



Release Bulletin

**Adaptive Server[®] Enterprise
Cluster Edition 15.7 SP100**

Linux

DOCUMENT ID: DC00761-01-1570100-02

LAST REVISED: February 21, 2014

Copyright © 2014 by Sybase, Inc. All rights reserved.

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

Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase trademarks can be viewed at the Sybase trademarks page at <http://www.sybase.com/detail?id=1011207>. Sybase and the marks listed are trademarks of Sybase, Inc. ® indicates registration in the United States of America.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

Java and all Java-based marks are trademarks or registered trademarks of Oracle and/or its affiliates in the U.S. and other countries.

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

IBM and Tivoli are registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

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

Sybase, Inc., One Sybase Drive, Dublin, CA 94568.

Contents

Product Summary	1
Adaptive Server Interoperability	1
Product Compatibility	5
Password Compatibility with Replication Server	5
Changes That Affect Existing Applications	5
Changes to Sybase Central	5
Backup Server and IBM Tivoli Storage Manager	6
Software Developer Kit (SDK) and Open Server	6
Command Line Version of ddlgen	6
Changed Functionality	6
XML Services and External Entity Reference	6
Support for db2 Server Class	7
Unsupported Features and Utilities	7
Installation and Configuration	8
Special Installation Instructions	8
Special Configuration Instructions	9
Configure Adaptive Server for Clients Using EPEP Algorithm	9
Special Upgrade and Downgrade Instructions	9
TIPSA Flag 0x4000 Error After Upgrade	10
Special Downgrade Instructions	10
Loading Database Dumps into Earlier Versions	11
Known Installation Issues for Adaptive Server Cluster Edition	11
Known Plug-In Issues for Adaptive Server Cluster Edition	15
Known Issues for Adaptive Server Cluster Edition	16

Documentation Updates	20
Enabling Custom Password Checks	20
isql -K keytab_file Parameter	23
Job Scheduler Parameter job_scheduler_interval	23
Job Scheduler Command sp_sjobdrop	24
Security Built-in Function Permissions	24
Shareable Temporary Tables	25
Hash-Based Update Statistics	25
Additional Information for Windows Failover in a High Availability System	25
Configure Adaptive Server for Failover on Windows	26
Setting maximum failed logins	28
Changing the maximum failed logins for specific roles	28
Changing the maximum failed logins for specific logins	29
Values for lock timeout pipe active	29
dbcc page Parameter logical	29
Behavior of Concurrent DDLs and reorg defrag	29
Updates to Third-Party Licensing	30
Obtaining Help and Additional Information	30
Technical Support	30
Downloading Product Updates	31
Product and Component Certifications	31
Accessibility Features	31

Product Summary

This release bulletin provides late-breaking information about Adaptive Server[®] Enterprise version 15.7 SP100. A more recent version may be available on the Web.

Adaptive Server Enterprise server and client components are distributed on separate CDs or DVDs.

See the Cluster Edition installation guide for operating system requirements, and <http://certification.sybase.com> for additional supported operating systems.

Installation Kit

The installation kit includes:

- Server media (CD or DVD)
- PC-Client media

Adaptive Server Interoperability

Interoperability of Adaptive Server against other Sybase[®] products, across different platforms, versions, and client products.

Interoperability between big-endian and little-endian platforms has been verified. Windows, Linux x86-32, Linux x86-64, Sun Solaris x86-32, and Sun Solaris x86-64 are little-endian platforms. IBM AIX, Linux on Power, Sun Solaris SPARC, and HP-UX on Itanium are big-endian platforms.

Note: To use new features of Adaptive Server Cluster Edition, make sure that your client supports them. See the client-specific documentation for information about the features your client supports. You may need to upgrade your client to use certain Cluster Edition features.

Table 1. Supported Client Platforms for Adaptive Server

Client Name	Version	Supported Platform
Open Client™/Open Server™	15.7	<ul style="list-style-type: none"> • AIX 32-bit • AIX 64-bit • HP-UX IA 32-bit • HP-UX IA 64-bit • Linux x86 32-bit • Linux x86-64 64-bit • Linux Pseries 32-bit • Linux Pseries 64-bit • Solaris SPARC 32-bit • Solaris SPARC 64-bit • Solaris x86 32-bit • Solaris x86-64 64-bit • Windows x86 32-bit • Windows x86-64 64-bit
	15.5	<ul style="list-style-type: none"> • 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 x86-64 64-bit • Linux Pseries 32-bit • Linux Pseries 64-bit • Solaris SPARC 32-bit • Solaris SPARC 64-bit • Solaris x86 32-bit • Solaris x86-64 64-bit • Mac OS X Intel 32-bit • Windows x86 32-bit • Windows x86-64 64-bit

Client Name	Version	Supported Platform
Adaptive Server Enterprise Extension Module for Python	15.7	<ul style="list-style-type: none"> • AIX 64-bit • HP-UX IA 64-bit • Linux x86-64 64-bit • Linux Pseries 64-bit • Solaris SPARC 64-bit • Solaris x86-64 64-bit • Windows x86-64 64-bit
Adaptive Server Enterprise Extension Module for PHP	15.7	<ul style="list-style-type: none"> • AIX 64-bit • HP-UX IA 64-bit • Linux x86-64 64-bit • Linux Pseries 64-bit • Solaris SPARC 64-bit • Solaris x86-64 64-bit • Windows x86-64 64-bit
Adaptive Server Enterprise Database Driver for PERL	15.7	<ul style="list-style-type: none"> • AIX 32-bit • HP-UX IA 32-bit • Linux x86-64 64-bit • Linux Pseries 32-bit • Solaris SPARC 32-bit • Solaris x86-64 32-bit • Windows x86-64 64-bit
jConnect™ for JDBC™	6.0.x, 7.0.x	<ul style="list-style-type: none"> • All
ODBC by Sybase	15.7	<ul style="list-style-type: none"> • Linux x86 32-bit • Linux x86-64 64-bit • AIX 64-bit • HP-UX IA 64-bit • Linux Pseries 64-bit • Solaris SPARC 64-bit • Solaris x86-64 64-bit • Windows x86 32-bit • Windows x86-64 64-bit

