource groups get deleted or corrupted), re-create the resource groups and configure them appropriately after you have upgraded the cluster system.

d Bring the resource groups online. This should bring primary and secondary companions online on their respective nodes.

e Run sp_companion resume as described in *Using Sybase Failover in a High Availability System*.

- Perform a minor upgrade, which does not involve a cluster downtime. In this case, each node is failed over to another node and upgraded one at a time. Assuming ASE1 is the primary companion on node N1, and ASE2 is secondary companion running on node N2:

  a Upgrade the primary companion:

    1 Fail over ASE1 from N1 to N2. You can achieve this by relocating the primary resource group to N2 or by shutting down ASE1.

    2 Upgrade the cluster subsystem on N1 following the upgrade instructions provided by the vendor.

    3 Fail back ASE1 from N2 to N1. See the appropriate cluster chapter in *Using Sybase Failover in a High Availability System* for more details on Adaptive Server failback.

  b Upgrade the secondary companion.

    - If you are using an asymmetric configuration:

      1 Take the secondary resource group offline and verify ASE2 is shut down. ASE2 is not available during this upgrade.

      2 Upgrade the cluster subsystem on N2 following the upgrade instructions from the vendor.

      3 Start ASE2 by bringing the secondary resource group online to N2.

    - If you are using a symmetric configuration, follow the steps in “upgrade the primary companion,” above, for ASE2.

4.3.2 Upgrading Adaptive Server configured with high availability

The following procedure explains how to upgrade an Adaptive Server configured with High Availability. These upgrade steps are only applicable for major upgrades such as 12.5.x to 15.x.x.

❖ **Upgrading a High Availability-enabled Adaptive Server in an active-active configuration:**

1 Drop the high availability companionship.

- Asymmetric configuration – on the secondary server, use isql:

```
sp_companion <primary-server-name>, "drop"
go
```

- Symmetric configuration – run the same command as above on both servers.

  Use isql to verify that both servers are in single-server mode:

```
sp_companion
go
```

2 Use the appropriate command for your cluster system to stop monitoring resources associated with Adaptive Server on each cluster node. You may want to offline or unmanage the resources and resource groups on some cluster systems to prevent unwanted failover during the upgrade.

3 Log in to the server using isql. Disable HA by entering:

```
sp_configure 'enable HA', 0
go
```

4 To complete the change, shut down and restart the Adaptive Server.

5 Upgrade each Adaptive Server Enterprise separately, following the instructions in the appropriate upgrade chapter of the installation guide for your platform.

6 Run the new Adaptive Server Enterprise installmaster script against the newly upgraded Adaptive Servers.

7 Enable the ASE HA property on the new server. Log in to the server using isql and configure the server to enable HA by entering:

```
sp_configure 'enable HA', 1
go
```

8 To complete the change, shut down and restart the Adaptive Server.

9 Run the new Adaptive Server *installhasvss* script against the newly upgraded Adaptive Servers.

10 Follow the instructions in *Using Sybase Failover in a High Availability System* to configure the permission and ownership for *$SYBASE/$SYBASE_ASE/bin/sybha* and *$SYBASE/$SYBASE_ASE/install/sybhauser*.

11 Modify high-availability related files such as the *RUN_server_file*, and the *SYBASE.csh* and *SYBASE.sh* files, if those files are required on the cluster platform.

12 Reconfigure each cluster resource associated with Adaptive Server, depending on platform-specific requirements. For example, on Veritas Cluster, update the *HAase* resource properties, the *RUN_server_file*, and *Sybase_home*.

13 Manually restart Adaptive Server on each cluster node using trace flag 2209. Use the Adaptive Server command line option -T2209.

14 Use the appropriate command for your cluster system to restart monitoring resources associated with Adaptive Server on each cluster node. You may need to online or manage the resources and resource groups if you offlined or unmanaged them in Step 2.

15 Reestablish companionship. See Using Sybase Failover in a High Availability System for information on how to configure companionship.

For an asymmetric configuration on the secondary server, use isql:

```
sp_companion <primary-server-name>, configure
go
```

If user databases exist on the secondary server, you may see one or more warning messages. You can safely ignore these messages, which look similar to:

```
Msg 18739, Level 16, State 1:
Server 'svr2', Procedure 'sp_hacmpcfgvrfy', Line 102:
Database 'svr2_db1': a user database exists. Drop this
database and retry the configuration again.
```

For a symmetric configuration run the sp_companion configure command as above on both servers. Use isql to verify that both servers are in single-server mode:

```
sp_companion
```

```
go
```

Warning! Do not use trace flag 2209 after the Adaptive Server companionship is re-established.

16 Use the appropriate cluster command to take offline, then bring back online, each resource group associated with Adaptive Server. Make sure you remove the -T2209 option from *run_server_file* if added. Onlining and offlining the Adaptive Server resource shuts down the server and restarts it using the *run_server_file*.

Use isql to connect to each Adaptive Server Enterprise and verify the correct server companionship:

```
sp_companion
go
```

In asymmetric mode, the output you see on the primary server is similar to the following:

```
Server 'svr1' is alive and cluster configured.
Server 'svr1' is configured for HA services.
Server 'svr1' is currently in 'Primary normal' mode.
(return status = 0)
```

The output you see on the secondary server is similar to the following:

```
Server 'svr2' is alive and cluster configured.
Server 'svr2' is configured for HA services.
Server 'svr2' is currently in 'Secondary normal'
mode.
(return status = 0)
```

In symmetric mode, the output you see on the primary server is similar to the following:

```
Server 'svr1' is alive and cluster configured.
Server 'svr1' is configured for HA services.
Server 'svr1' is currently in 'Symmetric normal'
mode.
(return status = 0)
```

The output you see on the secondary server is similar to the following:

```
Server 'svr2' is alive and cluster configured.
Server 'svr2' is configured for HA services.
Server 'svr2' is currently in 'Symmetric normal'
mode.
(return status = 0)
```

17 To verify failover and failback, use the cluster command to switch resources associated with Adaptive Server to another node and then switch back.

❖ **Upgrading High Availability-enabled Adaptive Servers in an active-passive configuration**

1 Use the appropriate cluster system command to take Adaptive Server Enterprise offline.

2 Use the appropriate cluster system command to stop monitoring resources associated with Adaptive Server Enterprise.

3 Disable the HA on the server by entering:

```
sp_configure 'enable HA', 0
go
```

4 To complete the change, shut down and restart the Adaptive Server.

5 Follow the steps in the appropriate upgrade chapter of the installation guide for your platform.

6 Run the new Adaptive Server *installmaster* script against the newly upgraded Adaptive Server.

7 If Adaptive Server is configured for active-passive HA on SunCluster3.x, log in to the server using isql and configure the server to enable HA by entering:

```
sp_configure 'enable HA', 2
go
```

If using other cluster systems, make sure ‘enable HA’ is set to 0.

8 To complete the change, shut down and restart the Adaptive Server.

9 Modify the high-availability related files such as the *RUN_server_file*, and the *SYBASE.csh* and *SYBASE.sh* files, if those files are required on the cluster platform you are using.

10 Use the appropriate cluster system command to reconfigure resource properties to reflect the new Adaptive Server installation location.

11 Manually restart Adaptive Server.

12 Use the appropriate cluster system command to restart monitoring resource associated with Adaptive Server.

13 Use the appropriate cluster command to take resources associated with Adaptive Server offline on the same node, then bring them back online.

14 To verify failover and failback, use the cluster command to switch resources associated with Adaptive Server to another node and then switch back.

5. Changes that affect existing applications

Following are changes in Adaptive Server that affect your existing applications.

5.1 Directio support

Adaptive Server version 15.0.1 supports directio on Linux platforms.

5.2 Directory changes

Sybase has made the following directory structure changes:

- Adaptive Server Enterprise components are in *ASE-15_0*, and Open Client and Open Server components are in *OCS-15_0*.
- In Adaptive Server Enterprise version 15.0.1, all product components use JRE-1_4, which is *jre142* in the *$SYBASE/shared* directory.
- The Job Scheduler component (*JS-12_5* in Adaptive Server version 12.5) is now in the *ASE-15_0/jobscheduler* directory.
- The Web Services offering directory has changed from *WS-12_5* to *WS-15_0*.
- The licensing component (SySAM) is now available in the *SYSAM-2_0* directory.

5.3 Changes to ODBC and OLE DB Driver support

Adaptive Server version 15.0 and later contains new ODBC and OLE DB drivers developed by Sybase. The third-party rebranded ODBC and OLE DB Driver Kits included with previous versions are no longer shipped.

The retired ODBC Driver Kit was installed in *%SYBASE%\ODBC*, and registered with the ODBC Driver Manager as "Sybase ASE ODBC Driver." The new ASE ODBC Driver by Sybase is installed in *%SYBASE%\DataAccess\ODBC*, and registered as "Adaptive Server Enterprise." The version shipping with Adaptive Server and SDK 15.0 is version 15.0.0.50.

