

Release Bulletin Sybase® IQ 12.7 for HP-UX

Document ID: DC74720-01-1270-07

Last revised: January 2010

Topic	Page
1. Accessing current release bulletin information	2
2. Product summary	2
3. Special installation and migration instructions	11
4. New features and behavior changes	26
5. Product compatibilities	87
6. Restrictions	88
7. Known problems	94
8. Documentation for this version	98
9. Documentation updates and clarifications	100
10. Technical Support	175
11. Other sources of information	177
12. Accessibility features	181
13. Checklist: information for Technical Support	183

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 product release, use the Sybase Product Manuals Web site.

❖ Accessing release bulletins at the Sybase Product Manuals Web site

- 1 Go to Product Manuals at <http://www.sybase.com/support/manuals/>.
- 2 Select a product and language and click Go.
- 3 Select a product version from the Document Set list.
- 4 Select the Release Bulletins link.
- 5 From the list of individual documents, select the link to the release bulletin for your platform. You can either download the PDF version or browse the document online.

2. Product summary

Enclosed is Sybase® IQ 12.7 for 64-bit HP-UX with multiplex capability. This product is compatible with the following platform and operating system configurations:

- HP 9000 Series 800 HP-UX 11.23
- HP 9000 Series 800 HP-UX 11.31
- HP Itanium HP-UX 11.23
- HP Itanium HP-UX 11.31

If you purchased the Sybase IQ Extended Enterprise Edition, the following products are also enclosed:

- Sybase IQ ETL Server
- Sybase IQ ETL Development – runs only on Windows systems, on the following platforms:
 - Microsoft Windows 2003 (x86) 32-bit
 - Microsoft Windows 2003 (x64) certified with Service Pack 1 (32-bit version on 64-bit platform)

- Microsoft Windows XP Professional (x86) certified with Service Pack 2 (32-bit version on 64-bit platform)

Note Sybase IQ is also referred to as Adaptive Server IQ. You will sometimes see the name Adaptive Server IQ in documentation and in the product.

Refer to the *New Features in Sybase IQ 12.7* document for descriptions of new features and behavior changes in Sybase IQ 12.7.

2.1 Installation kit

The Developer's and Enterprise Edition installation kits include the following CDs:

- Getting Started – Installation Guide, Release Bulletin.
- Sybooks™ – Core documentation. (See “Documentation for this version” on page 98.)
- Network Client – IQ Network Client (Windows 32, Windows Vista™. For Windows Vista, download Sybase IQ 12.7 ESD #3 and higher releases.)
- Developer's Kit or Enterprise Edition – IQ Server and related products. (See *Sybase IQ Installation and Configuration Guide* for details.)

The Extended Enterprise Edition installation kit includes the following CDs:

- Getting Started Extended Enterprise Edition – Installation Guide, Release Bulletin, User's Guide for Sybase IQ ETL.
- Sybooks – Core documentation. (See “Documentation for this version” on page 98.)
- Network Client on Linux – IQ Network Client (Linux 32).
- Extended Enterprise Edition – IQ Server and related products, IQ ETL Server. (See *Sybase IQ Installation and Configuration Guide* for details.)
- Extended Enterprise Edition Client – IQ Network Client, IQ ETL Development.

Only IQ Network Client on Windows 12.7 ESD #3 and higher releases support Windows Vista™. For details, see “Installing the Network Client on Windows Vista” on page 18.

2.2 Updated software components

Sybase IQ Server 12.7 installs the following software components, updated to the latest versions available at release time:

- Sybase jConnect JDBC Driver version 5.5. This component is optional but strongly recommended. If you install the jConnect driver, Java classes installed into a database can make JDBC calls to access and modify data.

Sybase jConnect JDBC Driver version 6.0 is also included with this release and is optional. For more information, see the *Sybase IQ Installation and Configuration Guide for HP-UX*.

- Java Runtime Environment (JRE) version 1.4.2 allows you to use the latest Sybase Central plug-in for Sybase IQ.
- Open Client Software Developer's Kit version 15.0

IMPORTANT!

Sybase strongly recommends that you check the online support Web site for software updates to these components before you install the software. If a software update (ESD or EBF) has been released, it contains bug fixes made after this product shipped. You must download the latest update and install it after you install Sybase IQ from the product CD. Read the cover letter included with the update for workarounds if you encounter problems not described in this Release Bulletin. For more information, see "Sybase EBFs and software maintenance" on page 179.

2.3 Configuring ETL client/server connectivity

For an overview of client/server connectivity for Sybase IQ, see "Configuring client connectivity," in Chapter 5, "Configuring Sybase IQ," in the *Sybase IQ Installation and Configuration Guide*. You must define Open Client server definitions and ODBC data source names (DSNs) for ETL projects or jobs on the appropriate machine. Server definitions and DSNs must be appropriate for the platform and operating system where the ETL project or job runs.

For example, if you build a project or job in the Windows ETL Development client that you plan to run on a remote GridNode engine on Solaris, you need to configure the connectivity on that remote machine with the GridNode. The ODBC DSN or Open Client server name must be defined in the Solaris *.odbc.ini* file or interfaces file.

❖ Configuring connectivity for ETL on Windows

Follow these steps to configure connectivity for Sybase IQ ETL Server on Windows.

- 1 Install Sybase IQ ETL Development from the Extended Enterprise Edition Client CD.
- 2 Open the ODBC Data Source Administrator (Start > Programs > Sybase > Data Access > ODBC Data Source Administrator.)
- 3 In the ODBC Data Source Administrator application, choose the System DSN tab.
- 4 Click Add and select the Adaptive Server IQ driver when prompted.
- 5 Name the data source (DSN) and complete the connection information.
See “Creating ODBC data sources,” in Chapter 5, “Configuring Sybase IQ,” in the *Sybase IQ Installation and Configuration Guide*.
- 6 Within the ETL client, create a project and add a DB component. In the component properties, select ODBC from the Interface drop-down menu and then specify the ODBC DSN in the Host name field, user id and password to connect.

❖ Configuring Open Client on Windows

Follow these steps to configure IQ ETL to be an Open Client Sybase IQ client on Windows.

- 1 Install the Sybase IQ Server from the Extended Enterprise Edition.
- 2 Install Sybase IQ SDK (Open Client).
- 3 Create a *sql.ini* file entry for Sybase IQ using the dsedit utility (Start > Sybase > Connectivity > Open Client Directory Services Editor).
For more details, see “Connecting using Open Client,” in Chapter 5, “Configuring Sybase IQ,” in the *Sybase IQ Installation and Configuration Guide*.
- 4 Click Add to define a server, name it, and provide the connectivity information.
- 5 Use the Ping utility to verify it.

- 6 Within the ETL client, create a project and add a DB component. In the component properties, select Sybase from the Interface drop-down menu and then specify the Open Client server name defined in the *sql.ini* file in the Host name field, user id and password to connect.

Note The Sybase IQ 12.7 Windows Client does not include Sybase native (Open Client) connectivity. In order to use Sybase native connectivity from Sybase IQ ETL to Sybase IQ, you must install the native libraries from another Sybase product. For example, the native libraries are available in the Sybase Software Development Kit (SDK), which is an installation option of the 12.7 Sybase IQ Server on Windows or as a separate product. Contact your Sybase Technical Support representative for information about the best way to obtain Sybase native connectivity for your site.

ETL products for Windows are 32-bit applications but can run on the Windows 64-bit platform. If the 64-bit Sybase IQ Server was installed on the same machine and the installation included the SDK, the 32-bit libraries would also be installed and could be used.