Product Summary

Client Name	Version	Supported Platform
	15.5.x	<ul style="list-style-type: none"> • Linux x86 32-bit • Linux x86-64 64-bit • Mac OS X Intel 32-bit • Windows x86 32-bit • Windows x86-64 64-bit
OLE DB by Sybase	15.7, 15.5.x	<ul style="list-style-type: none"> • Windows x86 32-bit • Windows x86-64 64-bit
ADO.NET	2.x, 4.x	<ul style="list-style-type: none"> • Windows x86 32-bit • Windows x86-64 64-bit
Replication Server [®]	15.7, 15.7.1	<ul style="list-style-type: none"> • IBM AIX POWER 64-bit • HP-UX IA 64-bit • Linux x86-64 64-bit • Linux on POWER 64-bit • Solaris SPARC 64-bit • Solaris x86-64 64-bit • Windows x86 32-bit • Windows x86-64 64-bit
	15.5, 15.6	<ul style="list-style-type: none"> • IBM AIX POWER 64-bit • HP-UX IA 64-bit • Linux x86 32-bit • Linux x86-64 64-bit • Linux on POWER 64-bit • Solaris SPARC 64-bit • Solaris x86-64 64-bit • Windows x86 32-bit • Windows x86-64 64-bit

Client Name	Version	Supported Platform
	15.2	<ul style="list-style-type: none"> • IBM AIX POWER 32-bit • IBM AIX POWER 64-bit • HP-UX IA 64-bit • Linux x86 32-bit • Linux x86-64 64-bit • Linux on POWER 64-bit • Solaris SPARC 32-bit • Solaris SPARC 64-bit • Windows x86 32-bit • Windows x86-64 64-bit

Note: Adaptive Server is supported on Sybase Control Center 3.2.8.

Product Compatibility

Learn about Adaptive Server compatibility.

Password Compatibility with Replication Server

There are some compatibility issues with replicating logins and roles between different versions of Adaptive Server.

You can replicate logins from Adaptive Server version 15.0.2 and later to earlier server versions only during the password downgrade period when **allow password downgrade** is set to 1.

You cannot replicate roles with passwords from Adaptive Server 15.7 and later to earlier server versions.

Changes That Affect Existing Applications

Learn about changes in Adaptive Server 15.7 SP100 that affect your existing applications.

Changes to Sybase Central

The Sybase Central ASE plug-in is no longer included in the Adaptive Server installation image. Use Sybase Control Center 3.2.8 to manage your Adaptive 15.7 SP100 servers.

Sybase no longer maintains Sybase Central. However, if required, you can download Sybase Central plug-ins by going to the <http://www.sybase.com/downloads> site, then clicking **Maintenance Releases and Bug Fixes (EBFs)**.

Backup Server and IBM Tivoli Storage Manager

Use a local backup server for dump and load to the IBM Tivoli Storage Manager; you cannot use a remote backup server for dumping and loading to the IBM Tivoli Storage Manager.

You can configure the IBM Tivoli Storage Manager on a different host machine than the local backup server.

Software Developer Kit (SDK) and Open Server

For information about changes that affect Software Developer Kit (SDK), see the most recent release bulletins for these products on the Sybase Web site.

For information about new features, see the *New Features Bulletin Open Server and SDK for Windows, Linux, and UNIX*.

Command Line Version of ddlgen

The location of the command line version of **ddlgen** has been changed.

In Adaptive Server versions 15.7 ESD #1 and earlier, the command line version of **ddlgen** is in:

- UNIX – \$SYBASE/ASEP/bin
- Windows – %SYBASE%\ASEP\bin

In Adaptive Server version 15.7 ESD #2 and later, **ddlgen** is located in:

- Unix – \$SYBASE/ASE-15_0/bin
- Windows – %SYBASE%\ASE-15_0\bin

Changed Functionality

Learn about late-breaking functionality changes in Adaptive Server version 15.7 SP100.

For detailed information, see the *Clusters Users Guide*.

XML Services and External Entity Reference

As of release 15.7 SP100, you must use **sp_configure 'enable xml', 2** to enable an external entity reference in XML documents.

In previous releases, executing **sp_configure 'enable xml', 1** enabled XML services with an external entity reference.

sp_configure 'enable xml', 1 enables XML services but not with external entity reference in XML documents.

Support for db2 Server Class

Adaptive Server does not support server class **db2**.

To use **db2**, migrate your **db2** server class to the **direct_connect** class.

Unsupported Features and Utilities

Learn about features that are not supported in the Cluster Edition.

These utility programs are not supported, or have minimum support:

- **sqlupgrade**
- **sqlloc**
- **srvbuild**

These Adaptive Server features are not supported:

- In-memory databases introduced in Adaptive Server 15.5.
- Adaptive Server disk mirroring.
- In the Cluster Edition, high-availability functionality is integrated into Adaptive Server.
- The Cluster Edition architecture provides for up to 32 instances in a cluster. However, Cluster Edition 15.5 and later is currently certified for, and supported on, 4 instances only.

These Replication features are not supported:

- Multi-Path Replication™
- Metadata reduction
- Repdef elimination
- In-row LOB replication
- Replication of compressed data
- Replication using ASE Replicator

Note: The ASE-15_0/bin directory includes a utility called **iofenceutil**, which duplicates the functionality of **qrmutil -fence_capable**, but does not require the quorum device, and also returns an exit code. This is an internal utility used by **sybcluster** and is not designed for regular use.

Installation and Configuration

Get last-minute information about installation, configuration, and upgrading and downgrading that was omitted from your installation guide, or that needs special emphasis.

Sybase strongly recommends that you install the Adaptive Server Enterprise Cluster Edition into its own directory. If you cannot do this, install the Cluster Edition last.

Warning! Sybase recommends that you do not install an older version of a Sybase product in a \$SYBASE directory that includes a newer Sybase product. For example, installing Replication Server® 15.0 on top of the Cluster Edition may make either the Cluster Edition or Replication Server inoperable, depending on the choices you select during installation.

Special Installation Instructions

Learn about special installation instructions for this version of Adaptive Server.

OCS Connectivity Driver Files Not Installed

When installing Adaptive Server, if the destination directory has an existing OCS directory that is a later version than the Adaptive Server version being installed, the Adaptive Server installer does not install another OCS directory, and the driver files required by Adaptive Server will not be available. The workaround is to install the product with the oldest OCS version first.

Installing the Adaptive Server Plug-in

Before installing the ASE plug-in on top of Adaptive Server, close all Java applications. If you do not, the installation will fail. The entry in the log file states that a problem occurred while attempting to overwrite a JRE 7 file because the file is open.

Install Adaptive Server and all associated plug-ins in the same location.

Amendments to the Response File for Silent Installation

For Adaptive Server 15.7 ESD#2 and later, the installer requires non-NULL passwords that are at least six characters long for the Adaptive Server sa login, and the Sybase Control Center logins uafadmin and sccadmin.