The retired OLE DB Driver Kit was installed in *%SYBASE%\OLEDB*, and used the provider short name of "Sybase.ASEOLEDBProvider" and the long name of "Sybase ASE OLE DB Provider." The new ASE OLE DB Provider by Sybase is installed in *%SYBASE%\DataAccess\OLEDB*, and uses provider short name "ASEOLEDB." The version shipping with SDK 15.0 is version 15.0.0.51.

Sybase recommends that you update your ODBC Driver and OLE DB Provider to the new Sybase ODBC Driver and OLE DB Provider as soon as possible. Specifically, you should migrate:

- When upgrading to Adaptive Server Enterprise version 15 or later versions of Adaptive Server Enterprise version 12.5.x
- For older versions of Adaptive Server Enterprise, before July 30, 2007

For your convenience during the migration process, the third-party driver and provider support connectivity to the 15.x version of Adaptive Server until November 29, 2007, but will not support new features added after Adaptive Server 12.5.3. Patches released after that date will perform a version check, and will not allow connection to Adaptive Server 15.x. This version-check behavior will only be incorporated in patches released after November 30, 2007. If patches released after November 30, 2007 are not applied, the third-party driver and provider can continue to connect to any versions of Adaptive Server available in the market.

See *New Features for Open Server 15.0 and SDK 15.0 for Windows, Linux and UNIX*, for instructions on migrating to the new drivers.

For more information on migrating to the new OLE DB Provider, see *Open Server 15.0, Open Client 15.0 and SDK 15.0* and *New Features Open Server™ 15.0 and SDK 15.0 for Microsoft Windows, Linux, and UNIX*.

5.4 PC-Client CD

The release of PowerDesigner PhysicalArchitect, included on the PC-Client CD with this release, no longer contains the DataDirect ODBC drivers that were previously bundled with the product. The Software Developer Kit (SDK), also included on the PC-Client CD, includes an ODBC Driver for Adaptive Server Enterprise.

5.5 Changes that affect application behavior

This section discusses changes that affect application behavior.

5.5.1 Long-identifier changes

- Adaptive Server Enterprise now supports long identifiers. There are new limits for the length of object names or identifiers: 255 bytes for regular identifiers, and 253 bytes for delimited identifiers. The new limit applies to most user-defined identifiers including table name, column name, index name and so on. Due to the expanded limits, some system tables (catalogs) and built-in functions have been expanded. For variables, "@" counts as 1 byte, and the allowed name is 254 bytes long.
- Change identifier names with corresponding application changes for binding values. Make sure your application is not binding names of identifiers with only 30 bytes, which was the previous limit. This may cause a variety of unexpected behaviors or error messages.

5.5.2 Query changes

- Some queries may now return general syntax error (message 102) instead of syntax error at line # (message 156).
- The order of result sets in Adaptive Server Enterprise version 15.0.1 differs unless there is an order by clause in the query.
- Query compilation time may increase as the query processing engine looks for more ways to optimize the query.
- See *Query Processor* for details about query processing in Adaptive Server Enterprise version 15.0.1.

5.5.3 Component Integration Services changes

- NULL behavior differs in Oracle, ASA, and IQ—you must override ANSI NULL behavior for it to work as in Adaptive Server Enterprise version 15.0.1.
- Adaptive Server Enterprise version 15.0.1 no longer pads char null, varchar, binary null, and varbinary datatypes for proxy tables.
- By default, cursors are READONLY. Declare cursors with for update to update through them.
- You must have an index if you are using updateable cursors.
- Component Integration Services engines no longer take special measures to re-declare cursors for back ends that close cursors on end tran.
- You must declare an explicit begin tran and end tran around cursor statements for DB2 servers.
- If you attempt to create a column constraint on java, text, image, and unitext datatypes, message 11074 is raised.
- create table, create existing table, or create proxy table statements that contain a location clause are restricted to be the only statement in a batch. This also applies to select into statements that contain a location clause.

5.5.4 Error message changes

- Many messages have been changed to specify “ASE” in the error message.
- When creating a temporary table that already exists, Adaptive Server Enterprise raises message 12822, instead of 2714.
- The identity column overflow message is now raised with message 587 instead of 4916.
- When you create a Java function that does not exist in the catalogs, message 14216 is raised instead of syntax error message 195.
- When a non-owner executes sp_procxmode to change the transaction mode associated with a stored procedure, error message 10354 is raised.
- Arithmetic overflow errors are now raised with message 3606 with severity 16.
- Message 2579 has been replaced with message 12907 in dbcc checktable output.

5.5.5 jConnect version 6.05 ships with Adaptive Server Enterprise

Starting with Adaptive Server Enterprise 15.0 only jConnect™ for JDBC™ version 6.05 is shipped. If you have applications that are dependent on jConnect 5.5, Sybase recommends you either migrate those applications to jConnect 6.0 or use an existing jConnect 5.5 release area.

5.5.6 SQL Remote no longer ships on the PC-Client CD

SQL Remote has been replaced with a more flexible and powerful technology called MobiLink that provides bidirectional synchronization between ASA/UltraLite clients and various back-end databases, including Adaptive Server Enterprise. Download the developer edition of MobiLink at http://www.ianywhere.com/developeredition.

5.5.7 Table changes

- sysindexes has the following new columns:
  - partitiontype
  - conditionid
  - status3
- The following sysindexes columns are now maintained in syspartitions:
  - doampg
  - ioampg
  - first
  - root

  These columns display 0 after you upgrade. The base_partition column is now obsolete and displays 0 after you upgrade.
- The syspartitions table is renamed to sysslices during the upgrade process, then the new table is empty and unused.
- syscomments gains a new column called partitionid.
- systabstats gains the following new columns:
  - partitionid
  - plldegree
  - statmoddate

- sysstatistics gains a new column called partitionid.
- There is a new type of object called a partition condition object, which has a row in sysobjects. A partition condition object is the representation of a tree for the partition table boundary conditions. The tree is stored in sysprocedures.
- sysobjects has the following new columns:
  - identburnmax
  - spacestate
  - erlchgts
- sysstatistics stores the data change counters with formatid=108. The space required by sysstatistics increases due to the additional rows stored.
- Most system catalogs have been converted to the datarows-locking scheme. However, DDLs continue to use the table-level locks. The row locking of the system catalogs can require an increase in the configuration parameter number of locks, depending on the DDLs in the application.
- The system catalogs converted to the datarows-locking scheme do not have a clustered index with indid equal to 1. The clustered indexes now have index ID greater than or equal to 2.

5.6 Changes that affect database administration operations

This section discusses changes that affect database administration operations.

5.6.1 Usage of system built-in functions

Built-in functions that provide space information, such as data_pgs, reserved_pgs, used_pgs, ptn_data_pgs, and rowcnt have been replaced with data_pages, reserved_pages, used_pages, and row_count, respectively. See *Reference Manual: Building Blocks* for detailed information.

5.6.2 DDL and DML changes

- alter table… unpartition is not allowed on tables with indexes (use alter table with 1 partition to remove the partitions).
- alter table on a partitioned table with max parallel degree less than the number of partitions may succeed without raising message 326.

- Creating a clustered index on an empty partitioned table returns a new informational message, 1936.
- null column names are not allowed during view creation.
- A select that contains an aggregate from a proxy table that is mapped to an RPC with a parameter fails with error message 201.

5.6.3 System stored procedure changes

- The output of sp_help *object* has been changed. Specifically, `Data_Located_on_segment` has been removed, while `Computed_Column_Object` and information related to partitions have been added.
- The order of index_keys and index_description in the output of sp_helpindex has been changed.
- Major changes have been made to the output of sp_helpartition and sp_helpsegment.
- sp_who returns “NULL” instead of blank in the hostname column for all system tasks.
- sp_who results in SQL command (insert or select) in the output. In earlier versions, sp_who returned select for the tasks that executed sp_who.
- sp_objectsegment now displays segment information for all partitions of the table.
- sp_lock output has a new column called partitionid, that currently has a value of 0, and is reserved for future use.
- sp_monitorconfig now accepts number of open partitions. When this procedure is executed with the parameter all, additional output for configuration parameter number of open partitions is displayed.
- The Metadata Cache Management section of sp_sysmon displays additional information pertaining to open partitions.

5.6.4 Common diagnostics changes including trace flag usage

- dbcc listoam output has changed substantially.
- The space state message printed in the old dbcc listoam output is now available as a new column spacestate in sysobjects.
- dbcc tablealloc output has changed.