❖ **Configuring connectivity for ETL on UNIX or Linux**

Follow these steps to configure connectivity for Sybase IQ ETL Server on UNIX or Linux.

- 1 Choose the Sybase IQ Server option from the Extended Enterprise Edition installation.
- 2 Install Sybase IQ Server.
- 3 Create a *.odbc.ini* file with any text editor and define a Data Source Name there. For details, see “Using ODBC data sources on UNIX,” in the *Sybase IQ System Administration Guide*.
- 4 Set the \$ODBCINI variable to point to the file by typing the following C shell command:

```
setenv ODBCINI  
/<Sybase_iq_installation_directory>/odbc.ini
```

For other shells, use the sh /export syntax.

- 5 Run *ASIQ-12_7.csh* or *ASIQ-12_7.sh* script from the root installation directory to set the variables.

- 6 Verify the ODBC configuration using the Sybase IQ component Interactive SQL. Change to the *\$ASDIR/bin* directory and type:

```
./dbisql -c "uid=<user_id>;pwd=<password>"
-ODBC -dataSource
<ODBC_data_source_name_from_.odbc.ini> -nogui.
```

On connection, a user name displays. Type `select @@version` to return the Sybase IQ versions string to verify the connection. If an error indicates that no driver and no valid data source are found, the `$ODBCINI` variable is not set properly. Check the `$ODBCINI` variable setting and the *.odbc.ini* file for any possible mistakes or omissions. Type `exit` to exit `dbisql`.

For more information, see “Storing connection information,” in Chapter 5, “Configuring Sybase IQ,” in the *Sybase IQ Installation and Configuration Guide*.

- 7 In the same console session where the `ASIQ` and `$ODBCINI` variables are set, change to the ETL server installation directory and run the *GridNode.sh* script.
- 8 Within the ETL client, create a project and add a DB component. In the component properties, select ODBC from the Interface drop-down menu and then specify the ODBC DSN from the *.odbc.ini* file in the Host name field to connect.

❖ **Configuring Open Client**

Follow these steps to configure the ETL server to be an Open Client to Sybase IQ on UNIX or Linux.

- 1 Choose the Sybase IQ Server option from the Sybase IQ Extended Edition Server installation media.
- 2 Choose to install Open Client from the Sybase IQ Server installation.
- 3 Run the sybase environment variable script in the root `$SYBASE` installation directory. For example, for the C shell, type: `./source SYBASE.csh`
- 4 Use `dsedit` to create an interfaces file in the root `$SYBASE` directory. Change to the `$SYBASE/OCS-15_0/bin` directory and type: `./dsedit`

Since `dsedit` is a GUI application, you may need to set the `$DISPLAY` variable to the machine where you want `dsedit` to display. For example, in the C shell, type: `setenv DISPLAY mypc:0.0`
- 5 Specify the interfaces file. Use the default file in the root `$SYBASE` installation directory or create your own. Click OK.

- 6 When the Directory Service Session panel appears, click the Add new server entry button to define a new server.
- 7 In the Server Entry Editor panel, enter the server and click the Add new network transport button.
- 8 In the Network Transport Editor panel, provide the following information:
 - Transport type: Select tcp from the drop-down menu.
 - Host name: Enter the machine name of the IQ host.
 - Port number: Enter the port on which the IQ server is running. For example, the default port for the Sybase IQ demo server is 2638. in the interfaces file.

Click OK.

- 9 Click Cancel as necessary and Exit to close the dsedit.
- 10 In the same console with the Sybase variables set, and within the `$SYBASE/OCS-15_0/bin` directory, use the `isql` application to verify the Open Client server definition. Type the following command:

```
./isql -U<user_name> -P<password>  
-S<server_name_defined_in_interfaces>
```

For example, for the Sybase IQ demo server, use: `./isql -Udba -Psql -Sasiqdemo`. The application connects and displays a 1> prompt.

- 11 To return the Sybase IQ version string for verification, type: `select @@version`, press return and then enter `go`. If the connection fails, use any messages to troubleshoot the problem and verify the dsedit entry
- 12 Type `exit` at the command prompt to close `dbisql` and the connection.
- 13 In the same session where the `$SYBASE` variables are set, change to the ETL server installation directory and run the `configure.sh` file to update the ETL environment with the Sybase environment information.
- 14 Run the `GridNode.sh` file.
- 15 Within the ETL client, create a project and add a DB component. In the component properties, select Sybase from the Interface drop-down menu and then specify the Open Client server name defined in the `sql.ini` file in the Host name field to connect.

2.4 ODBC and database drivers for Sybase IQ ETL

The Sybase IQ ETL environment has been tested, evaluated, and verified thoroughly to comply with many interface drivers of the supported database systems. While most drivers used in a current Windows environment exhibit no problems, some drivers might cause unexpected results when dealing with Unicode character sets.

If you encounter unexpected results that might be related to driver incompatibility, install one of the following versions for your interface driver.

Table 1: Interface driver versions for Sybase IQ ETL Server

Driver	Version
Sybase native (via ct-lib)	15.0
Sybase ASE ODBC	15.00.00.50 (Windows only)
Sybase IQ ODBC	9.00.02.1023
DB/2 native	9
DB/2 ODBC	9.01.00.356 (Windows only)
Microsoft Access ODBC	12.00.4518.1014 (Windows only)
Microsoft SQL Server ODBC	2005.90.3042.00 (Windows only)
Oracle native (via oci)	10.2.0.3
Oracle ODBC	10.2.0.3.0 (Windows only)

2.5 Network Client supported platforms

The Sybase IQ Network Client is available for Windows and Linux. Both the Windows and the Linux versions of the Sybase IQ Network Client operate with the Sybase IQ server on any supported platform.

The Sybase IQ Network Client for Windows is included with Sybase IQ 12.7 on all platforms. The Sybase IQ Network Client for Linux can be downloaded or ordered separately.

- **Linux** – The Sybase IQ Network Client for Linux can be installed on the following configurations:
 - Red Hat Enterprise Linux 3.0 IA32, Advanced Server or Workstation Edition, distribution of kernel version 2.4.21-4.0.1.ELhugemem #1 SMP and glibc version 2.3.2-95.39, or
 - Red Hat Enterprise Linux 4.0, Advanced Server or Workstation Edition, distribution of kernel version 2.6.9-22.0.2.ELsmp #1 SMP and glibc version 2.3.4-2.13, or

- Red Hat Enterprise Linux 5.0, Advanced Server or Workstation Edition, or
- SuSE Linux Enterprise Server 8.0, or
- SuSE Linux Enterprise Server 9.0, with kernel version 2.6.5-7.244-pseries64 #1 SMP and glibc version 2.3.3.-98.61, or
- SuSE Linux Enterprise Server (SLES) 10.0, with kernel version 2.6.16.21-0.25-ppc64 #1 SMP and glibc version 2.4-31.5.
- **Windows** – The Sybase IQ Network Client for Windows can be installed on a Windows 2000 SP2, Windows 2003, Windows XP SP2, or Windows Vista system.

Sybase IQ Network Client for Windows on 32-bit platforms includes both 32-bit and 64-bit ODBC Windows drivers. Windows 64-bit ODBC drivers, provided as of ESD #4, allow you to connect to Sybase IQ from third-party packaged 64-bit Windows applications and 64-bit Windows applications written in C++.

The Sybase IQ installation for UNIX and Linux platforms includes ODBC drivers needed for the client and DBISQL.