To accommodate this, the response file should include these additional rows:

```
SY_CFG_ASE_PASSWORD=<ASE sa password>
```

```
CONFIG_SCC_CSI_SCCADMIN_PWD=<Sybase Control Center admin password>
```

```
CONFIG_SCC_CSI_UAFADMIN_PWD=<Sybase Control Center agent admin  
password>
```

The passwords for sccadmin and uafadmin logins need not be the same as the password for sa login.

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 directory.

Installing Adaptive Server Enterprise Cluster Edition with Sybase IQ 15.1

Sybase recommends that you do not install Sybase IQ™ 15.1 in a directory that includes the Cluster Edition.

Special Configuration Instructions

Special configuration instructions for this version of Adaptive Server.

Configure Adaptive Server for Clients Using EPEP Algorithm

You can use **sp_configure** to configure an Adaptive Server to require the Extended Plus Encrypted Password (EPEP) login protocol.

The **sp_configure** parameter **net password encryption reqd** supports a value of 3, which indicates the server should only allow incoming clients that are using EPEP login protocol. The values 0, 1, and 2 also allow EPEP login protocol to be used when a client that supports the login protocol attempts to use it with an Adaptive Server that implements the EPEP login protocol.

Setting the value to 2 or 3 increases network memory to support the maximum configured connections using this protocol. The **additional network memory** configuration parameter dynamically adds more memory to the network memory pool used by EPEP. When the value is set to 3, the KPP Handler goes into sleep status, because there is no need to provide new RSA key pair for every connection. You can use the **sp_who** command to check the KPP Handler status.

Note: Adaptive Server supports two versions of the login protocol using RSA asymmetric encryption. See "Securing login passwords on the network" in the *Security Administration Guide*.

Special Upgrade and Downgrade Instructions

Get last-minute instructions for upgrading, downgrading, and migrating different versions of Adaptive Server.

TIPSA Flag 0x4000 Error After Upgrade

The flag 0x4000 in TIPSA is set in Adaptive Server versions earlier than 15.7. This flag can cause unexpected errors after an upgrade.

The TIPSA 0x4000 flag is used by the LOB compression feature in Adaptive Server 15.7. After an upgrade to Adaptive Server 15.7 and later, an LOB column for which this flag was set in earlier versions is mistakenly treated as a compressed LOB column, which results in unexpected errors. To correct the flag issue after an upgrading to Adaptive Server 15.7 and later, use:

```
dbcc rebuild_text()
```

Special Downgrade Instructions

Learn about special downgrading instructions for this version of Adaptive Server.

Use `sp_downgrade_esd` to Downgrade From Adaptive Server 15.7 SP100 to Earlier Versions

If you are downgrading from Adaptive Server 15.7 SP100 to 15.7 ESD #4, 15.7 ESD #3, 15.7 ESD #2, 15.7 ESD #1 and 15.7, use **sp_downgrade_esd**(not **sp_downgrade**) to downgrade both Adaptive Server and any databases you may have upgraded to 15.7 SP100. See *Downgrading to an Earlier Version of Adaptive Server 15.7* in the installation guide for your platform.

Downgrading Cluster Edition From 15.7 ESD #2 to ESD #1 with sybcluster

Use the following step to downgrade Cluster Edition 15.7 ESD #2 to ESD #1 using sybcluster:

1. Back up \$SYBASE for 15.7 ESD #2.
2. Restart the cluster in single-user mode.
3. Run **sp_downgrade_esd** on each database.

Note: For temporary databases, running **sp_downgrade_esd** is optional.

4. Install 15.7 ESD #1 on top of the current \$SYBASE.

Note: Do not configure SCC from the installer.

5. Replace SCC-3_2 from the backup copy of 15.7 ESD #2 \$SYBASE.

Note: This allows you to keep the 15.7 ESD #2 version of SCC for use after the downgrade.

6. Replace SDCADMIN-15_0 from the backup copy of 15.7 ESD #2 \$SYBASE.
This allows you to keep the 15.7 ESD #2 version of sybcluster for use after the downgrade.
7. Start the SCC agent on each node.
8. Run **sybcluster** and execute **start instance <instance 1>** to bring up the first instance.
9. Run **installmaster** on the first instance.

10. Run **sybcluster** and execute **shutdown cluster**.

11. Run **sybcluster** and execute **start cluster**.

Loading Database Dumps into Earlier Versions

You cannot dump a database from a 15.7 SP100 version of Adaptive Server and load it into a database on a server running a release earlier than 15.7 SP100.

Known Installation Issues for Adaptive Server Cluster Edition

Learn about known installer issues and workarounds. Known issues are listed in descending order of Change Request (CR) numbers.

CR #	Description
739079	<p>Update of Adaptive Server returns "Arithmetic overflow occurred" error for database size check.</p> <p>If any system database is greater than 2GB, updating an Adaptive Server with the Adaptive Server installer, or updatease (on UNIX) or updatease.exe (on Windows) utility, will fail with the error "Arithmetic overflow occurred."</p> <p>Workaround: To update an Adaptive Server, manually run the updatease utility, or for Windows the updatease.exe utility with the -k argument. The -k option tells the utility to skip the database size check.</p> <p>The utilities are located here:</p> <ul style="list-style-type: none"> • UNIX – \$SYBASE/ASE-15_0/bin/updatease • Windows – %SYBASE%\ASE-15_0\bin\updatease.exe

CR #	Description
736300	<p>Replication Manager Plug-in (RMP) cannot be loaded after Adaptive Server is installed upon RMP</p> <p>Installing Adaptive Server and Sybase Central into the same directory as Replication Server and Replication Manager Plug-in can cause the Replication Manager Plug-in to not load properly.</p> <p>Workaround: Unset the SYBASE_JRE7 environment variable that is set in the SYBASE {csh env sh} file before starting Sybase Central.</p> <p>Rather than unsetting the SYBASE_JRE7 environment variable in the SYBASE file by commenting out the line, unset it on the command line in the same window that is used to start Sybase Central. This prevents other processes from being affected. Regardless of platform, do the following:</p> <ol style="list-style-type: none"> 1. Open a new window. 2. Source the SYBASE file. 3. Unset the SYBASE_JRE7 environment variable. See commands below. 4. Execute the runSybaseCentral script. <ul style="list-style-type: none"> • for C-Shell (SYBASE.csh): unsetenv SYBASE_JRE7 • for Bourne Shell (SYBASE.sh): SYBASE_JRE7=export SYBASE_JRE7 • for Bash(?) Shell (SYBASE.env): SYBASE_JRE7=
692496	<p>Configuring tempdb device for HA</p> <p>When installing an Adaptive Server to be used in an HA configuration, the sp_companion stored procedure configure command may fail due to duplicate device names for the tempdbdev logical device.</p> <p>Workaround: Do not specify the "Tempdb Device" (set the value of this field to blank) when configuring the secondary server during ASE installation.</p> <ul style="list-style-type: none"> • If additional tempdb space is required on either server, log in to this server after installation is complete and use the alter database command to increase the size of tempdb. • If additional disk space is required to allocate the tempdb database, use the disk init command to create a new database device for tempdb. • If a new device is created, the logical device name should be unique between the two servers.