- dbcc page output now displays the print partition ID instead of the object ID.
- dbcc checktable output on a partitioned table has been changed to provide partition-level information.
- dbcc checktable/tablealloc/indexalloc has new syntax for partition support.
- Output for this command has changed:

  REORG RECLAIM_SPACE <*tablename*> with RESUME
- reorg rebuild fails with message 11051 when table is in use by other tasks.
- Output of sp_dbcc_faultreport has been enhanced to show partition ID information.
- The hostname, program_name, hostprocess, and cmd columns in sysprocesses have been changed to varchar(30) and made nullable. select from sysprocesses returns NULL for these columns instead of space for system tasks.
- The “first” column in sysindexes has been moved to the “firstpage” column in syspartitions.
- The name column in syspartitions has been changed from index_name+_+tableid to index_name+_+ptn_id for partition table with indexes.
- The basic cost of the optimizer now includes CPU cost.
- Parallel costing is now performed only on base tables and indexes that are larger than 20 pages.
- For information about trace flags and diagnostics, see the *Query Processor* document.

5.6.5 System resource changes

- Adaptive Server Enterprise version 15.0.1 uses more procedure cache for several reasons, including:
  - The query processing engine now looks for additional ways to optimize the query.
  - The execution engine avoids materialization of worktables and evaluates aggregations in memory as much as possible.

    - The data change counters maintained for the datachange() function allocate memory from procedure cache. The partition condition tree is cached in the procedure cache. Partition boundary values are allocated in the procedure cache, resulting in a required increase in procedure cache resources.

- A configuration parameter, max repartition degree, has been added. This parameter controls the maximum degree to which an intermediate data stream can be repartitioned. The default value of this parameter is 1, which indicates that repartitioning is not set and is bound by the number of online engines configured for Adaptive Server Enterprise. For a query with a large number of tables, Adaptive Server Enterprise version 15.0.1 can put an increased demand on auxiliary scan descriptors.

  Set the value of max repartition degree to a value lower than the number of engines to decrease resource usage. You may also need to configure a larger value for the auxiliary scan descriptor pool.

- Adaptive Server Enterprise version 15.0.1 avoids worktable materialization and incurs more resources in auxiliary scan descriptors.

- During the upgrade process, max memory is increased by Adaptive Server Enterprise if the new total logical memory is greater than max memory. The new value of max memory is set to the new value of total logical memory.

- During the upgrade process, number of open partitions is set to the same value as number of open indexes, resulting in increased memory usage. An open partition requires approximately 950 bytes.

5.6.6 Unpartitioning of user tables during upgrade

In version 15.0, each table partition must have a different partitionid. To avoid the expense of changing the partitionid for each page during the upgrade process, Adaptive Server Enterprise unpartitions any partitioned tables. Repartition these tables when the upgrade process completes.

5.6.7 Space management changes

- System databases have increased in size by 2 allocation units. The master database has increased in size by 14 allocation units. The default user database size has increased by 2 allocation units.

- The upgrade process requires free space to convert system catalogs to the datarows-locking scheme. The typical requirement for free space is 125 percent of the largest system catalog. If you are upgrading from a 32-bit to 64-bit binary, additional space of 55 percent of sysprocedures is required.

5.6.8 Changes to transaction dump content

You can load a transaction dump from Adaptive Server Enterprise version 11.9.2, 12.0x, or 12.5x over a database dump or transaction dump that has preceded it in the load sequence.

However, if that transaction dump contains a create index transaction, then load transaction recovery treats that transaction's sort record as the end of the log. Adaptive Server Enterprise does not process any more log records in that transaction, and does not allow load transaction dumps in the load sequence. This message prints:

```
The pre-15.0 log for database '%.*s' includes a create
index incompatible with this server version and
operating environment. Only transactions completed
before create index will be recovered.
```

This behavior is much like the point-in-time recovery that is affected by the load transaction until_time option.

In contrast, Adaptive Server Enterprise should successfully load any database dump from one of the earlier versions, regardless of its content.

There is no content limitation in load database or load transaction when using a dump from a version 15.0.1 Adaptive Server Enterprise.

5.7 Open Client and Open Server changes

For information about changes that affect Open Client and Open Server, see the most recent release bulletins for these products at the Sybase Web site at http://sybooks.sybase.com.

6. Known problems

This section describes known problems and workarounds for Adaptive Server Enterprise.

Note Sybase no longer includes System Problem Reports (SPR) and Closed Problem Reports (CPR) with your copy of Adaptive Server Enterprise. Known problems, which used to be in the SPR, are now described in this section of the release bulletin. Closed issues that were documented in the CPR are now searchable from Solved Cases at the Sybase Web site at http://www.sybase.com. To access the Solved Cases search form, click Support and Services | Solved Cases.

6.1 Binding tempdb to a named cache

[CR #438815] In rare cases, when you bind tempdb to a named cache, and when Adaptive Server is using a HASH operator to execute a query, a stacktrace is reported in the error log. The error message looks similar to `current process infected with 11` and may include the following modules:

- `LeHashContext::HashVtuple`
- `LeHashContext::GetNext`
- `LeHashContext::ProcessTuple`
- `LeAlignedDataRow::readRow`
- `LeHashBufferReader::GetNext`

Workaround: Should you encounter this error, do not bind tempdb to a named cache.

6.2 Using *sp_setreptable* for replication

[CR #439024] When a table with a JAVA ADT computed column is marked for replication using sp_setreptable, the JAVA ADT columns are internally set to no-computed column, which causes DML on this column to be return with inaccurate results.

Workaround: Use sp_setreplicate and sp_setrepcol instead of sp_setreptable to mark a table with a computed column for replication. To do so:

1 Set the table replicated using sp_setreplicate:

```
sp_setreplicate object_name, {true | false}
```

2 Set the text, image, or off-row object column replicated using sp_setrepcol, by entering:

```
sp_setrepcol table_name [, column_name
[, {do_not_replicate | always_replicate |
replicate_if_changed} [, use_index ] ] ]
```

6.3 Creating Job Scheduler schedules for Adaptive Server with Simplified Chinese

[CR #439404] When you attempt to create a Job Scheduler schedule using the Adaptive Server plug-in for an Adaptive Server localized for Simplified Chinese, you occasionally receive an error indicating that the varchar datatype can not be converted to a date.

Workaround: Localize Adaptive Server for Simplified Chinese before configuring Job Scheduler.

6.4 Configuring Job Scheduler Self Management for Simplified Chinese

[CR #439406] When you configure the Job Scheduler Self Management feature for an Adaptive Server localized for Simplified Chinese, you occasionally receive an error saying that a varchar datatype can not be converted to a date. The Job Scheduler configuration process stops at this point.

Workaround: Do not configure the Self Management feature for Simplified Chinese.

6.5 Installing an ESD overlay with Unified Agent

[CR #427275] ESD overlay installations do not replace the Unified Agent's *ARL.xml* file. As a result, any developer updates or changes to *ARL.xml* are not available.

Workaround: Rename the old *ARL.xml* before you run the installer.

6.6 Formatting the options file during a silent installation

[CR #437646] When you run the Installer to generate an options file for the silent installation, a series of formatting error messages display.

Workaround: You can safely ignore these error messages. An accurate options file has been generated.

6.7 Using SYBASE_LICENSE_FILE

[CR #438342] On Linux AMD 64-bit platforms, if you use the SYBASE_LICENSE_FILE environment variable during start-up, Adaptive Server fails to start.

Workaround: Copy the licenses or create pointer licneses in the default licnese directory.

6.8 Installing correct fonts when displaying Chinese characters

[CR #434986, 440097] If the fonts installed on you target machine did not match the default font configuration in *font.properties*,the InstallShield GUI interface and Sybase Central do not display Chinese characters properly.

Workaround: Set the *font.properties* file manually for the precise language and platform you are using. To set the *font.properties* file manually:

1 Check whether the locale is set correctly into Chinese (using the locale command to see the correct locale information). If not, set locale to Chinese.

2 Check whether the Chinese fonts are installed correctly. If not, reinstall them.

3 Select one available Chinese font you want to use to display Chinese characters. Remember the name and the absolute path of the font's file.

4 Go to the directory where you execute the setup program to install Adaptive Server. Go into the subdirectory *JVM\lib*.

Look for the *font.properties.zh_CN.<platform>* file for you platform. For example, if you platform is SuSE 9.0, look for *font.properties.zh_CN.SuSE* or *font.properties.zh_CN.SuSE9.0*. If the file does not exist, create it.

5 Open *font.properties.zh_CN.<platform>*, and configure the file content correctly. the following two links should be shown to customers. For information about the format of the font.properties file, see the Sun Web site at http://java.sun.com/j2se/1.4.2/docs/guide/intl/fontprop.html. For information about how to configure the font.properties files, see the Sun Web site at http://java.sun.com/j2se/corejava/intl/reference/faqs/index.html

6.9 Migration Tool GUI truncation issues with Japanese characters

[CR #436783] Some screens in the Migration Tool GUI are truncated in Adaptive Server version 15.0.1 (Japanese). The following list identifies the GUI screen and component that has been truncated and provides the intended label.