2.6 Required HP-UX operating system patches

Sybase recommends that you keep your system current with the latest patch level. Refer to the Hewlett-Packard Company Web site for patch information.

When installing patches, log in as “root” (superuser) and follow the directions provided with the patch. All tasks for adjusting the operating system configuration must be performed as “root”. Failure to install the necessary patch bundle and patches before starting Sybase IQ will require reinstalling Sybase IQ.

The following minimum operating system patch bundles and patches are required to run Sybase IQ 12.7 on HP-UX

- HP-UX 11.23 – patch level B.11.23.0512.034
- HP-UXi 11.23 – patch level B.11.22.0306.4

2.6.1 Listing version and patch information on HP-UX

Displaying the level of your base system software	The following commands return information about your operating system. The uname -sr operating system command shows the level of your base system software: <pre>% uname -sr</pre>
	The command returns the following result: <pre>HP-UX B.11.23</pre>
Listing individual patches	To list individual patches installed on your system, log in as “root” superuser and issue the following operating system command: <pre>% /usr/sbin/swlist -l product grep PH</pre>
Listing the patch bundle	To list the currently installed patch bundle with month, year, and version number, issue the following operating system command: <pre>% /usr/sbin/swlist -l bundle</pre>

3. Special installation and migration instructions

This section includes installation and migration information that is new for this release or needs emphasis. For complete installation and migration instructions, see the *Sybase IQ Installation and Configuration Guide for HP-UX*. Be sure to read the section “Before you install” before installing Sybase IQ 12.7.

Sybase IQ 12.6 and higher enforces check constraints. If you have existing check constraints, see the section “Before you install” in the chapter “Installing Sybase IQ” in the *Sybase IQ Installation and Configuration Guide* before installing.

“Before you install” also contains important information on upgrading LONG BINARY columns.

IMPORTANT!

You must upgrade all existing databases after you install *any* Sybase ESD. To upgrade, run ALTER DATABASE UPGRADE. See the *Sybase IQ Installation and Configuration Guide* for details.

Before you run a new version of Sybase IQ, see “Restrictions” on page 88 for the most current requirements.

Sybase strongly recommends that you check the online support Web site for software updates to these components before you install the software. If a software update (ESD or EBF) has been released, it contains bug fixes made after this product shipped. Read the cover letter included with the update for workarounds if you encounter problems not described in this Release Bulletin. You must download the latest update and install it after installing IQ from the product CD. See “Sybase EBFs and software maintenance” on page 179.

3.1 Installing Sybase IQ server updates

For server updates, Sybase IQ 12.7 ESD #8 must be installed (or have been installed) before Sybase IQ 12.7 ESD #9 or later ESDs are installed.

3.2 Windows 64-bit ODBC drivers added to 32-bit Network Client on Windows [CR 492213]

Sybase IQ Network Client for Windows now includes 64-bit ODBC drivers and libraries when installed on a 64-bit Windows machine. Windows 64-bit ODBC drivers, provided as of ESD #4, allow you to connect to Sybase IQ from third-party packaged 64-bit Windows applications and 64-bit Windows applications written in C or C++.

3.3 Windows Network Client installs silently [CR 468594 and CR 491779]

Sybase IQ Network Client for Windows now installs silently on all Windows platforms, including Windows Vista. (For an overview of silent installation, see “Installing without user interaction,” in Chapter 3, “Installing Client Components on Windows,” in the *Sybase IQ Installation and Configuration Guide*.) A new GA refresh is available on the Sybase download site at <http://downloads.sybase.com>. Clicking *win32-iq127_client_refresh.exe* in the browser returns the prompt Run or Save As File. Choose Save As File, then run *win32-iq127_client_refresh.exe* from the command prompt to install the refresh, as described in the following procedures. To use the silent install on Windows Vista, you must install this new version of the GA client and then install ESD #4.

For a silent install, most users create a response file on a single source system and run the silent install on multiple target systems. If you need use the same machine for both source and target, uninstall the Sybase IQ Network Client after creating the response file and reboot before running the silent install.

❖ Creating response files

Response files provide the answers to user prompts from the installation procedure.

- 1 Uninstall any previous version of Sybase IQ.
- 2 On the source system, create a response file for the GA refresh using the following syntax:

```
win32-iq127_client_refresh /R  
/F1"<DirectoryName>\<FileName>.iss"
```

For example:

```
win32-iq127_client_refresh /R /F1"C:\ga_setup.iss"
```

The F1 parameter must contain the full path.

- 3 Close the current window and, in a new window, create the response file for the ESD using the following command:

```
win32-iq127_esd4_client /R /F1"C:\esd_setup.iss"
```

- 4 Close the current window.

You now have two response files for use by the silent install: *ga_setup.iss* and *esd_setup.iss*.

❖ Installing Sybase IQ Network Client silently

Before you begin, make sure that you have created response files (see previous procedure).

- 1 Log on to the target system and uninstall any previous version of Sybase IQ.
- 2 Copy the response files from the source system to the root directory (C:\) on the target system.
- 3 Open a new window and run the GA refresh in silent mode as follows:

```
Win32-iq127_client_refresh /S /F1"C:\ga_setup.iss"  
-save_log
```

The optional *-save_log* parameter creates a log for troubleshooting purposes in *%TEMP%\Sybase_IQ.install.log*.

- 4 Please wait patiently for the install to complete; it should take no more than two minutes. You may monitor its progress on the task manager or installation log, but do not perform other tasks while the install runs in the background. After the silent install completes, close the current windows.

- 5 Open a new window and run the ESD install in silent mode as follows:

```
Win32-iq127_esd4_client /S /F1"C:\esd_setup.iss"  
-save_log
```

- 6 After the silent install completes, close the current windows.
- 7 Reboot the target system.

3.4 Running iq_bcp with the Sybase IQ Network Client on Windows [CR 488573]

The iq_bcp utility, included with Sybase IQ as of 12.7 ESD #1, requires the Sybase utility bcp to be installed and in the path. The bcp utility is part of the Sybase Software Development Kit (SDK) and Adaptive Server Enterprise, which are separately licensed products not included with Sybase IQ Network Client. To purchase either product, contact your Sybase sales representative.

3.5 Sybase IQ certified on HP-UX 11.31 [CR 476205]

As of Sybase IQ 12.7 ESD #3, the following section updates “Adjusting kernel parameters” in Chapter 1, “Installing Sybase IQ,” in the *Sybase IQ Installation and Configuration Guide*.

Adjusting kernel parameters

You also need to adjust certain HP-UX kernel parameters to run Sybase IQ, as follows:

- For HP-UX 11.31, invoke the HP System Management Homepage (SMH) to verify and set parameters.
- For HP-UX 11.23, the HP System Administration Manager (SAM) allows verification and setting of these parameters. Generally superuser or “root” privileges are needed to run SAM. Use the Kernel Configuration submenu to change the parameters listed in the following table.