CR #	Description
688101	<p>Cannot start Sybase Central for Sybase IQ 15.4 after uninstalling Adaptive Server 15.7.</p> <p>If you install Adaptive Server Enterprise 15.7 and Sybase IQ 15.4 in the same directory, then uninstall Adaptive Server, you cannot then start Sybase Central for IQ. This problem occurs because shared files are removed by the Adaptive Server uninstaller.</p> <p>Workaround: If the version of Sybase Central installed with IQ 15.4 cannot be started after you uninstall Adaptive Server 15.7:</p> <ol style="list-style-type: none"> 1. Re-create the \$SYBASE/shared/JavaHelp-2_0 if necessary. 2. Copy \$SYBASE/IQ-15_4/java/jh.jar to the \$SYBASE/shared/JavaHelp-2_0 directory. 3. Reload or register the IQ plug-in in Sybase Central.
678912	<p>Cannot use DBISQL from Sybase Central after installing RMP on top of Adaptive Server 15.7.</p> <p>Workaround: To enable DBISQL to launch, add the following into scjview.sh java arguments.</p> <pre>-Disql.repositoryDirectory="\$SYBROOT/DBISQL/bin"</pre>
671351	<p>Installing Adaptive Server 15.7 on top of Replication Agent™ 15.2 causes SySAM errors.</p> <p>Workaround: Before installing Adaptive Server 15.7, modify the SySAM file permissions to include write permissions.</p>
639623	<p>Cannot Start Adaptive Server Interactive SQL from Sybase Central.</p> <p>If you install Replication Manager into the same directory as Adaptive Server, the Interactive SQL cannot start from Sybase Central, and issues: Interactive SQL could not load the "SQL Anywhere plugin", its "saip11.jar" file has been removed or has been deleted. You will not be able to connect to the databases handled by that plugin.</p> <p>Workaround:</p> <ul style="list-style-type: none"> • Before installing Replication Manager, rename shared/sybcentral600/scjview.sh to shared/sybcentral600/scjview.sh.save. • After installing Replication Manager, rename shared/sybcentral600/scjview.sh.save back to shared/sybcentral600/scjview.sh.

CR #	Description
589483	<p>If the <code>df</code> command stops responding, the installation stops responding before the pre-installation summary pane appears.</p> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Identify the network file system (NFS) mounts that may have problems: <code>strace -e statfs, statfs64 df</code> 2. Use the umount command to unmount an identified NFS mount: <code>umount -l <path></code> 3. Repeat the above steps until there are no more NFS mounts that have stopped responding. <p>The umount command requires root permission.</p> <p>If umount command failed to unmount the offending NFS, you may need to fix the problem at the source or restart the machine.</p>
588793	<p>Installing Adaptive Server 15.5 and later versions that use Install Anywhere on older directories that have been installed with InstallShield Multiplatform can lead to incompatibility issues.</p> <p>These issues occur while using products installed with different installer technologies:</p> <ul style="list-style-type: none"> • If you install products using Install Anywhere or InstallShield Multiplatform on top of a product installed using the other installer, the same files included in both installers are silently overwritten by the later installation. • If you run either the Install Anywhere or InstallShield Multiplatform uninstaller, the same files installed by both installers are removed without any indicating messages. <p>Workaround: Do not install products that use Install Anywhere and products that use InstallShield Multiplatform in the same directory.</p>
583979	<p>The installer does not validate feature names specified in the response file when you install in silent mode.</p> <p>Workaround: Ensure that the specified feature names are correct.</p>

Known Plug-In Issues for Adaptive Server Cluster Edition

Learn about known issues and workarounds for Interactive SQL and the Adaptive Server Enterprise plug-in. Known issues are listed in descending order of Change Request (CR) numbers.

CR#	Description and Workaround
713282	<p>Refresh of Sybase Control Center instance deletes the Adaptive Server CE plug-in.</p> <p>When you run SYBASE/SCC-3_2/bin/sccinstance -refresh to refresh Sybase Control Center instances, the Adaptive Server Cluster Agent plug-in files, <code>\$\$SYBASE/SCC-3_2/instances/SCC_instance/plugins/cluster_name</code> are deleted.</p> <p>Workaround: Before you run sccinstance -refresh, backup the <code>\$\$SYBASE/SCC-3_2/instances/SCC_instance_name/plugins/cluster_name</code> directory.</p> <hr/> <p>Note: Do not copy or back up the files to the same directory.</p> <hr/> <p>After running \$\$SYBASE/SCC-3_2/bin/sccinstance -refresh, restore the backup files, and then copy these files:</p> <ol style="list-style-type: none"> 1. Copy <code>\$\$SYBASE/SCC-3_2/templates/plugins/com.sybase.ase.cluster/15.0.1/lib/*</code> to <code>\$\$SYBASE/SCC-3_2/instances/SCC_instance_name/plugins/cluster_name/lib</code>. 2. Copy <code>\$\$SYBASE/SCC-3_2/templates/plugins/com.sybase.ase.cluster/15.0.1/ar1.xml</code> to <code>\$\$SYBASE/SCC-3_2/instances/SCC_instance_name/plugins/cluster_name/ar1.xml</code> 3. Copy <code>\$\$SYBASE/SCC-3_2/templates/plugins/com.sybase.ase.cluster/15.0.1/mbean-descriptor.xml</code> to <code>\$\$SYBASE/SCC-3_2/instances/SCC_instance_name/plugins/cluster_name/mbean-descriptor.xml</code>

CR#	Description and Workaround
709646	<p>Unable to authenticate SCC agents.</p> <p>Due to security service changes in SCC 3.2.7, the following related features in Adaptive Server plug-in have been removed:</p> <ul style="list-style-type: none"> • Adaptive Server to Unified Agent role mappings on the Roles Folder and Role Properties windows. • Unified Agent operations and attributes on the Server and ASE Plug-in Properties windows (Show agent attributes and operations check box on the Preferences page). • Authentication and Security Modules on the Server Properties Agent page.
693548	<p>A new cluster cannot be agent enabled.</p> <p>After creating a cluster using the ASE plug-in, the cluster does not appear as agent enabled.</p> <p>Workaround: Shut down and restart the ASE plug-in.</p>

Known Issues for Adaptive Server Cluster Edition

Learn about known issues and apply workarounds for Adaptive Server Cluster Edition. Known issues are listed in descending order of Change Request (CR) numbers.