- In the Login Account Options dialog, the password file label and the Cancel button are truncated.
- In the Setup Progress dialog, the label to the left of the Task Status field should display “Current Task field.”
- In the Thread information dialog, the Combo Box label should display “Migration Path,” the second Edit Box label should display “Index Thread.” The truncated button should display “Cancel.”
- In the Report Path dialog, the truncated button should display “Run Report.”

Workaround: There is no workaround for this issue.

6.10 Highlighted known installation issues

This section contains information about known installation problems and their workarounds.

Warning! If you have multiple Sybase products installed in the same *$SYBASE* directory, uninstalling one of them *may* make the other products inoperable. If this happens, you may need to reinstall those remaining products.

6.10.1 ClearCase causes InstallShield to stop responding

InstallShield may stop responding, and display this message:

```
Please wait while the installer checks your system
information.
```

There are two known causes:

[CR #404068] The JRE used by InstallShield is in conflict with ClearCase.

Workaround: Turn off ClearCase before installing Adaptive Server Enterprise. See your ClearCase administrator for details. You can turn ClearCase back on after installation.

[CR #402530] NFS-mounted device is inaccessible because the remote server is not responding. Verify this by executing `df -P` from a command prompt. You may instead see a message similar to: `NFS3 server not responding still trying`.

Workaround: Unmount the inaccessible device and retry the installation. You may need to restart the server.

Note You cannot use these workarounds on a 64-bit hardware system. There are no workarounds for the 64-bit systems.

6.10.2 EAServer manager and Security manager

[CR #400839] The Adaptive Server Enterprise Sybase Central 15.0.1 install overwrites the EAServer Manager and the Security Manager. When you cannot access the EAServer Manager you cannot access Jaguar Manager.

Workaround: Copy the information about EAServer Manager and Security Manager from file *C:\Sybase\WorkSpace\DevRuntimes\Shared\Sybase Central 4.3\.scRepository* to the end of file *C:\sybase\Shared\Sybase Central 4.3\.scRepository,* save it. You must restart EAServer Manager.

6.11 Cross-platform bulk copying of parsed XML images

[CR # 332012] You cannot use bcp or replicate parsed XML image data between a big-endian platform and a little-endian platform, for example between a Solaris platform (big-endian) and a Linux or Windows platform (little-endian).

The term “parsed XML image data” refers to data of datatype image that is generated by the xmlparse function.

Workaround: When you transfer XML data between platforms, transfer the character form of the XML documents rather than the parsed XML form. If you have not stored the character form of the XML documents, you can regenerate it from the parsed form. For example, if column xmlindexed of the xmldocuments table is an image column containing parsed XML image data, the following two commands append a new text column to the table and populate it with a character form of the documents contained in the xmlindexed column:

```
alter table xmldocuments add xmltext text
update xmldocuments set xmltext = xmlextract('/', xmlindexed)
```

You can then use bcp or replicate the xmltext column from one platform to the other.

6.12 Running on RH3.0 64-bit distribution occasionally stops responding

Due to a Red Hat 3.0 update 3 anomaly, Adaptive Server Enterprise occasionally stops responding when kernel asynchronous I/O is enabled. When this happens, Adaptive Server Enterprise generally displays a high CPU load—approximately 99%—without I/O. Adaptive Server Enterprise can still accept connections, but cannot recover; you must restart the system to clear the process.

This problem has been reported to Red Hat and addressed in kernel 2.4.21-27.6.EL.

For more information, see the Red Hat Web site at https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=138905.

6.13 Interactive SQL and Adaptive Server Enterprise plug-in issues

This section discusses Interactive SQL and Adaptive Server Enterprise plug-in issues and their workarounds.

6.13.1 Opening XML files with Interactive SQL

[CR #400825] Interactive SQL displays a stack trace when you open an XML file.

Workaround: There is no workaround for this issue.

6.13.2 Connecting Adaptive Server Enterprise plug-in to RepConnector

[CR #400709] The Adaptive Server Enterprise plug-in cannot connect to Adaptive Server Enterprise if RepConnector™ 2.5 and the Adaptive Server Enterprise plug-in 15.0.1 are installed in the same *$SYBASE* location.

Workaround: Install RepConnector and the Adaptive Server Enterprise plug-in into different *$SYBASE* directories.

6.13.3 Query support in Interactive SQL

[CR #398435] Interactive SQL does not support queries with a compute clause.

Workaround: Use isql from the command line instead.

6.14 Statistics in system tables

[CR #399624] When you create or configure a server in, or upgrade a server to, Adaptive Server Enterprise 15.0.1, index and table-level statistics are incorrect.

Workaround: Execute update index statistics on the following system tables:

- sysobjects
- sysindexes
- syscolumns
- systypes
- syslogins
- sysusers

6.15 Monitor GUI does not start when the LDAP server is used

[CR #400784] When you start the Monitor Client GUI in an environment using LDAP instead of an *interfaces* file, the Monitor Client GUI cannot connect to any servers.

Workaround: Create an *interfaces* file with entries for the servers that are monitored and use this with the Monitor Client GUI.

6.16 *allow backward scan* does not work

[CR #401543] The configuration parameter allow backward scan does not work in Adaptive Server.

Workaround: If you are encountering increased deadlocks while executing backward scans, separate problematic scans:

1 Select the required rows into a temporary table in ascending order.

2 Select from the temporary table in descending order.

6.17 *set statistics io* does not display I/O generated by worker processes

[CR# 401560] After executing a parallel query, displays only show the alpha thread scan, and logical and physical I/O counts. set statistics io does not display I/O counts generated by worker processes.

Workaround: Use set statistics plancost on to display the I/O counts from alpha thread scan and any child threads in a graphical format.

6.18 Replication Server compatibility issues

[CR # 382874] rs_init in Replication Server version 12.6 and earlier cannot create an RSSD database in Adaptive Server Enterprise version 15.0.1.

Workaround: Create the RSSD database manually before using rs_init to create a new Replication Server if you do not use an Embedded RSSD. See the *Replication Server Administration Guide* for additional information.

6.19 Blank spaces in an *ldap* server entry

[CR #333307] If you have a blank space after your ldap server entry, dscp defaults to using the *interfaces* driver and does not connect to an ldap server.

Workaround: To use dscp to make an entry into an ldap server, enable ldap by editing the *$SYBASE/OCS-15_0/config/libtcl.cfg* file to add the ldap server you plan to use.

6.20 Dumping or loading databases with asynchronous I/O

[CR #335852] On an IA32 running Red Hat, running a dump or load database command can cause Backup Server to stop responding when using asynchronous I/O. Backup Server uses asynchronous I/O by default.

Workaround: Start Backup Server using trace flag -D32 to force a synchronous I/O.

6.21 *cis connect timeout* and *enable SNMP* are not implemented

[CR #323177] The cis connect timeout and enable SNMP configuration parameters are not currently implemented.

Workaround: Do not use these configuration parameters.

6.22 Handling multibyte character sets during migration

[CR #353079] If you configure multibyte character sets after migrating data from system catalogs sysattributes and sysxtypes, the text columns in these catalogs are inconsistent with the multibyte character sets.

Workaround: Manually run dbcc fix_text on sysattributes and sysxtypes to make the text columns consistent with the multibyte character sets.

6.23 Sybase Central

[CR #428292] Using the Adaptive Server Plug-in with an older Sybase Central release results in "SybPasswordField not found" stack trace.

Workaround: Update the existing Sybase Central version, or install Sybase Central in a new location.

6.24 File names

[CR #416569] When updating the filename column of a proxy table mapped to a directory of files, Sybase recommends that you qualify the update to specific files by having the filename included in the where clause of the update. Since some file systems may implement a rename as a deletion followed by the creation of a new directory entry, the same filename could be updated multiple times if the update were not restricted.

For example, this statement could cause multiple updates:

```
update t1 set filename=filename + 'old' where filetype
= 'REG'
```

Workaround: Add a clause such as `"and filename like "%.c"`.

6.25 *sp_configure current audit table*

[CR #432345] sp_configure 'current audit table' <*number*> does not generate an error if the audit table is not empty and the previous and current audit table numbers are the same. You see an informational message:

```
18932, "Resulting configuration value and memory use
have not changed from previous values: new configuration
value %1!, previous %2!."
```

7. Product compatibilities

This section lists the Sybase components that are compatible with Adaptive Server Enterprise version 15.0.1. For information about operating system requirements, see individual component documentation.

Note The DirectConnect for Informix, DirectConnect for Microsoft SQL Server, and DirectConnect for DB2 UDB do not support LDAP.

Note Component Integration Services and SSL are not supported on Adaptive Server Enterprise version 15.0 for Linux.