Table 2: HP-UX parameters

Parameter	HP-UX 11.31 suggested value	HP-UX 11.23 Itanium suggested value	HP-UX 11.23 RISC suggested value
STRMSGSZ	0	0	0
bufpages	n/a	n/a	n/a
max_thread_proc	4096	4096	4096
maxfiles	2048	2048	2048
maxfiles_lim	2060	2060	2060
maxswapchunks	n/a	n/a	n/a
maxusers	n/a	n/a	n/a
nfile	n/a	20000	20000
nkthread	8416	8416	4096
sema	n/a	n/a	n/a
semmap	n/a	n/a	n/a
semمني	7082	7082	7082
semممنس	14164	14164	14164
shmmax	2147483647 (0X7FFFFFFF)	2147483647 (0X7FFFFFFF)	2147483647 (0X7FFFFFFF)
swapmem_on	0	0	0
unlockable_mem	n/a	10240	10240
maxdsiz_64bit	17179869184 (0x400000000)	17179869184 (0x400000000)	17179869184 (0x400000000)
maxrsessiz_64bit	1073741824	1073741824	n/a
maxssiz_64bit	1073741824	1073741824	268435456 (0x10000000)
maxtsiz_64bit	4294967296 (0x100000000)	4294967296 (0x100000000)	4294967296 (0x100000000)

After setting these parameters as required, reboot your system.

Notes

- The swapmem_on parameter is set to 0 because setting it to 1 increases paging and deactivation activity.
- The maxswapchunks parameter is increased to raise the limit of configurable swap space beyond the default.
- For HP-UX 11.23 Itanium, set the nkthread value greater than (nproc + 100).

3.6 Files locked during installation [CR 465663]

On some UNIX systems (particularly AIX), files that the installer tries to update could be held in memory and locked. This problem usually resolves itself in subsequent runs of the installer. However, if the locks continue to be held, run `sybinstall` with the `-unlock` parameter added in 12.7 ESD #3. This parameter installs the files in a shadow directory, links to the correct location, then removes the shadow directory, thereby avoiding lock problems.

3.7 Back up database and log files [CR 465000, CR 424572]

As a backup, copy the `.db` and `.log` files for the database immediately before running `ALTER DATABASE UPGRADE`. Perform this backup before starting the server. Because `ALTER DATABASE UPGRADE` modifies only the catalog, a full backup is unnecessary.

3.8 Windows Vista support issues

For Sybase IQ 12.7 ESD #3 and later releases, the IQ Network Client is certified on the Windows Vista operating system. Following are some issues relating to running Sybase IQ software on Vista:

- Windows Vista security

Windows Vista incorporates a new security model. User Account Control (UAC) is enabled by default and may affect the behavior of programs that expect to be able to write files, especially when the computer supports more than one user. Depending on where and how files and directories are created, a file created by one user may have permissions that do not allow another user to read or write to that file. If you install Sybase IQ in the default directories, files and directories that require read/write access for multiple users are set up appropriately.

- Sybase IQ elevated operations agent

In Vista, certain actions require privilege elevation to execute when run under User Account Control. The programs `dbelevate9.exe` and `iqdsn.exe` may require elevation in Sybase IQ.

The following dll files require elevation when they are registered or unregistered: `dbodbc9.dll` and `dboledb9.dll`.

On a Vista system with User Account Control activated, you may receive an elevation prompt for the Sybase IQ elevated operations agent. The prompt is issued by the Vista User Account Control system to confirm that you want to continue running the identified program (if logged on as an administrator) or to provide administrator credentials (if logged on as a non-administrator).

- Deployment changes

The program *dbelevat9.exe* is used internally by Sybase IQ components to perform operations that require elevated privileges. This executable must be included in deployments of Sybase IQ. Administrator privileges are required to run *dbelevat9.exe*.

- Sybase IQ executables signed

Sybase IQ executables on Vista are signed by Sybase, Inc.

- Using an AWE cache

To use an AWE (Address Windowing Extensions) cache on Windows Vista, you must run the database server as administrator. Starting a non-elevated database server with an AWE cache results in a warning that the database server must be run as an administrator to use AWE. See “-cw server option” in Chapter 1, “Running the Database Server,” in the *Sybase IQ Utility Guide*.

- Samples

Samples now correctly handle Sybase IQ installation path names that contain one or more spaces.

- Windows services

Vista-compliant services are not allowed to interact with the desktop. On Windows Vista, no Sybase IQ services interact with the desktop (even if Allow Interaction with Desktop is enabled in the service definition). Sybase IQ database servers can be monitored using the *dbconsole* utility or from Sybase Central. Sybase Central disables the option to allow service to interact with desktop when running on Windows Vista.

3.9 Adding ODBC Data Source Names on 64-bit Windows systems [CR 480906]

On 64-bit Windows systems, Sybase IQ installs a 32-bit ODBC driver. To add new Data Source Names, invoke the ODBC Administrator one of two ways:

- Run Start > Programs > Sybase > Data Access > ODBC Data Source Administrator
- Run `C:\WINDOWS\SysWOW64\odbcad32.exe`

Other methods of launching ODBC Administrator from the Control Panel or Run box do not return desired results. If you type “odbcad32” instead of the full path in the Run text box, ODBC lists only one driver (SQL Server driver) and prevents adding Data Source Names. The Version or Company column in the driver list may include the value “(Not Marked).”

Attempts to add a DSN starting the ODBC Administrator incorrectly may return one of the following errors:

- The setup routines for the Adaptive Server IQ ODBC driver could not be loaded due to system error code 126.
- Driver's ConfigDSN, ConfigDriver, or ConfigTranslator failed: Could not load the setup or translator library.
- The setup routines for the Adaptive Server IQ ODBC driver could not be found. Please reinstall the driver.

3.10 Installing the Network Client on Windows Vista

Windows Vista's multiple security levels affect Sybase IQ Network Client installation and operation in several areas:

- Sybase IQ digitally signs all `.exe` and `.dll` files. Digital signing guarantees that the executable was built by the signing company and has not been modified.
- Only users with administrator privileges can write to the Program Files directory. If another user attempts to write to Program Files, the write is redirected (without the user's knowledge) or blocked (returning an error).
- Sybase IQ files that need to be modified by programs are moved to new user writable directories:
 - User public directory (%PUBLIC%)
`<MainDisk>:\Users\Public\Documents\Sybase IQ 12.7`

- Program data directory
<MainDisk>:\ProgramData (%ALLUSERSPUBLIC%)

On the IQ Network Client on the Windows Vista platform, the sample files are relocated from <MainDisk>:\Program Files\SYBASE\ASIQ-12_7\Samples to <MainDisk>:\Users\Public\Documents\Sybase IQ 12.7\Samples.

See “Before you install Sybase IQ Extended Enterprise Edition on Windows Vista” if installing Extended Enterprise Edition.

3.11 Before you install Sybase IQ Extended Enterprise Edition on Windows Vista

If you changed any of the default Vista settings, installing the GA version of Sybase IQ Extended Enterprise Edition may require a lower security setting.

❖ Setting signed executables option

The following procedure does not apply to the Enterprise Edition or Developer's Kit.

- 1 Select Control Panel > Administrative Tools > Local Security Policy > Local Policies > Security Options.
- 2 Select User Account Control: Only elevate executables that are signed and validated.
- 3 If enabled, set to Disabled until Sybase IQ installation completes.

3.12 Upgrading write and query servers [CR 474126]

The order in which you upgrade multiplex servers depends on the versions of Sybase IQ on the servers.

If you plan to run a mixed-version multiplex, you must upgrade query servers before the write server as described in the section “Running a mixed-version multiplex,” in Chapter 4, “Migrating data,” in the *Sybase IQ Installation and Configuration Guide*.

If you plan to run one version across the multiplex, you may upgrade servers in any order. You must shut down all multiplex servers, and after the write server is upgraded, synchronize query servers. See the section “Upgrading servers from 12.6 to 12.7” in the same chapter.