Sybase does not include system problem reports (SPRs) and closed problem reports (CPRs) with Adaptive Server Enterprise. You can search the Web site for solved cases. Click **Support > Services > Solved Cases**.

CR #	Description
739773	<p>The load database ... with listonly=load_sql command may generate load commands which, if executed, may require manually altering the target database for the load to succeed.</p> <p>This issue can occur in cases where an alter database off command was run on the source database.</p> <p>Workaround: Manually alter the target database.</p>
739080	<p>The show_condensed_text function returns a NULL value for queries in the statement cache when the enable_literal_autoparam configuration parameter is enabled.</p> <p>Workaround: Do not use the show_condensed_text function when enable_literal_autoparam is enabled.</p>
738744	<p>A cumulative dump cannot be used as a meta-data file for the sybdumptran utility.</p> <p>Workaround: None.</p>

CR #	Description
738675	<p>load database with verify does not support automatic physical database rearrangement.</p> <p>If load database with verify is executed, the automatic remapping between the source and target database to ensure that the data segment is on data devices and the log segment is on log devices is not performed. This may result in the behaviour of releases prior to 15.7 SP100, which is the data segment may be mapped to log devices and the log segment may be mapped to data devices in the target database.</p> <p>Workaround: Sybase recommends that you use the dump database with verify command to verify the database.</p>
696072	<p>In some Cluster Edition configurations, Adaptive Server 624 and 69x errors may occur during database replication of a database or replication of tables within a database.</p> <p>In active-active or active-passive cluster configurations, you can configure RepAgent to run on only one node of the cluster. In rare circumstances, when RepAgent reads an older image of database log pages that is on disk while the latest image is on another node, you may see Adaptive Server 624 and 69x errors.</p> <p>Workaround: Sybase recommends that for:</p> <ul style="list-style-type: none"> • Active-passive configuration – run RepAgent on the active node to ensure that data and log pages are on the same node as RepAgent. • Active-active configuration – start Adaptive Server with the 16872 trace flag to prevent log page corruption and 69x errors. However, using this flag degrades the server performance.
695625	<p>Confidentiality service causes error.</p> <p>This message is reported by Open Client applications connecting to Adaptive Server Enterprise:</p> <pre>ct_send(): network packet layer: internal Client Library error: State error: trying to write when connection is expecting a read.</pre> <p>This message is reported in some cases when Kerberos message confidentiality service is turned on.</p> <p>Workaround: Reconnect to the server with confidentiality service turned off.</p>
684556	<p>The select for update semantics introduced in version 15.7 do not support datapages-locked tables.</p> <p>select for update support exists only for datarows-locked pages.</p>

CR #	Description
680041	<p>The value of the server network name as defined by the <code>srvnetname</code> column for a loopback server must be the name of the cluster.</p> <p>For the <code>sp_opt_querystats</code> stored procedure to succeed, the value of the <code>srvnetname</code> column for a loopback remote server must be the name of a cluster in the cluster interfaces file and not the name of a cluster instance.</p> <p>Workaround: Make sure that the <code>srvnetname</code> column for the loopback remote server is the name of the cluster.</p>
639813	<p>An error is raised when renaming a read only database</p> <p>Error (3906) can occur when executing <code>sp_renamedb</code> on a database that has 'read only' status. Also, when attempting to turn off the 'read only' status on a database that you have attempted to rename, error 3501 is raised.</p> <p>Workaround: Execute <code>sp_renamedb</code> stored procedure from master database context as follows:</p> <pre data-bbox="319 682 626 808">use master go sp_renamedb <dbname>, <new_dbname> go</pre>
627114	<p>On Linux RHEL 5, not all shared memory is released when hugepages are in use.</p> <p>When Adaptive Server is configured with hugepages on the Linux platform, not all of the shared memory is released when shutting down Adaptive Server. See the Linux bug description here: https://bugzilla.redhat.com/show_bug.cgi?id=593131.</p> <p>Workaround: Install the Linux kernel patch for this bug to resolve the memory leak.</p>
615204	<p>If enable DTM is set, the distributed recovery feature is disabled and all databases are recovered only on the cluster coordinator, rather than on different instances of the cluster.</p> <p>Workaround: None.</p>
595923	<p>There are two known issues with transfer table:</p> <ul data-bbox="323 1229 1185 1433" style="list-style-type: none"> • When creating a unique index that is using the <code>ignore_dup_key</code> property, importing a row with a key that is already present in the table leads to an error, and the import is aborted. This differs from when a duplicate key is inserted with <code>bcp</code> or with a regular <code>insert</code> statement, since in those cases, the row is discarded, and the transaction continues. • When an <code>insert</code> trigger exists, the trigger is not fired when data is inserted through <code>transfer table...from</code>. <p>Workaround: None.</p>

CR #	Description
578921	<p>On Linux SUSE Enterprise 11 systems, configuring an Adaptive Server fails with SEGV while retrieving user ID and password from Network Information Service (NIS).</p> <p>While creating the PCA-JVM, Adaptive Server can retrieve ID and password information using Network Security Services (NSS), but not NIS.</p> <p>Workaround: Use either of these:</p> <ul style="list-style-type: none"> • In <code>/etc/nssswitch.conf</code>, specify the retrieval order to be either "compat" or "files NIS." • Add the ID and password for the user starting the Adaptive Server to the local files <code>/etc/passwd</code> and <code>/etc/shadow</code>. This ensures that the ID and password information is available locally, and that contacting NIS is unnecessary. <p>If you cannot add the user information to local files, ensure that the Name Service Cache Daemon (NSCD) is running on the server host machine. The NSCD maintains a cache of the ID and password information that is consulted prior to dispatching a request to the NIS for this information.</p>
576652	<p>SySAM2 enabled products incorrectly determine the number of cores and chips on quad-core processors.</p> <p>This may be as a result of your Adaptive Server small business edition or chip license not being activated and having gone into a grace period.</p> <p>Workaround: Remove this line from your SySAM properties file, if it exists:</p> <pre>cpuinfo.mechanism=NOAPICCHECK</pre> <p>If the problem still remains, set the following environment variable, then run the products:</p> <pre>(sh) export SYBASE_SAM_CPUINFO=CPUID (csh) setenv SYBASE_SAM_CPUINFO_CPUID</pre> <p>If the problem still remains, please contact Sybase Technical Support.</p>