The following components have been tested for compatibility with Adaptive Server Enterprise version 15.0.1:

- Sybase Character Sets 3.0
- Sybase Central viewer 4.3
- ECDA option for ODBC 12.6
- ECDA option for Informix 12.6
- MainframeConnect DirectConnect for z/OS 12.6
- ECDA option for Oracle 12.6
- ECDA option for Microsoft SQL Server 12.6
- ECDA option for DB2 Universal Database 12.6
- OpenSwitch 12.5, OpenSwitch 12.5.1, ESD #2

  Note OpenSwitch 12.5 does not support SSL.

- jConnect for JDBC 5.5
- jConnect for JDBC 6.05
- Open Server™ 12.5.1
- Open ServerConnect™ (CICS, IMS/MVS) 4.0
- Open ClientConnect™ (CICS, IMS/MVS) 4.0
- InfoMaker™ 7.0.3
- Replication Server® 12.6

- PhysicalArchitect™ 8.0
- XA-Library™ for CICS/ENCINA/Tuxedo 12.5.1
- PowerTransfer
- ADO.NET 1.15.x Windows x86 32-bit
- ADO.NET 1.1x Windows x86 32-bit
- ASE ODBC Driver by Sybase 12.5.x Window x86 32-bit
- ASE ODBC Driver by Sybase 15.0.x Window x86 32bit and x64 64-bit
- ASE OLE DB Provider by Sybase 12.5.x Windows x86 32-bit
- ASE OLE DB Provider by Sybase 15.x Windows x86 32-bit
- ASE ODBC Driver by Sybase Linux x86 32-bit

7.1 Adaptive Server interoperability

The following tables show the interoperability of Adaptive Server version 15.0.1 against other Sybase products across different platforms and versions, as well as client products. For specific platform or operating system level information, perform a search at the Sybase Certification Web page at http://certification.sybase.com/ucr/search.do for the appropriate platform certification.

Note 1 Even though two or more products might be interoperable, features introduced in a newer version of a product are not likely to be supported with older versions of the same or other products.

Note 2 Interoperability between big-endian and little-endian platforms has also been verified. Windows and Linux-32 platforms are little-endian. IBM AIX, Sun Solaris, Linux on Power, and HP-UX are big-endian platforms.

Table 1: Adaptive Server version 15.0.x interoperability with other products

Platforms		OC/OS 12.0x	OC/OS 12.5X	OC/OS 15.0X	jConnect 5.5	jConnect 6.0.x	ASE ODBC by Sybase 12.5x	ASE ODBC by Sybase 15.0x	ASE OLEDB by Sybase 12.5x	ASE OLEDB by Sybase 15.0x	ADO.NET 1.1x	OEM ODBC Driver Kit 5.00.0096	OEM OLEDB Driver Kit 2.70.0063	Replication Server 12.6	Replication Server 15.0.x
AIX	32-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
HP-UX	32-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	IA 64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
Linux	x86 32-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	Pseries 64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	x64 64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	IA 64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
Solaris	32-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
Windows	x86 32-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	x86 64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X

+ Available until November 2007. See “Changes to ODBC and OLE DB Driver support” on page 16 for details.

Table 2 shows the supported platforms of client products listed in Table 1.

Table 2: Supported client platforms

Client name	Version	Supported platform	
Open Client/Open Server	12.0.x	• AIX 32-bit • HP-UX 32-bit • Solaris 32-bit	• TruUNIX (Alpha) 64-bit • Windows x86 32-bit
	12.5.x	• AIX 32-bit • AIX 64-bit • HP-UX 32-bit • HP-UX 64-bit • HP-UX IA 32-bit • HP-UX IA 64-bit • Linux x86 32-bit • Linux Pseries 32-bit • Linux Pseries 64-bit • Linux x64 32-bit • Linux x64 64-bit • Linux IA 64-bit	• Mac OSX 64-bit • SGI 32-bit • SGI 64-bit • Solaris 32-bit • Solaris 64-bit • Solaris x86 32-bit • Solaris Opteron 64bit • TruUnix (Alpha) 64-bit • Windows x86 32-bit
	15.0.x	• AIX 32-bit • AIX 64-bit • HP-UX 32-bit • HP-UX 64-bit • HP-UX IA 32-bit • HP-UX IA 64-bit • Linux x86 32-bit • Linux x64 32-bit • Linux x64 64-bit • Linux Pseries 32-bit • Linux Pseries 64-bit	• Solaris 32-bit • Solaris 64-bit • Solaris Opteron 32-bit • Solaris Opteron 64-bit • Windows x86 32-bit • Windows x64 64-bit
jConnect	5.5	• All	
	6.0x	• All	
ASE ODBC by Sybase	12.5.x	• Linux x86 32-bit • Linux x64 32-bit • Mac OSX 32-bit	• Windows x86 32-bit • Windows x64 64-bit
	15.0.x	• Linux x86 32-bit • Linux x64 32-bit	• Windows x86 32-bit • Windows x64 32-bit • Windows x64 64-bit

Client name	Version	Supported platform	
ASE OLEDB by Sybase	12.5.x	• Windows x86 32-bit	• Windows x64 32-bit
	15.x	• Windows x86 32-bit	• Windows x64 32-bit
ADO.NET	1.1.x	• Windows x86 32-bit	• Windows x64 32-bit
OEM ODBC Driver Kit	5.00.0096	• Windows x86 32-bit	
OEM OLEDB Driver Kit	2.70.0063	• Windows x86 32-bit	
Replication Server	12.6	• AIX 32-bit • HPUX 32-bit • HPUX IA 32-bit • Linux x86 32-bit • Linux x64 32-bit	• Mac OSX 32-bit • SGI 32-bit • Solaris 32-bit • TruUnix (Alpha) 64-bit • Windows x86 32-bit
	15.0.x	• AIX 32-bit • HPUX 32-bit • HPUX IA 32-bit	• Linux x86 32-bit • Linux x64 32-bit • Solaris 32-bit • Windows x86 32-bit

7.2 Adaptive Server Replicator 15.0 support of Adaptive Server Enterprise 12.5 datatypes

Adaptive Server Replicator 15.0 does not support large identifiers, or these datatypes that are introduced in Adaptive Server Enterprise version 15.0: bigint, unsigned bigint, unsigned int, unsigned smallint, unitext, computed columns, XML, and encrypted columns.

7.3 Known compatibility issues

This section discusses known compatibility issues.

7.3.1 Installing Enterprise Connect Data Access (ECDA) or MainframeConnect DirectConnect for z/OS with other Sybase software

Sybase strongly recommends you install the ECDA DirectConnect option or MainframeConnect DirectConnect for z/OS, including DirectConnect Manager, into its own product directory.

7.3.2 jConnect 6.0 and *sqldbgr*

Due to missing components in jConnect 6.0, sqldbgr does not run correctly if you select a "Typical" installation. If you plan on using sqldbgr, select the "Full" installation option to install jConnect6.0.

7.3.3 Installing Replication Server with other Sybase software

Sybase strongly recommends that you install Replication Server, including Replication Server Manager (RSM), in its own product directory, separate from other Sybase products.

7.3.4 Using Job Scheduler templates

Sybase recommends that you upgrade to the Job Scheduler templates provided with Adaptive Server Enterprise 15.0. The templates contain important changes in support of new functionality, and use new functions reserved_pages and data_pages. For information on upgrading the templates, see Chapter 5, "Upgrading Adaptive Server," in the installation guide for your platform.

Note If you have existing jobs created from templates delivered with a pre-12.5.2 Adaptive Server Enterprise, see the 12.5.2 release bulletin for your platform and follow the upgrade steps there for the Job Scheduler templates. Then you can upgrade Job Scheduler templates as explained above.

7.3.5 Copy menu

Job scheduler context menus are nonstandard. The options on the Copy menu do not copy to the clipboard, but create new objects in sybmgmtdb.

Edit | Duplicate As is identical to the Copy command.

8. Documentation updates and clarifications

This section describes updates and clarifications for Adaptive Server Enterprise documentation.

8.1 New Features Guide for 12.5.4

The *New Features Guide for Adaptive Server Release 12.5.4* states that Linux Itanium 64-bit supports the Native XML feature. This is incorrect: the Adaptive Server 12.5.4 release for Linux Itanium 64-bit does not support the Native XML feature.

8.2 Transact SQL User's Guide

Add the following to the "Overview" section of the "Partitioning Tables and Indexes chapter of the *Transact SQL User's Guide*.