3.13 Reboot to synchronize message and SQL log files [CR 467129]

If you run Sybase IQ in a country that observes Daylight Savings Time, you must reboot IQ servers after the change back to standard time. The reboot corrects a time difference between the IQ message log and the SQL log generated by specifying the `-zr` switch on the server startup. IQ servers also require the reboot after the change to daylight savings time.

3.14 Correction to “Installing ODBC Drivers” [CR 446798]

The following paragraph appeared erroneously in the section “Installing ODBC Drivers” in Chapter 5, “Configuring Sybase IQ” of the *Sybase IQ Installation and Configuration Guide*.

If you are using ODBC with UNIX or Linux, see “Using ODBC without the driver manager” in Chapter 4 of the *Adaptive Server Anywhere Programming Interfaces Guide* to ensure that you are using the correct driver.

The *Programming Interfaces Guide* does not exist in the Adaptive Server Anywhere documentation set. The relevant section is now “Linking ODBC applications on UNIX” in Chapter 7 of the *Adaptive Server Anywhere Programming Guide*. This section is cross-referenced earlier in the same section of the *Sybase IQ Installation and Configuration Guide*. Another reference is unnecessary.

3.15 Migrating across hardware platforms [CR 445754]

This section was omitted from Chapter 4, “Migrating Data from Previous Versions,” in the *Sybase IQ Installation and Configuration Guide*.

Sybase IQ supports migrating your database from one platform to another, as long as both have the same endian structure.

Platforms with big-endian structure are:

- AIX64
- HP-UX64 PA-RISC
- HP-UX64 Itanium
- IBM Linux on POWER
- SunOS64

Platforms with little-endian structure are:

- Linux32**
- Linux64
- Windows 32
- WinAMD64
- SunAMD64

Sybase IQ 12.6 ESD #2 and higher releases support migration between Windows and Linux.

IMPORTANT!

** If you created your Sybase IQ database on a Linux 32-bit version prior to Sybase IQ 12.6 ESD #2, you must first install IQ 12.6 ESD #2 for Linux 32-bit systems and create a new data backup before migrating to another platform.

❖ **Migrating a database from one platform to another (same endian structure)**

- 1 Back up the database.
- 2 Shut down the Sybase IQ server.
- 3 Install the Sybase IQ server on the new platform. Your migration can take place on the same or a different machine.
- 4 Start the Sybase IQ server on the new hardware platform.
- 5 Connect to the utility database, *utility_db*.
- 6 Restore the database from the backup you created in Step 1.
- 7 Shut down the server and restart it against the restored database. If the current version of Sybase IQ is higher than the version on which you were previously running, you need to upgrade databases, and therefore restart the server in a way that restricts user connections. Sybase recommends using two server startup options:
 - Use -gd DBA so that only users with DBA authority can start and stop databases.
 - Use -gm 1 to allow a single connection plus one DBA connection above the limit so that a DBA can connect and drop others in an emergency.

An alternate way to restrict connections is to specify

```
sa_server_option 'disable_connections', 'ON'
```

on the connection where you intend to perform the upgrade and

```
sa_server_option 'disable_connections', 'OFF'
```

on the same connection after upgrading. *The disadvantage is that this method precludes emergency access from another DBA connection.*

- 8 Start Interactive SQL and issue the database upgrade statement. For example:

```
ALTER DATABASE UPGRADE
```

If the database was created with the Java options off, append the keywords `JAVA OFF JCONNECT OFF` to the preceding command.

For more information, see “Upgrading non-multiplex databases,” in Chapter 4, “Migrating Data,” in the *Sybase IQ Installation and Configuration Guide*.

3.16 Moving data in one endian format to a system with a different endian format [CR 445754]

The following information was omitted from the 12.7 *Sybase IQ Release Bulletin*.

This section documents a procedure for moving data from a database in big-endian format to a database in little-endian format. This procedure moves table definitions but does not include migration of database objects, such as stored procedures or events, which you must recreate.

For example, Sybase IQ databases built on Sun64 SPARC systems store binary data in big-endian (Most Significant Byte first) format. Because Sun Solaris x64 is a little-endian system, Sybase IQ databases built on Sun64 SPARC cannot be upgraded with `ALTER DATABASE UPGRADE` to run on Sun Solaris x64 systems.

To move data for each database across hardware platforms of different endian structures, you must:

- Copy the database schema from the source platform (tables, indexes, etc.).
- Create a new database on the target platform.
- Perform a binary data dump from the source database.
- Load data into the new target database.

The following steps describe this process in detail.

❖ **Moving data from big-endian to little-endian systems or the reverse**

Note Before you begin, make sure that you have a process for capturing your database and table schema.

The following example loads a table named `lineitem` and identifies one extract file on UFS (file system) called `lineitem_binary.inp`.

Check operating system documentation for the maximum file size for your system. For example, an extract file on Sun Solaris x64 has a maximum size of 512GB.

- 1 Activate the extract utility:

```
SET TEMPORARY OPTION Temp_Extract_Name1 =  
'lineitem_binary.inp'  
SET TEMPORARY OPTION Temp_Extract_Name2 = ''
```

- 2 Set up a binary extract of the `lineitem` table:

```
SET TEMPORARY OPTION Temp_Extract_Binary = 'on'  
SET TEMPORARY OPTION Temp_Extract_Swap = 'off'
```

- 3 Place output in the file `lineitem_binary.inp`:

```
SELECT * FROM lineitem
```

- 4 Turn off the extract utility:

```
SET TEMPORARY OPTION Temp_Extract_Name1 = ''
```

- 5 Create a duplicate of your database on the target system.

- 6 Assuming table *lineitem* as defined below, load the *lineitem* table as follows:

```
LOAD TABLE lineitem
( l_orderkey      BINARY WITH NULL BYTE,
  l_partkey       BINARY WITH NULL BYTE,
  l_suppkey       BINARY WITH NULL BYTE,
  l_linenumbers   BINARY WITH NULL BYTE,
  l_quantity      BINARY WITH NULL BYTE,
  l_extendedprice BINARY WITH NULL BYTE,
  l_discount      BINARY WITH NULL BYTE,
  l_tax           BINARY WITH NULL BYTE,
  l_returnflag    BINARY WITH NULL BYTE,
  l_linestatus    BINARY WITH NULL BYTE,
  l_shipdate      BINARY WITH NULL BYTE,
  l_commitdate    BINARY WITH NULL BYTE,
  l_receiptdate   BINARY WITH NULL BYTE,
  l_shipinstruct  BINARY WITH NULL BYTE,
  l_shipmode      BINARY WITH NULL BYTE,
  l_comment       BINARY WITH NULL BYTE )
FROM 'C:\\mydata\\lineitem_binary.inp'
FORMAT BINARY
STRIP OFF
QUOTES OFF
ESCAPES OFF
PREVIEW ON
BYTE ORDER HIGH
COMMIT
```

Note particularly two clauses:

- BINARY WITH NULL BYTE is required when loading a binary file.
- BYTE ORDER HIGH specifies the byte order from the system where the data *originated*. The source database in this example is a big-endian platform; therefore, this data requires byte order HIGH. (Little-endian databases require byte order LOW.)

When loading a multiplex database, *use absolute (fully-qualified) paths in all file names*. Do not use relative path names.

3.17 Avoiding port number conflicts [CR 358218]

To avoid product conflicts, change the IQ port number in *default.cfg* and other configuration files (for example, *asiqdemo.cfg*) if Adaptive Server Anywhere is installed on the same system as Sybase IQ. Both products use the default port 2638.