CR #	Description
<p>558020</p>	<p>dsedit fails with translation table syntax errors.</p> <p>When you try to enter a new server entry or a port number, the dsedit utility fails and displays parsing errors similar to:</p> <pre data-bbox="323 326 1174 583">Warning: translation table syntax error: Unknown keySYM name: osfActivate. Warning: ... found while parsing ':<Key>osfActivate: ManagerParentActivate()'. Warning: String to TranslationTable conversion encountered errors. Warning: translation table syntax error: Unknown keySYM name: osfBeginLine. Warning: ... found while parsing ':<Key>osfBeginLine: ManagerGadgetTraverseHome()'</pre> <p>These errors indicate that the X-server program that renders graphics is not cached or set up as the default for the environment.</p> <p>Workaround: Set the environment variable XKEYSYMDB:</p> <pre data-bbox="323 718 1042 748">% setenv XKEYSYMDB "/usr/share/X11/XKeysymDB"</pre>
<p>483651</p>	<p>Incorrect cluster failover may occur during diagnostic shared-memory dump.</p> <p>If you start an instance while the rest of the cluster is performing a diagnostic shared memory dump, the instance that is starting may incorrectly perform a cluster takeover. This occurs only if automatic cluster takeover is set to 1, or if the --cluster_takeover option is passed to the data server. In environments without i/o fencing enabled, this may lead to data corruption.</p> <p>Workaround: Do not start an instance during a diagnostic shared-memory dump. Configure automatic cluster takeover to 0.</p>

Documentation Updates

Read about updates, corrections, and clarifications to the documentation released with Adaptive Server Cluster Edition.

Enabling Custom Password Checks

This information supplements the "Enabling Custom Password Checks" documentation in the Security Administration Guide.

In Adaptive Server 15.7 SP100 and later, **sp_extrapwdchecks** allows NULL values for **caller_password** and **loginame** parameters.

The **caller_password** parameter is NULL when:

- The system security officer creates a new login account using `create login` command.
- The system security officer modifies the login account's password using **alter login ... modify password** command.

The **loginame** parameter is NULL when:

- The system security officer creates a new login account using the **create login** command.

To implement password history checks, create a new user table to store password histories:

```
create table pwdhistory
(
    name varchar(30)not null, -- Login name.
    password varbinary(30)not null, -- old password.
    pwdate datetime not null, -- datetime changed.
    changedby varchar(30)not null -- Who changed.
)
go
```

Create a new stored procedure **master.dbo.sp_extrapwdchecks** which saves previously used passwords in an encrypted form in the `pwdhistory` table and disallows reuse of used passwords. The **sp_extrapwdchecks** user-defined stored procedure is called by Adaptive Server automatically when either the **create login** or **alter login ... modify password** commands are executed. The following is an example of the implementation of **sp_extrapwdchecks**:

```
create proc sp_extrapwdchecks
(
    @caller_password varchar(30) = NULL, -- the current password of caller
    @new_password    varchar(30), -- the new password of the target acct
    @loginame        varchar(30) = NULL -- user to change password on
)
as
begin
    declare @current_time datetime,
            @encrypted_old_pwd varbinary(30),
            @encrypted_new_pwd varbinary(30),
            @salt varchar(8),
            @changedby varchar(30),
            @cutoffdate datetime

    select @changedby = suser_name()
    select @salt = null

    -- NOTE : caller_password and/or loginame arguments can be null.
    -- In these cases, password history checks should be skipped.

    -- @loginame is null when SSO creates a new login account
    -- using "create login" command.

    -- @caller_password is null when

    -- 1. SSO creates a new login account using
    -- "create login" command.
```

Documentation Updates

```
-- 2. SSO modifies the login account's password using
-- "alter login ... modify password" command.

-- Business logic for custom password checks should be
-- implemented here.

-- If there is no need to maintain password history, return
-- from here.

if (@loginame is NULL)
begin
    return 0
end

-- Change this line according to the needs of your installation.

-- This checks below keep history of 12 months only.

select @current_time = getdate(),@cutoffdate = dateadd
    (month, -12, getdate())

delete master..pwdhistory
where name = @loginame
and pwdate < @cutoffdate

select @salt = substring(password, 1, 8) from master..pwdhistory
    where pwdate =
        (select max(pwdate) from master..pwdhistory where
            name=@loginame)and name=@loginame

if @salt is null
begin
    select @salt = substring(hash
        (password_random(8), 'sha1'), 1, 8)
end

select @encrypted_new_pwd = @salt + hash
    (@salt + @new_password, 'sha1')

if not exists ( select 1 from master..pwdhistory
    where name = @loginame and password = @encrypted_new_pwd )

begin
    -- new password has not been used before

    if (@loginame != @changedby)
    begin
        return 0
    end

    -- Save old password
    select @encrypted_old_pwd = @salt + hash
        (@salt + @caller_password, 'sha1')
```

```

        insert master..pwdhistory
            select @loginame,
                @encrypted_old_pwd,@current_time, @changedby
        return (0)
    end
else
begin
    raiserror 22001 --user defined error message
end
end
end
go

```

isql -K keytab_file Parameter

The documentation for **isql** incorrectly describes the **-K keytab_file** parameter.

The description in the Utility Guide for the **isql -K keytab_file** parameter currently reads:

-K keytab_file

specifies the path to the keytab file used for authentication in DCE. `keytab_file` specifies a DCE keytab that contains the security key for the user name you specify in `-U`. Create keytab files using the DCE `dcecp` utility. See your DCE documentation. If you do not specify `-K`, the `isql` user must be logged in to DCE with the same user name specified in `-U`.

The correct description for the **-K keytab_file** parameter is:

(used only with Kerberos security) specifies a Kerberos keytab file that contains the security key for the user name specified with the `-U` option. To create a keytab, see your Kerberos documentation. If you do not specify the `-K` option, the `isql` user must be logged in to Kerberos with the same user name as specified with the `-U` option.

Job Scheduler Parameter job scheduler interval

The configuration parameter **job scheduler interval** incorrectly describes the value of units as minutes.

The documentation provides the value of units for the configuration parameter as follows:

```
job scheduler interval Default value : 1 (in minutes)
```

```
Range of values : 1 - 600
```

The correct unit value is seconds.

Job Scheduler Command `sp_sjobdrop`

The `sp_sjobdrop` name parameter incorrectly states that the name of a job or a schedule can be used for the name argument. You must use the ID of scheduled job, job, or schedule for the name argument.

Security Built-in Function Permissions

The documentation regarding permissions for several built-in functions is incomplete.

Table 2.