8.2.1 Range-partitioned tables and sort order change

If you change the sort order, Adaptive Server marks range-partitioned tables with character-based partition keys as "suspect" in sysindexes. These restrictions apply when you run commands on tables that are marked suspect:

- You cannot run DML commands or use cursors with partitions marked "suspect."
- You can run commands like select on tables marked "suspect," but the query processor treats these tables as round-robin partitioned tables, so optimizations using partition conditions such as "partition pruning" do not apply.
- The only alter table command you can run on partitions marked suspect is alter table...repartition.
- You cannot create index on | from or drop index on | from on tables marked with suspect partitions.
- You can fix tables with suspect partitions using:
  - alter table...repartition – if you need to change the partition condition after a sort order change, or
  - reorg rebuild – if you know the partition condition is correct.

  When a table is marked as having suspect partitions and suspect indexes, you can run alter table...repartition or reorg rebuild to fix the partition and the index.

8.3 *Reference Manual: Building Blocks*

This section discusses updates and clarifications to *Reference Manual: Building Blocks*.

8.3.1 Update in *derived_stat*

The information for the derived_stat built-in function in *Reference Manual: Building Blocks* is incomplete. Replace the entire reference section for derived_stat with the following:

Description

Returns derived statistics for the specified object and index.

Syntax

derived_stat("*object_name*" | *object_id*,
index_name | *index_id*,
["*partition_name*" | *partition_id*,]
"*statistic*")

Parameters

- *object_name* – is the name of the object you are interested in. If you do not specify a fully qualified object name, derived_stat searches the current database.
- *object_id* – is an alternative to *object_name*, and is the object ID of the object you are interested in. *object_id* must be in the current database
- *index_name* – is the name of the index, belonging to the specified object that you are interested in.
- *index_id* – is an alternative to *index_name*, and is the index ID of the specified object that you are interested in.
- *partition_name* – is the name of the partition, belonging to the specific partition that you are interested in. *partition_name* is optional. When you use *partition_name* or *partition_id*, Adaptive Server returns statistics for the target partition, instead of for the entire object.
- *partition_id* – is an alternative to *partition_name*, and is the partition ID of the specified object that you are interested in. *partition_id* is optional.
- "*statistic*" – the derived statistic to be returned. Available statistics are:

Value	Returns
data page cluster ratio or dpcr	The data page cluster ratio for the object/index pair
index page cluster ratio or ipcr	The index page cluster ratio for the object/index pair
data row cluster ratio or drcr	The data row cluster ratio for the object/index pair
large io efficiency or lgio	The large I/O efficiency for the object/index pair
space utilization or sput	The space utilization for the object/index pair

Examples

Example 1 Selects the space utilization for the titleidind index of the titles table:

```
select derived_stat("titles", "titleidind", "space utilization")
```

Example 2 Selects the data page cluster ratio for index ID 2 of the titles table. Note that you can use either `"dpcr"` or `"data page cluster ratio"`:

```
select derived_stat("titles", 2, "dpcr")
```

Example 3 Statistics are reported for the entire object, as neither the partition ID nor name is not specified:

```
1> select derived_stat(object_id("t1"), 2, "drcr")
2> go

----------------------------
                      0.576923
```

Example 4 Reports the statistic for the partition tl_928003396:

```
1> select derived_stat(object_id("t1"), 0, "t1_928003306", "drcr")
2> go

 ----------------------------
                    1.000000

(1 row affected)
```

Example 5 Selects derived statistics for all indexes of a given table, using data from syspartitions:

```
select convert(varchar(30), name) as name, indid,
    convert(decimal(5, 3), derived_stat(id, indid, 'sput')) as 'sput',
    convert(decimal(5, 3), derived_stat(id, indid, 'dpcr')) as 'dpcr',
    convert(decimal(5, 3), derived_stat(id, indid, 'drcr')) as 'drcr',
    convert(decimal(5, 3), derived_stat(id, indid, 'lgio')) as 'lgio'
from syspartitions where id = object_id('titles')
go

 name                          indid   sput     dpcr     drcr     lgio
 ------------------------------ ------ -------- -------- -------- --------
 titleidind_2133579608              1    0.895    1.000    1.000    1.000
 titleind_2133579608                2    0.000    1.000    0.688    1.000

(2 rows affected)
```

Example 6 Selects derived statistics for all indexes and partitions of a partitioned table. Here, mymsgs_rr4 is a roundrobin partitioned table that is created with a global index and a local index.

```
1> select * into mymsgs_rr4 partition by roundrobin 4 lock datarows
```

```
2> from master..sysmessages
2> go

(7597 rows affected)

1> create clustered index mymsgs_rr4_clustind on mymsgs_rr4(error, severity)
2> go
1> create index mymsgs_rr4_ncind1 on mymsgs_rr4(severity)
2> go
1> create index mymsgs_rr4_ncind2 on mymsgs_rr4(langid, dlevel) local index
2> go

2> update statistics mymsgs_rr4
1>

2> select convert(varchar(10), object_name(id)) as name,
3>      (select convert(varchar(20), i.name) from sysindexes i
4>       where i.id = p.id and i.indid = p.indid),
5> convert(varchar(30), name) as ptnname, indid,
6> convert(decimal(5, 3), derived_stat(id, indid, partitionid, 'sput')) as
'sput',
7> convert(decimal(5, 3), derived_stat(id, indid, partitionid, 'dpcr')) as
'dpcr',
8> convert(decimal(5, 3), derived_stat(id, indid, partitionid, 'drcr')) as
'drcr',
9> convert(decimal(5, 3), derived_stat(id, indid, partitionid, 'lgio')) as
'lgio'
10> from syspartitions p
11> where id = object_id('mymsgs_rr4')

name                          ptnname                      indid sput dpcr drcr
lgio
----------------------------- ---------------------------- ----- ---- ---- ----
----
mymsgs_rr4 mymsgs_rr4           mymsgs_rr4_786098810           0 0.90 1.000 1.00
1.000
mymsgs_rr4 mymsgs_rr4           mymsgs_rr4_802098867           0 0.90 1.000 1.00
1.000
mymsgs_rr4 mymsgs_rr4           mymsgs_rr4_818098924           0 0.89 1.000 1.00
1.000
mymsgs_rr4 mymsgs_rr4           mymsgs_rr4_834098981           0 0.90 1.000 1.00
1.000
mymsgs_rr4 mymsgs_rr4_clustind mymsgs_rr4_clustind_850099038 2 0.83 0.995 1.00
1.000
mymsgs_rr4 mymsgs_rr4_ncind1   mymsgs_rr4_ncind1_882099152   3 0.99 0.445 0.88
1.000
mymsgs_rr4 mymsgs_rr4_ncind2   mymsgs_rr4_ncind2_898099209   4 0.15 1.000 1.00
```

```
1.000
mymsgs_rr4 mymsgs_rr4_ncind2   mymsgs_rr4_ncind2_914099266   4 0.88 1.000 1.00
1.000
mymsgs_rr4 mymsgs_rr4_ncind2   mymsgs_rr4_ncind2_930099323   4  0.877 1.000
1.000    1.000
mymsgs_rr4 mymsgs_rr4_ncind2   mymsgs_rr4_ncind2_946099380   4  0.945 0.993
1.000 1.000
```

Example 7 Select derived statistics for all allpages-locked tables in the current database:

```
2> select convert(varchar(10), object_name(id)) as name
3>     (select convert(varchar(20), i.name) from sysindexes i
4>      where i.id = p.id and i.indid = p.indid),
5> convert(varchar(30), name) as ptnname, indid,
6> convert(decimal(5, 3), derived_stat(id, indid, partitionid, 'sput')) as
'sput',
7> convert(decimal(5, 3), derived_stat(id, indid, partitionid, 'dpcr')) as
'dpcr',
8> convert(decimal(5, 3), derived_stat(id, indid, partitionid, 'drcr')) as
'drcr',
9> convert(decimal(5, 3), derived_stat(id, indid, partitionid, 'lgio')) as
'lgio'
10> from syspartitions p
11> where lockscheme(id) = "allpages"
12>   and (select o.type from sysobjects o where o.id = p.id) = 'U'

 name                          ptnname                       indid sput dpcr drcr
lgio
----------- ----------------- ---------------------------- ----- ---- ---- ----
----
 stores     stores            stores_18096074               0 0.276 1.000 1.000
1.000
 discounts  discounts         discounts_50096188            0 0.075 1.000 1.000
1.000
 au_pix     au_pix            au_pix_82096302               0 0.000 1.000 1.000
1.000
 au_pix     tau_pix           tau_pix_82096302              255 NULL NULL NULL
NULL
 blurbs     blurbs            blurbs_114096416              0 0.055 1.000 1.000
1.000
 blurbs     tblurbs           tblurbs_114096416             255 NULL NULL NULL
NULL
 t1apl      t1apl             t1apl_1497053338              0 0.095 1.000 1.000
1.000
 t1apl      t1apl             t1apl_1513053395              0 0.082 1.000 1.000
```

```
1.000
 t1apl      t1apl             t1apl_1529053452             0 0.095 1.000 1.000
1.000
 t1apl      t1apl_ncind       t1apl_ncind_1545053509       2 0.149 0.000 1.000
1.000
 t1apl      t1apl_ncind_local t1apl_ncind_local_1561053566 3 0.066 0.000 1.000
1.000
 t1apl      t1apl_ncind_local t1apl_ncind_local_1577053623 3 0.057 0.000 1.000
1.000
 t1apl      t1apl_ncind_local t1apl_ncind_local_1593053680 3 0.066 0.000 1.000
1.000
 authors    auidind           auidind_1941578924           1 0.966 0.000 1.000
1.000
 authors    aunmind           aunmind_1941578924           2 0.303 0.000 1.000
1.000
 publishers pubind            pubind_1973579038            1 0.059 0.000 1.000
1.000
 roysched   roysched          roysched_2005579152          0 0.324 1.000 1.000
1.000
 roysched   titleidind        titleidind_2005579152        2 0.777 1.000 0.941
1.000
 sales      salesind          salesind_2037579266          1 0.444 0.000 1.000
1.000
 salesdetai salesdetail       salesdetail_2069579380       0 0.614 1.000 1.000
1.000
 salesdetai titleidind        titleidind_2069579380        2 0.518 1.000 0.752
1.000
 salesdetai salesdetailind    salesdetailind_2069579380    3 0.794 1.000 0.726
1.000
 titleautho taind             taind_2101579494             1 0.397 0.000 1.000
1.000
 titleautho auidind           auidind_2101579494           2 0.285 0.000 1.000
1.000
 titleautho titleidind        titleidind_2101579494        3 0.223 0.000 1.000
1.000
 titles     titleidind        titleidind_2133579608        1 0.895 1.000 1.000
1.000
 titles     titleind          titleind_2133579608          2 0.402 1.000 0.688
1.000

(27 rows affected)
```