- 1 Add the following line to *\$ASDIR/scripts/default.cfg* with an unused port number, for example:

```
-x tcpip{port=4444}
```

- 2 Look for a port number definition in each configuration file. For example, *\$ASDIR/demo/asiqdemo.cfg* contains the following line:

```
-x tcpip{port=2638}
```

- 3 Edit the line and replace the default port number with the new one, for example:

```
-x tcpip{port=4444}
```

- 4 Save each file when finished.

If Adaptive Server Anywhere is on the same subnet as Sybase IQ, the server names must be unique.

3.18 Adjusting HP-UX parameters

You need to adjust certain HP-UX kernel parameters to run Sybase IQ. The HP System Administration Manager (SAM) allows verification and setting of these parameters. See your *Sybase IQ Installation and Configuration Guide* for a complete list of kernel parameter settings.

3.19 Upgrading database options

Sybase IQ 12.7 has many new and modified database options. For information on these options, see *New Features in Sybase IQ 12.7*.

If you have explicitly set database options that adjust performance and want to be sure that the new settings are appropriate for your environment, you may find it helpful to run `sp_iqcheckoptions`, which displays:

- Database options whose value is different from the default
- Current value of these options
- Default value for the connected user

It is a good idea to run `sp_iqcheckoptions` and capture the output before and after running `ALTER DATABASE UPGRADE`.

3.20 Megaphone Telco Demo

The Megaphone Telco Demo is larger than the *asIQdemo* sample database and shows how fast Sybase IQ can be as the amount of data grows. It features sample queries, a star schema, and a fact table with over one million rows. You can download the Megaphone Telco Demo from the Sybase Software Download Page at <http://downloads.sybase.com>.

4. New features and behavior changes

For new features and behavior changes, see *New Features in Sybase IQ 12.7*, which you should read before you install Sybase IQ 12.7. This section describes changes omitted from that book or the documentation set.

4.1 Highlights of new features in Sybase IQ 12.7 ESD #9

Sybase IQ 12.7 ESD #9 includes new features in the areas of administration and performance.

- Sybase IQ now provides a mechanism to verify an existing database backup. See “Verification of a database backup [CR 487181]” on page 36.
- Sybase IQ now allows you to control file system buffering on Linux, Linux IBM, and AIX platforms, in addition to Solaris and Windows platforms. See “Controlling file system buffering [CR 487762]” on page 85.

4.2 Highlights of new features in Sybase IQ 12.7 ESD #4

Sybase IQ 12.7 ESD #4 includes new features in the areas of data load, security, and Network Client for Windows.

- Sybase IQ now supports loading a default value into a column as specified in the `LOAD TABLE` statement using the new `DEFAULT` clause in the column specification.

See “`LOAD TABLE` supports loading default values [CR 396843]” on page 47.

- The Sybase IQ server now supports TDS password encryption.

See “Server support of TDS password encryption [CR 484702]” on page 74.

- Sybase IQ Network Client for Windows on 32-bit platforms now includes a 64-bit ODBC driver.

See “Windows 64-bit ODBC drivers added to 32-bit Network Client on Windows [CR 492213]” on page 12.

- Sybase IQ Network Client for Windows now installs silently on all Windows platforms.

See “Windows Network Client installs silently [CR 468594 and CR 491779]” on page 12.

4.3 Highlights of new features in Sybase IQ 12.7 ESD #3

Sybase IQ 12.7 ESD #3 includes new features in the areas of large object (LOB) support, operational management, ease of use, and performance.

- INSERT...LOCATION now supports retrieving large object data (LOB) up to 20MB in length from a remote database.

See “INSERT...LOCATION supports 20MB large object (LOB) data [CR 453439]” on page 44.

- WORD (WD) indexes are now supported on LONG VARCHAR character large object (CLOB) columns. The CONTAINS predicate is also supported on LONG VARCHAR (CLOB) columns with a WORD (WD) index.

See “WORD index and CONTAINS predicate support for character large object (CLOB) columns [CR 415547]” on page 61.

- String search functions for LONG VARCHAR (CLOB) and LONG BINARY (BLOB) data are now supported.

See “String search functions for large object columns [CR 415543]” on page 64.

- Support of the LIKE predicate on LONG VARCHAR (CLOB) columns is now included.

See “LIKE predicate for character large object (CLOB) columns [CR 415546]” on page 63.

- The LOCK TABLE statement supports acquiring WRITE locks on a set of tables. The LOCK TABLE statement also allows enqueueing for READ, WRITE, and EXCLUSIVE locks for a specified time period.

See “LOCK TABLE support [CR 480880, CR 472791, CR 463196, 444921]” on page 30.

- Authorized users can now display query plans in the Java-based Interactive SQL (dbisql) plan window. You can also save and print query plans from dbisql, instead of accessing the *.iqmsg* file or query plan files on the server.

See “Using query plans with Interactive SQL for Java [338070]” on page 67.

- Workload monitoring enhancements provide a mechanism to analyze a workload to determine the tables/columns/indexes in use, so that unused objects can be dropped to save space, improve DML performance, and decrease backup time.

See “Workload monitoring [CR 472513]” on page 50.

- The login auditing feature, present in previous ESDs, lets you log connection details for a given user in the transaction log file. A new example added to documentation shows how to use the auditing feature to examine attempts to access unauthorized information.

See “Login auditing capability” on page 108.

- The IQ Network Client is now certified on the Windows Vista operating system.

See “Product summary” on page 2.

4.4 Data Definition Language (DDL) changes

This section contains new features and changes related to DDL.

4.4.1 Behavior change declaring variables in procedures [CR 584334]

In Sybase IQ 12.7 and later releases, new procedures are invalid when there are duplicate declarations of a variable within a compound statement. This error is reported:

```
There is already a variable named '@variable_name'  
SQLCODE=-261, ODBC 3 State="42000"
```

Ensure there is only a single instance of a variable in a compound block.

4.4.2 Database created with JAVA ON now returns correct error [CR 469250]

ESD #3 corrects a problem where Sybase IQ returned the same error (Error -845) for two different situations.

Now a database created with JAVA ON correctly reports a different error for incorrect dotted references than a database created with JAVA OFF.

In a database created with JAVA OFF, the following statement reports error -845, `SQLSTATE_INVALID_COLUMN_QUALIFICATION`, "Owner used in a qualified column reference does not match correlation name."

```
CREATE TABLE dot ( a int, b int );
INSERT dot VALUES ( 1, 1 );
COMMIT;
SELECT * FROM dot WHERE
dot..a = 100;
```

In a database created with JAVA ON, the statement reports error -706, `SQLSTATE_OMNI_SERVER_NOT_CAPABLE`, "Remote server does not have the ability to support this statement."

4.4.3 Unique table names on a connection [CR 455684]

In Sybase IQ 12.7 ESD #3, an attempt to create a base table or a global temporary table will fail, if a local temporary table of the same name exists on that connection, as the new table cannot be uniquely identified by *owner.table*.

You can, however, create a local temporary table with the same name as an existing base table or global temporary table. References to the table name access the local temporary table, as local temporary tables are resolved first.

For example, consider the following sequence:

```
CREATE TABLE t1 (c1 int);
INSERT t1 VALUES (9);

DECLARE LOCAL TEMPORARY TABLE t1 (c1 int);
INSERT t1 VALUES (8);

SELECT * FROM t1;
```

The result returned is 8. Any reference to t1 refers to the local temporary table t1 until the local temporary table is dropped by the connection.