Functions	Permissions
<code>asehostname</code>	<p>With granular permissions enabled, you must be granted select on <code>asehostname</code> or have <code>manage server</code> permission to execute <code>asehostname</code>.</p> <p>With granular permissions disabled, you must be granted select on <code>asehostname</code> or be a user with <code>sa_role</code> to execute <code>asehostname</code>.</p>
<code>migrate_instance_id</code>	Any user can execute <code>migrate_instance_id</code> .
<code>show_cached_text</code>	<p>With granular permissions enabled, you must be a user with <code>mon_role</code>, or have <code>monitor qp performance</code> permission to execute <code>show_cached_text</code>.</p> <p>With granular permissions disabled, you must be a user with <code>mon_role</code> or <code>sa_role</code> to execute <code>show_cached_text</code>.</p>
<code>show_cached_text_long</code>	<p>With granular permissions enabled, you must be a user with <code>mon_role</code>, or have <code>monitor qp performance</code> permission to execute <code>show_cached_text_long</code>.</p> <p>With granular permissions disabled, you must be a user with <code>mon_role</code> or <code>sa_role</code> to execute <code>show_cached_text_long</code>.</p>
<code>show_plan</code>	<p>With granular permissions enabled, you must be a user with <code>monitor qp performance</code> permission to execute <code>show_plan</code>.</p> <p>With granular permissions disabled, you must be a user with <code>sa_role</code> to execute <code>show_plan</code>.</p>
<code>user_id</code>	Any user can execute <code>user_id</code> .

Functions	Permissions
workload_metric	<p>With granular permissions enabled, you must have <code>manage cluster</code> permission or be a user with <code>ha_role</code> to execute workload_metric.</p> <p>With granular permissions disabled, you must be a user with <code>sa_role</code> or <code>ha_role</code> to execute workload_metric.</p>

Shareable Temporary Tables

The Transact-SQL Users Guide incorrectly states that a shared temporary table exists until the current session ends, or until its owner drops it using **drop table**.

While hash temporary tables exist until the current session or scope is exited, shared temporary tables exist until they are explicitly dropped. You can find the incorrect information in *Creating Databases and Tables > Creating Tables > Using temporary tables*.

Hash-Based Update Statistics

The New Features Guide for 15.7 ESD #2 describes the **update statistics** configuration parameter **hashing**, which enables Adaptive Server to gather hash-based statistics but an example showing that three parameters are required for the option is missing.

This is an example of using **update statistics** using the hashing option:

```
sp_configure 'update statistics hashing', 0, 'on'
```

Additional Information for Windows Failover in a High Availability System

This information supplements the Adaptive Server 15.7 document *Using Failover in a High Availability System > Configuring Adaptive Server for Failover on Windows*.

Two Windows systems are required. Microsoft Cluster Server must be installed on both systems and they must be configured together as a Microsoft cluster. Sybase recommends that you set up and test failover of a simple application such as Notepad on the cluster to ensure that the basic failover capabilities of the cluster are working properly. All configuration must be performed from an account that has Domain Administrator privilege.

Configure Adaptive Server for Failover on Windows

Start the failover cluster manager and use the `syconfig` utility to configure the Adaptive Server.

Prerequisites

1. Download and install Microsoft .NET runtime package V4.0
2. The .NET runtime includes a utility called `installutil.exe`. Locate this and copy it to `%SYBASE%\ASE-15_0\install`
3. Run the command `sybcnprn -s` on each cluster node. This installs the “Sybase Companion Server” resource type and Sybase Cluster Administrator extensions.

Task

1. Select **Start > Administrative tools**, then right click **Failover Cluster Manager** and choose **Run as administrator**.

Each instance of an Adaptive Server companion server primary/secondary pair requires a dedicated shared disk that can be failed over between the cluster nodes (you can configure additional shared devices later). The shared disk must be online to the cluster node on which you are configuring Adaptive Server. You might want to create a temporary group and assign the shared disk to this group. When the Adaptive Server companion server setup is complete, it automatically creates its own group, and you can then move the shared disk can then to the new Adaptive Server group.

2. When `syconfig` prompts you to create the various devices, alter these default values:
 - a) Increase the `sysystemprocs` size by 30MB over the default value
 - b) Alter all the path names so that they are stored on the shared-cluster disk
 - c) Alter the name of the `tempdb` device to be unique on each Adaptive Server instance in the cluster

3. Test that the two Adaptive Servers can execute remote commands:

Assuming our servers are called ASE1 and ASE2:

- In ASE1 isql connection, enter: `ASE2...sp_who` ;
- In ASE2 isql connection, enter: `ASE1...sp_who`;

Verify that these commands both succeed. If they do not, verify:

- That the Windows global variables such as `%SYBASE%`, and so on are set.
- The `%SYBASE%\%SYBASE_ASE%\ini\sql.ini` file entries
- You can use `isql` to connect to the ASE1 from m1 and m2. If connecting from m1 fails, then there is a problem with the m1. If connecting from m2 fails, and you find that there is no problem with the m2 environment, try connecting from a third machine. If the connection fails from a third machine, check the firewall settings, as it may prevent you from logging in from outside of m1.

Note: You might want to make backup copies of your master, sybsystemprocs and other devices at this time so you can easily get back to this point.

4. Configure the asymmetric failover companionship.

a) On ASE2 enter:

```
sp_companion ASE1, "configure" , null, sa, <sa pswd>, sa, <sa pswd>
```

As part of the companion configuration, a new group called ASE1_GRP is created. The group contains the Adaptive Server service name which is the same as the server name (ASE1). ASE1 becomes the primary server and ASE2 the secondary.

- b) Move the shared disk containing ASE1's master and other devices into the new group ASE1_GRP.
- c) Inside the new group, ASE1_GRP, right-click the service name ASE1, select **Properties**, and choose the **Dependency** tab. Add a new dependency of ASE1 on the shared disk that has just been moved. This ensures that the cluster brings the disk online before attempting to start Adaptive Server.
- d) Right-click on ASE1_GRP and choose **Properties**. On the General tab, in the Preferred owner subwindow, choose **m1**. For failback m1 must be the location of ASE1.
- e) To set values based on tasks, click the **Maximum failures in the specified period** tab. For example (these values are for reference only; determine the actual values based on your specific needs): 50-100 for test, 10-20 for production.
- f) If you need to perform a failback, choose **Allow failback**; otherwise, choose **Prevent failback**.

To configure symmetric failover companionship, perform the same set of above steps on server ASE1, substituting ASE2 for ASE1 in the various commands.