Usage

- derived_stat returns a double precision value.
- The values returned by derived_stat match the values presented by the optdiag utility.

- If the specified object or index does not exist, derived_stat returns NULL.
- Specifying an invalid statistic type results in an error message.
- Using the optional *partition_name* or *partition_id* reports the requested statistic for the target partition; otherwise, derived_stat reports the statistic for the entire object.
- If you provide:
  - Four arguments – derived_stat uses the third argument as the partition, and returns derived statistics on the fourth argument.
  - Three arguments – derived_stat assumes you did not specify a partition, and returns derived statistic specified by the third argument.

Standards: ANSI SQL – Compliance level: Transact-SQL extension.

Permissions: Only the table owner can execute derived_stat.

See also: **Document** *Performance and Tuning Guide* for:

- "Access Methods and Query Costing for Single Tables"
- "Statistics Tables and Displaying Statistics with optdiag"

Utility optdiag

8.3.2 Correction in *xa_bqual* example

Example 1 in the reference pages for the xa_bqual built-in function shows an incorrect binary translation. Replace the example with the following.

Example 1 Returns "0x227f06ca80", the binary translation of the branch qualifier for the Adaptive Server Enterprise transaction ID "0000000A_IphIT596iC7bF2#AUfkzaM_8DY6OE0". The Adaptive Server Enterprise transaction ID is first obtained using sp_transactions:

```
1> sp_transactions

xactkey                          type        coordinator starttime
state             connection dbid   spid  loid
failover                    srvname                          namelen
xactname
------------------------------ ----------- -----------
--------------------
----------------- ---------- ------ ------ -----------
------------------------- ------------------------------ -------
------------------------------------------------------------------
0x5316000006000000017e4885b0700 External    XA          Dec 9 2005 5:15PM
In Command        Attached      7    20    877
```

```
Resident Tx                NULL                              39
0000000A_IphIT596iC7bF2#AUfkzaM_8DY6OE0


1> select xa_bqual("0000000A_IphIT596iC7bF2#AUfkzaM_8DY6OE0", 0)
2> go

...
------------------------------------------------------------------
0x227f06ca80
```

8.4 *Reference Manual: Commands*

This section discusses updates and clarifications to *Reference Manual: Commands*.

8.4.1 Correction in *create table*

The Usage section of the create table reference pages includes the following information, which erroneously lists the limit of user-defined columns per table as 250:

- There can be as many as 2,000,000,000 tables per database and *250* user-defined columns per table. The number of rows per table is limited only by available storage.

The correct number of user-defined columns per table is *1024*.

8.4.2 Clarification in *dbcc*

Example 13 for the dbcc command is described as being an abbreviated output for the dbcc traceon(3604) command. You can obtain the same abbreviated version using dbcc serverlimits; running dbcc traceon(3604) results in a complete listing of limits in the server.

8.4.3 Additional usage information for *reorg rebuild*

Add these items to the Usage section of the reorg command description:

- Earlier versions of Adaptive Server restricted you from using reorg rebuild on all-pages locked tables (APL). As of Adaptive Server 15.0, you can run reorg rebuild on the entire table that uses APL. reorg rebuild rebuilds the entire table, copying the data to new sets of pages, and rebuilds all indexes.
- You cannot use the reorg rebuild sub commands (for example, compact, reclaim_space, and forwarded_rows) on APL tables.

- You can use reorg rebuild *table_name index_name* on APL tables.

8.4.4 Correction in *select*

In the reference section for the select command, a bulleted item describes how to add a new IDENTITY column and use precision to specify the numeric datatype:

- A specification to add a new IDENTITY column to the result table:

  column_name = identity(precision)

Beginning with Adaptive Server Enterprise version 12.5.4, you can specify the IDENTITY column with int, smallint, and tinyint, as well as precision.

- A specification to add a new IDENTITY column to the result table:

  column_name = identity(int | smallint | tinyint | precision)

  If you specify int, smallint, or tinyint, the resulting column is an integer. If you specify precision, the result is a numeric datatype.

8.4.5 Correction in *set proc_output_params off*

The example provided for set proc_output_params off in the set command reference pages includes two errors in the output:

In the first portion, the value of @x returned from the command is erroneously listed as "2." The value is 1, and appears in bold in the following corrected version:

In the second portion, the return parameters output is erroneously published as "3." The value is 2, and appears in bold in the following corrected version:

Example 14 Suppresses the output of parameter information:

```
1> create procedure sp_pout (@x int output) as select @x = @x + 1
2> go
1> set proc_output_params off
2> go
1> declare @x int
2> select @x = 1
3> exec sp_pout @x output
4> print "Value of @x returned from sproc is: %1!", @x
5> go
(1 row affected)
(return status = 0)

Value of @x returned from sproc is: 1
```

If you do not perform set proc_output_params off, the output after `(return status = 0)` includes the following:

```
Return parameters:
 -----------
           2
```

8.4.6 *setuser*

The setuser command has no effect when *creating* a database.

8.5 *Reference Manual: Procedures*

This section describes updates and clarifications in *Reference Manual: Procedures*.

8.5.1 *sp_monitorconfig* example updated with new parameters

Replace example 8 for sp_monitorconfig in the book with the following, which incorporates new parameters introduced in version 12.5.3:

Example 8 Using the optional parameter *result_tbl_name* to create a user table, saves the sp_monitorconfig result to this table:

```
1> create table sample_table
2> (Name varchar(35),Config_val int, System_val int, Total_val int,
3> Num_free int, Num_active int, Pct_act char(6), Max_used int,
4> Num_Reuse int, Date varchar(30))
5> go
```

The name of the table created becomes the second parameter of sp_monitorconfig:

```
1> sp_monitorconfig "locks", sample_table
2> go

(return status = 0)

1> sp_monitorconfig "number of alarms", sample_table
2> go

(return status = 0)

1> select * from sample_table
2> go

 Name                Config_val System_val Total_val Num_free Num_active
Pct_act Max_used Num_Reuse   Date
------------------- ----------- ---------- --------- --------- ----------
```

```
------- -------- ---------   ------------------------------
 number of locks          5000         684      5000      4915          85
1.70        117          0    Aug 23 2006  6:53AM
 number of alarms              40         0        40        28          12
30.00        13          0    Aug 23 2006  6:53AM
```

The result set saved to the table accumulates until you delete or truncate the table.

Note If sample_table is in another database, you must provide its fully qualified name in quotes.

8.5.2 *sp_sproc_columns* column_type description

In the reference pages for the sp_sproc_columns stored procedure, the table listing the result set includes the following description of the column_type column:

This field always returns a value:

- 0 = SQL_PARAM_TYPE_UNKNOWN
- 1 = SQL_PARAM_TYPE_INPUT
- 2 = SQL_PARAM_TYPE_OUTPUT
- 3 = SQL_RESULT_COL
- 4 = SQL_PARAM_OUTPUT
- 5 = SQL_RETURN_VALUE

This information is erroneous and should be removed from the description.

8.5.3 sp_poolconfig

The reference pages for sp_poolconfig have been revised. Replace the pages for sp_poolconfig with the following.