4.4.4 LOCK TABLE support [CR 480880, CR 472791, CR 463196, 444921]

Sybase IQ now allows users to reserve WRITE locks in advance to prevent failure of DML update statements like INSERT, DELETE, SYNCHRONIZE JOIN INDEX, UPDATE, and TRUNCATE due to version errors. Sybase IQ now also allows connections to enqueue for READ, WRITE, and EXCLUSIVE locks in case of lock unavailability.

As of 12.7 ESD #3, the LOCK TABLE statement supports acquiring WRITE locks on a set of tables as well as enqueueing on all the lock types. The following updates apply to existing documentation:

Syntax	LOCK TABLE <i>table-list</i> [WITH HOLD] IN { SHARE WRITE EXCLUSIVE } MODE [WAIT [<i>time</i>]]
Parameters	<i>table-list</i>: [<i>owner</i>.] <i>table-name</i> [, [<i>owner</i>.] <i>table-name</i>, ...] <i>time</i>: <i>string</i>
Examples	For example, the following statement obtains a WRITE lock on the customer and employee tables, if available within 5 minutes and 3 seconds: <pre>LOCK TABLE customer, employee IN WRITE MODE WAIT '00:05:03'</pre> The following statement waits indefinitely, until the WRITE lock on the customer and employee tables, if available or an interrupt occurs: <pre>LOCK TABLE customer, employee IN WRITE MODE WAIT</pre> table-name The table must be a base table, not a view. WRITE mode is only valid for IQ base tables. LOCK TABLE either locks all tables in the table list, or none. If obtaining a lock for a SQL Anywhere table, or when obtaining SHARE or EXCLUSIVE locks, you may only specify a single table. Standard Sybase IQ object qualification rules are used to parse <i>table-name</i>. For related details, see the section “Identifiers” in Chapter 3 “SQL Language Elements,” in <i>Reference: Building Blocks, Tables, and Procedures</i> and “Types of tables,” in Chapter 5, “Working with Database Objects,” in the <i>Sybase IQ System Administration Guide</i>. WRITE A WRITE lock on an IQ base table prevents it from being updated by concurrent transactions. WRITE mode lock on a table is compatible with SHARE locks but not with WRITE and EXCLUSIVE locks in other connections.

LOCK TABLE IN WRITE MODE unconditionally commits the connection's outermost transaction and returns successfully if all locks are available within the specified wait time. The snapshot of the database to be used by the transaction is established not by the LOCK TABLE IN WRITE MODE statement, but by the execution of the next command processed by Sybase IQ.

For all LOCK TABLE statement write mode failures, the user's current transaction commits unless there is a syntax error.

A WRITE mode lock on an IQ table X that participates in a join index also locks:

- The top table of the join index hierarchy in WRITE mode when X is a non-top table
- The corresponding join virtual table (JVT)

LOCK TABLE grants write locks only if the user has permission to execute a DML update statement (INSERT/DELETE/UPDATE/TRUNCATE). The user must have one of these permissions on each table in the table-list. Even after write locks are granted, IQ checks for the appropriate permissions for subsequent commands. For example, INSERT permissions are checked when executing an INSERT statement.

Write locks can never be held across transactions (WITH HOLD mode is not supported). Once acquired, locks are released only when the transaction commits or rolls back, or when the connection disconnects. If necessary, the DBA can issue a DROP CONNECTION statement, causing the transaction to roll back.

WAIT time Wait options specify maximum blocking time for all lock types. This option is mandatory when lock mode is WRITE. When a time argument is given, the server locks the specified tables only if available within the specified time. The time argument can be specified in the format *hh:nn:ss:sss*. If a date part is specified, the server ignores it and converts the argument into a timestamp. When no time argument is given, the server waits indefinitely until a WRITE lock is available or an interrupt occurs.

LOCK TABLE on views is unsupported. Attempting to lock a view acquires a shared schema lock regardless of the mode specified in the command. A shared schema lock prevents other transactions from modifying the table schema.

The Transact-SQL (TSQL) stored procedure dialect does not support LOCK TABLE. For example, the following statement returns Syntax error near LOCK:

```
CREATE PROCEDURE tproc()  
AS  
BEGIN  
COMMIT  
LOCK TABLE t1 IN SHARE MODE  
INSERT INTO t1 VALUES (30)  
END
```

The Watcom-SQL stored procedure dialect supports LOCK TABLE. The default command delimiter is a semicolon (;). For example:

```
CREATE PROCEDURE wproc()  
BEGIN  
COMMIT;  
LOCK TABLE t1 IN SHARE MODE;  
INSERT INTO t1 VALUES (20);  
END
```

4.4.5 Change to ROLLBACK TO SAVEPOINT behavior [CR 472791]

As of ESD #3, rollback to savepoint does not release WRITE locks. For example:

```
SAVEPOINT sp1  
INSERT INTO table1 [...]  
ROLLBACK TO SAVEPOINT sp1
```

In ESD #2 and earlier releases, a rollback to savepoint released WRITE locks.

4.4.6 ALTER DBSPACE returns out of space error messages [CR 435501]

In Sybase IQ 12.7 ESD #2, if a dbspace is out of space and an ALTER DBSPACE is attempted on that dbspace without adding new space, or if an ALTER DBSPACE is attempted on any other dbspace, the command fails with error -1010015 "The MAIN store is out of space. Space must be added to the MAIN store first" or -1010016 "The TEMPORARY store is out of space. Space must be added to the TEMPORARY store first."

When error -1010015 is returned, space must be added to the MAIN store before space can be added to any other store or before any other dbspace operation can be performed.

When error -1010016 is returned, space must be added to the TEMPORARY store before space can be added to any other store or before any other dbspace operation can be performed.

4.4.7 NEWID() function in column default values [CR 435215]

A Sybase IQ 12.7 ESD #1 enhancement allows values generated by the NEWID() function to be used as default values for columns.

4.4.8 Limiting number of rows fetched [CR 424766]

In ESD #1, Sybase Central lets you limit the number of rows fetched when displaying table data.

- 1 Right-click Sybase IQ (the plug-in name) in the left pane.
- 2 Choose Preferences from the drop-down menu.
- 3 Click the Table Data tab on the Plug-in Preferences dialog.
- 4 To specify a row limit, type a positive integer value in the text field to limit the number of rows displayed. A zero value (the default) displays all rows.
- 5 Click OK, Apply, or Enter.

4.4.9 Out-of-space status for sp_iqstatus [CR 416006]

In ESD #3, sp_iqstatus displays out-of-space status for main, temporary and local stores. If a store runs into an out-of-space condition, sp_iqstatus shows Y in the store's out-of-space status display value.

4.5 Startup and connection changes

This section contains new features and changes related to startup and connection.

4.5.1 IQ Agent startup takes list as argument [CR 465686]

The IQ Agent startup command accepts a list of alternate host names or IP addresses, separated by commas, in addition to a single argument. Blank spaces are disallowed in the list. IP aliases may be used on a single host that supports multiple networking cards with unique addresses.