5. To test the asymmetric HA companionship failover and failback:

- a) On ASE1, issue **shutdown with nowait**
- b) On m1 taskmgr, issue **kill ASE1 process**
- c) On m1 service, issue **stop ASE1**
- d) On the Failover Cluster Manager, ASE1_GRP – ASE1 offline
- e) On the Failover Cluster Manager, ASE1_GRP – ASE1 simulating failure on the ASE
- f) On the Failover Cluster Manager, ASE1_GRP – disk1 simulating failure on the disk
- g) Shut down machine m1
- h) Turn off the power to machine m1

The above actions should result in a ASE1 shared disk failover to ASE2. On ASE2 will show the ASE1 databases and devices. ASE1 client connections with -Q failover property should also failover to ASE2.

Note: Asymmetrically shutting down of ASE2 does not cause a failover because the group ASE2 has not been created in this case.

6. HA test for a completed failback:

a) On ASE2, enter:

```
sp_companion ASE1 , "prepare_failback"
```

b) On m1 service start ASE1

c) On the Failover Cluster Manager: ASE1_GRP | ASE1: bring online

d) Turn on the power to machine m1.

To test the symmetric companionship, you can perform these steps on both machines, or only on ASE1.

Also see "CR 692496, Configuring tempdb device for HA" in the Known Installation Issues for Adaptive Server.

The following `syconfig` utility fields must be set correctly.

- Connection protocol – TCP
- Connection value–hostname,portnumber
The host and port number must separated by a comma.
- If `syconfig` is not started in the ASE configuration file location, the parameter field must specify `-c <ase>.cfg`, where `<ase>.cfg` is the full path of the file.

Setting maximum failed logins

The documentation provides incorrect examples for the maximum number of failed logins.

- Use this example:

```
create role intern_role with passwd "temp244", max failed_logins  
20
```

Disregard this example:

```
create role intern_role with passwd "temp244", maximum failed  
logins 20
```

Changing the maximum failed logins for specific roles

The documentation provides incorrect examples for removing overrides or changing the maximum number of failed logins for a role.

The correct example for removing the overrides for the maximum number of failed logins for all roles is:

```
alter role "all overrides" set max failed_logins -1
```

The correct example for changing the maximum failed logins allowed for “physician_role” to 5 is:

```
alter role physician_role set max failed_logins 5
```


Changing the maximum failed logins for specific logins

The documentation provides an incorrect example for changing the maximum number of failed logins for a login.

The correct example is:

```
alter login joe modify max failed attempts 40
```

Values for lock timeout pipe active

The default value and the range of values for **lock timeout pipe active**, as described in the System Administration Guide are incorrect.

The correct values are:

Default value 0 (off) Range of values 1 (on), 0 (off)

dbcc page Parameter logical

The documentation for the **dbcc page** command incorrectly describes the **logical** parameter.

The correct description for the **logical** parameter is:

If `cache = 1`, then value of this parameter is considered as 1 and value passed for **pageno** parameter is treated as logical page number.

If `cache = 0`, then value of this parameter is taken as the virtual device number and value passed for **pageno** parameter will be treated as virtual page number. The virtual device number can be obtained from the `master..sysusages` table.

Behavior of Concurrent DDLs and reorg defrag

This information supplements the Adaptive Server 15.7 SP100 documentation for Incremental Reorganization and describes how concurrent DDLs behave when **reorg defrag** is in progress and how **reorg defrag** behaves during concurrent DDLs.

- When **reorg defrag** is in progress on a table, invocation of any schema changing utility or data reorganization utility fails with error 11051. In rare cases of race condition between **reorg defrag** and the other utility, one of them might have to wait for the other to complete.
- When **reorg defrag** utility is in progress on a table, another invocation of **reorg defrag** on the same table will fail with error 8233.

Obtaining Help and Additional Information

- When a schema changing utility or data reorganization utility is in progress, invocation of **reorg defrag** would wait for the utility to complete before proceeding to defragment the data.

Updates to Third-Party Licensing

Check the Free Download Terms document for recent updates to third-party licensing details.

The document is at: http://www.sybase.com/softwarelicenses/third_party_legal.

Obtaining Help and Additional Information

Use the Product Documentation site and online help to learn more about this product release.

- Product Documentation at <http://sybooks.sybase.com/> – online documentation that you can access using a standard Web browser. You can browse documents online, or download them as PDFs. The Web site also has links to other resources, such as white papers, community forums, maintenance releases, and support content.
- Online help in the product, if available.

To read or print PDF documents, you need Adobe Acrobat Reader, which is available as a free download from the *Adobe* Web site.

Note: A more recent release bulletin, with critical product or document information added after the product release, may be available from the Product Documentation Web site.

Technical Support

Get support for SAP® products.

If your organization has purchased a support contract for this product, then one or more of your colleagues is designated as an authorized support contact. If you have any questions, or if you need assistance during the installation process, ask a designated person to contact Technical Support as specified in your contract:

- Sybase Technical Support or the Sybase subsidiary in your area
- SAP Technical Support

Customers with an SAP support contract can obtain support for this product at the SAP support site, <http://service.sap.com/sybase/support>. You can also find information about Sybase support migration at this location (login may be required).

Customers with a Sybase support contract can obtain support for this product at <http://www.sybase.com/support> (login required).

Downloading Product Updates

Get maintenance releases, support packages and patches, and related information.

- If you purchased the product directly from Sybase or from an authorized Sybase reseller:

- a) Go to <http://www.sybase.com/support>.
- b) Select **Support > EBFs/Maintenance**.
- c) If prompted, enter your MySybase user name and password.
- d) (Optional) Select a filter, a time frame, or both, and click **Go**.
- e) Select a product.

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

- f) Click the **Info** icon to display the EBF/Maintenance report, or click the product description to download the software.
- If you ordered your product under an SAP contract:
 - a) Go to <http://service.sap.com/swdc> and log in if prompted.
 - b) Select **Search for Software Downloads** and enter the name of your product. Click **Search**.

Product and Component Certifications

Certification reports verify Sybase product performance on a particular platform.

To find the latest information about certifications:

- For partner product certifications, go to http://www.sybase.com/detail_list?id=9784
- For platform certifications, go to <http://certification.sybase.com/ucr/search.do>

Accessibility Features

Accessibility ensures access to electronic information for all users, including those with disabilities.

Documentation for this product is available in an HTML version that is designed for accessibility.

Obtaining Help and Additional Information

Vision impaired users can navigate through the online document with an adaptive technology such as a screen reader, or view it with a screen enlarger.

HTML documentation has been tested for compliance with accessibility requirements of Section 508 of the U.S Rehabilitation Act. 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 may 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 the Sybase Accessibility site: <http://www.sybase.com/products/accessibility>. The site includes links to information about Section 508 and W3C standards.

You may find additional information about accessibility features in the product documentation.