Description

Creates, drops, resizes, and provides information about memory pools within data caches.

Syntax

To create a memory pool in an existing cache, or to change pool size:

sp_poolconfig *cache_name* [, "*mem_size* [P|K|M|G]", "*config_pool*K" [, "*affected_pool*K"]]

To change a pool's wash size:

sp_poolconfig *cache_name*, "*affected_poolK* ", "wash=*size*[P|K|M|G]"

To change a pool's asynchronous prefetch percentage:

```
sp_poolconfig cache_name, "affected_poolK ",
    "local async prefetch limit=percent "
```

Parameters

- *cache_name* – is the name of an existing data cache.
- *mem_size* – is the size of the memory pool to be created or the new total size for an existing pool with the specified I/O size. The minimum size of a pool is 256 logical server pages. For a 2K logical page size server, the minimum size is 256K. Specify size units with P for pages, K for kilobytes, M for megabytes, or G for gigabytes. The default is kilobytes.
- *config_pool* – is the I/O size performed in the memory pool where the memory is to be allocated or removed.

  Valid I/O sizes are multiples of the logical page size, up to four times the amount.
- *affected_pool* – is the size of I/O performed in the memory pool where the memory is to be deallocated, or the pool's attributes such as 'wash size' and 'prefetch limit' are to be modified. If *affected_pool* is not specified, the memory is taken from the lowest logical page size memory pool.
- wash=*size* – 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 Creates a 16K pool in the data cache pub_cache with 10MB of space. All space is taken from the default 2K memory pool:

```
sp_poolconfig pub_cache, "10M", "16K"
```

Example 2 Creates 16MB of space to the 32K pool from the 64K pool of pub_cache:

```
sp_poolconfig pub_cache, "16M", "32K", "64K"
```

Example 3 Reports the current configuration of pub_cache:

```
sp_poolconfig pub_cache
```

Example 4 Removes the 16K memory pool from pub_cache, placing all of the memory assigned to it in the 2K pool:

```
sp_poolconfig pub_cache, "0K", "16K"
```

Example 5 Changes the wash size of the 2K pool in pubs_cache to 508K:

```
sp_poolconfig pub_cache, "2K", "wash=508K"
```

Example 6 Changes the asynchronous prefetch limit for the 2K pool to 15 percent:

```
sp_poolconfig pub_cache, "2K", "local async prefetch limit=15"
```

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 logical page size 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 for a 2K logical page size server.
- The combination of cache name and I/O size must be unique. In other words, you can specify only one pool of a given I/O size in a particular data cache in sp_poolconfig commands.
- 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 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 1: Data cache with default and user-defined caches

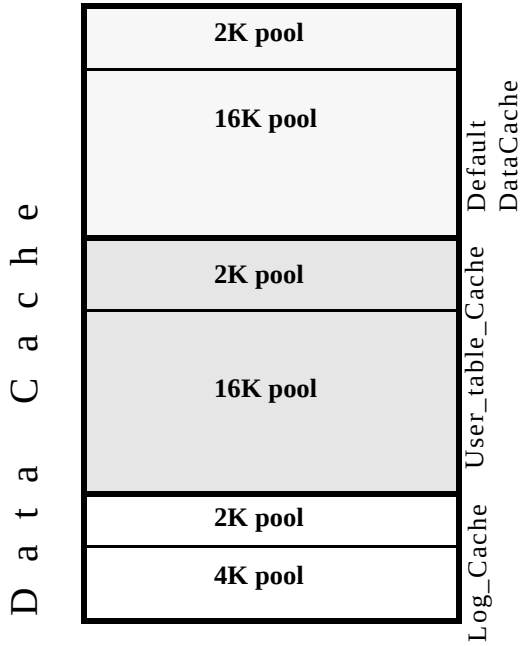

- You can create pools with I/O sizes up to 16K in the default data cache for a 2K page size server.
- The minimum size of a memory pool is 256 logical pages (for example, a 2K logical page size server, the minimum size is 512K). You cannot reduce the size of any memory pool in any cache to less than 256 pages by transferring memory to another pool.
- Two circumstances can create pools 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 the 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.
- Some dbcc commands and drop table perform only logical page size I/O. dbcc checkstorage can perform large I/O, and dbcc checkdb performs large I/O on tables and logical page size I/O on indexes.
- Most Adaptive Servers perform best with I/O configured for transaction 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. Run sp_sysmon to look for recommendations, or 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.
- See the *Performance and Tuning Guide* for information on the performance impact of changes to the asynchronous prefetch limit.

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.

Auditing

Values in event and extrainfo columns from the sysaudits table are:

Event	Audit option	Command or access audited	Information in extrainfo
38	exec_procedure	Execution of a procedure	• *Roles* – Current active roles • *Keywords or options* – NULL • *Previous value* – NULL • *Current value* – NULL • *Other information* – All input parameters • *Proxy information* – Original login name, if set proxy in effect

See also

System procedures sp_cacheconfig, sp_helpcache, sp_logiosize, sp_unbindcache, sp_unbindcache_all

8.6 Reference Manual: Tables

This section describes updates and clarifications in *Reference Manual: Tables*.

8.6.1 *syscomments*

In the reference page for syscomments, the description for the status column is blank. Insert the following information:

Name	Datatype	Description
status	smallint null	Bits indicating the status of the objects: • 0x1 – SYSCOM_TEXT_HIDDEN indicates that the text is hidden • 0x2 – Reserved for internal use • 0x4 – SYSCOM_QUOTED_ID_ON indicates that quoted identifiers were on when the object was created

8.6.2 *sysdatabases*

The reference page for sysdatabases includes Table 1-8, which lists the status3 control bits for sysdatabases. In it, the 0x0001 status was left blank. Insert the following description:

A normal or standard database, or a database without a proxy update in the create statement.

8.7 Installation guide

This section describes updates and clarifications in the installation guide for your platform.

8.7.1 Determining devices for a database

[CR# 407828] The installation guide gives incorrect information on how to determine which device a database is using. The infocenter.sybase.com search function may show many hits but it does not show the documents where the hits occur. See Appendix 2 of the installation guide for complete details.

Workaround:

- In versions of Adaptive Server Enterprise prior to 15.0, you use sysdevices to determine which device has a low through high virtual page range that includes the *vstart*.The device fragment whose *vstart* you used is on that device.
- Select the vdevno from sysusages matching the dbid retrieved in step1.

8.8 Performance and Tuning Guide: Monitoring

The behavior for sp_sysmon has changed so that you no longer have to specify the noclear option to prevent sp_sysmon from clearing the monitor counters. The noclear option is now the default behavior for sp_sysmon when you specify a sample interval.

If you run sp_sysmon using the begin_sample and end_sample options to begin and end the sample period, sp_sysmon always clears the monitor counters. Adaptive Server issues an error message if you run sp_sysmon with begin_sample or end_sample and the noclear option.

If you need to clear the monitor counters, use sp_sysmon with the clear option. For compatibility reasons, Adaptive Server accepts the noclear option as a valid parameter, but it does not affect the behavior of sp_sysmon.

9. Technical support

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you have any questions about this installation or if you need assistance during the installation process, ask the designated person to contact Sybase Technical Support or the Sybase subsidiary in your area.

10. Other sources of information

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

- The Getting Started CD contains release bulletins and installation guides in PDF format, and also contains other documents or updated information not included on the SyBooks CD. It is included with your software. To read or print documents on the Getting Started CD, you need Adobe Acrobat Reader, which you can download at no charge from the Adobe Web site using a link provided on the CD.
- The SyBooks CD contains product manuals and is included with your software. The Eclipse-based SyBooks browser allows you to access the manuals in an easy-to-use, HTML-based format.

  Some documentation can be provided in PDF format, which you can access through the PDF directory on the SyBooks CD. To read or print the PDF files, you need Adobe Acrobat Reader.

Refer to the *SyBooks Installation Guide* on the Getting Started CD, or the *README.txt* file on the SyBooks CD for instructions on installing and starting SyBooks.

- The Sybase Product Manuals Web site is an online version of the SyBooks 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 Sybase Product Manuals Web site, go to Product Manuals at http://www.sybase.com/support/manuals/.

10.1 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. Click Certification Report.
3. In the Certification Report filter select a product, platform, and time frame and then click Go.
4. Click a Certification Report title to display the report.

❖ **Finding the latest information on component certifications**

1. Point your Web browser to Availability and Certification Reports at http://certification.sybase.com/.
2. Either select the product family and product under Search by Base Product; or select the platform and product under Search by Platform.
3. Select Search to display the availability and certification report for the selection.

❖ **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.

10.2 Sybase EBFs and software maintenance

❖ **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.

11. Accessibility features

This document is available in an HTML version that is specialized for accessibility. You can navigate the HTML with an adaptive technology such as a screen reader, or view it with a screen enlarger.

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

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

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