For example:

```
S99SybaseIQAgent12 -host
"fiona.sybase.com,10.20.30.40,ciaran,12.20.30.41"
```

To avoid a known issue when binding to fully qualified host names, start the IQ Agent using either the short host name only or all three structures. For example:

```
S99SybaseIQAgent12 -host ciaran -port 2099
```

or

```
S99SybaseIQAgent12 -host
"ciaran,ciaran.sybase.com,10.50.7.70" -port 2199
```

Note that the IQ Agent binds in the following order:

- 1 The local host name
- 2 The host IP address
- 3 The item(s) specified in the -host command line parameter

For example:

```
$ASDIR/bin/S99* -port 2099 -host "10.50.7.70,njal,njal.sybase.com"
Agent Log File: /ABC-IQPROD/test/logfiles/SybaseIQAgent2099.001.log
Starting Agent ...
-----
07/24 14:27:38 Agent: Initializing security manager
07/24 14:27:38 Agent: Startup Timestamp = 07/24 14:27:37.027
07/24 14:27:38 setProperty: asiqalthost (10.50.7.70,njal,njal.sybase.com)
07/24 14:27:38 setProperty: asdir (/ABC-IQPROD/iq127/ASIQ-12_7)
07/24 14:27:38 setProperty: host (PRDNJA)
07/24 14:27:38 setProperty: asiqalthost (10.50.7.71,10.50.7.70,njal,njal.sybase.com)
07/24 14:27:38 getProp: asdir = /ABC-IQPROD/iq127/ASIQ-12_7|
07/24 14:27:38 Agent: Started RMI server
07/24 14:27:38 Agent: Binding to //PRDNJA:2099/Agent
07/24 14:27:38 getProp: asiqalthost =10.50.7.71,10.50.7.70,njal,njal.sybase.com
07/24 14:27:38 Agent: Binding to //10.50.7.71:2099/Agent
07/24 14:27:38 Agent: Binding to //10.50.7.70:2099/Agent
07/24 14:27:38 Agent: Binding to //NJAL:2099/Agent
07/24 14:27:38 Agent: Binding to //NJAL.SYBASE.COM:2099/Agent

07/24 14:27:38 Agent: Agent Started Successfully
07/24 14:27:38 Agent: version 12.7 - 070706
```

4.5.2 Stopping server processes on UNIX and Linux systems [CR 461665]

The stop_asiq utility has been improved in the 12.7 ESD #3 release. It now searches the process table by login, user id (UID), and/or truncated user login. This functionality was added to accommodate UNIX and Linux systems where users with logins longer than eight characters exist. The -user parameter lets you specify a particular user on the stop_asiq command to easily stop server processes no matter how the table displays them.

The following syntax updates current stop_asiq syntax in the *Sybase IQ Utility Guide*.

Syntax **stop_asiq [-agent | -cleanup] [-stop [one | all]] [-wait <seconds>] [-user <user_name>] [-help]**

The user parameter performs two functions.

- If, due to truncation or substitution, stop_asiq cannot find the server/agent owned by the current user, the system manager can specify the name/id found in the process table in the -user argument to shut down the server/agent with the stop_asiq utility.
- A user with root privileges can shut down another user's server/agent without having to log in as that user. The stop_asiq utility has no 'su' or root powers, so a nonprivileged user cannot shut down a server owned by another user.

Example % stop_asiq -user joe

Checking system ...

The following 2 server(s) are owned by other users.

##	Owner	PID	Started	CPU Time	Additional Information
1	bill	15939	Jul.03	522:34	SVR:mpx_w DB:mpx_mult PORT:6722
asiqsrv12 -n mpx_w @/duncan/users/bill/mpx_a/mpx/w/params.cfg					
2	jim	16180	Jul.03	553:45	SVR:mpx_r1 DB:r1 PORT:6732
asiqsrv12 @asiqdemo.cfg asiqdemo.db -gn 25 -o /ciaran/users/jim/127test					

The following 1 server(s) are owned by 'joe'

##	Owner	PID	Started	CPU Time	Additional Information
1:	joe	28155	Mar.30	152:19	SVR:joe_asiqdemo DB:asiqdemo
PORT:1746					
asiqsrv12 @asiqdemo.cfg asiqdemo -gn 25 -o /fiona/users/joe/12.6/ASIQ-12					

Please note that 'stop_asiq' will shut down a server completely without regard for users, connections, or load process status.

For more control, use the 'dbstop' utility, which has options that control stopping servers based on active connections.
Do you want to stop the server displayed above <Y/N>?

4.5.3 TDS connects through default database [CR 449141]

Sybase IQ now allows connections to continue through to the default database even if the Open Client login server name does not match the name of the default database, provided that the connection string does not start a database (that is, there is no DBF=parameter) and the server is running only one database. If both of these conditions are met, then the Open Client connection to the default database is allowed.

4.6 Administration and troubleshooting improvements

This section contains new features and changes related to administration and troubleshooting.

4.6.1 Raw partition sector size for backup and restore [CR 531331]

Sybase IQ 12.7 ESD #5 and 12.7 ESD #5.x allow backup and restore using raw partitions formatted with a sector size of 4096 bytes.

4.6.2 Verification of a database backup [CR 487181]

As of version 12.7 ESD #9, Sybase IQ provides a mechanism to verify an existing Sybase IQ version 12.6 or 12.7 database backup using the VERIFY clause of the RESTORE SQL statement. The verification process checks the specified archive for the same errors a restore process checks, but performs no write operations. All status messages and detected errors are written to the server log file.

Syntax

```
RESTORE DATABASE 'db_file'
FROM 'archive_device' [ FROM 'archive_device' ]...
[ CATALOG ONLY ]
[ KEY key_spec ]
[ [ RENAME dbspace-name TO 'new-dbspace-path']...
  | VERIFY [ COMPATIBLE ] ]
```

See “RESTORE statement” in Chapter 6, “SQL Statements” in *Sybase IQ Reference Manual* for the complete RESTORE syntax, parameter descriptions, and usage.

Example

```
RESTORE DATABASE <database_name.db>
FROM '/sys1/dump/dmp1'
FROM '/sys1/dump/dmp2'
VERIFY
```

Usage

The VERIFY clause of the RESTORE command directs the server to validate the specified Sybase IQ database backup archives for a full, incremental, incremental since full, or virtual backup without performing any write operations.

You cannot use the RENAME clause with the VERIFY clause; the error “VERIFY and RENAME cannot be used together in a RESTORE statement” is reported.

The backup verification process can run on a different host than the database host.

If the COMPATIBLE clause is specified with VERIFY, then the compatibility of an incremental archive is checked with the existing database files. If the database files do not exist on the system on which RESTORE...VERIFY COMPATIBLE is invoked, an error is returned. If COMPATIBLE is specified while verifying a full backup, the keyword is ignored; no compatibility checks need to be made while restoring a full backup.

Note The verification of a backup archive is distinct from the database consistency checker (DBCC) verify mode (`sp_iqcheckdb 'verify...'`). RESTORE VERIFY validates the consistency of the backup archive to be sure it can be restored, whereas DBCC validates the consistency of the database data.

You should run `sp_iqcheckdb 'verify...'` before taking a backup. If an inconsistent database is backed up, then restored from the same backup archive, the data continues to be in an inconsistent state, even if RESTORE VERIFY reports a successful validation.

Verification of an incremental backup

If RESTORE VERIFY is specified without COMPATIBLE for an incremental restore, Sybase IQ does not look for any dbspaces and does not perform any compatibility checks. No warning is reported, even if the files do not exist. The compatibility check is performed only if the COMPATIBLE clause is specified.

If RESTORE VERIFY COMPATIBLE is specified for an incremental restore and the IQ catalog store or any of the Sybase IQ dbspaces do not exist, the compatibility check cannot be made; an error is reported and the operation fails.

During the validation of an incremental backup, the RESTORE VERIFY COMPATIBLE process opens the Sybase IQ dbspaces in readonly mode in order to do consistency checking. No dbspaces are modified during the verification process.

In the case of incremental restores, if the database has been modified or the particular incremental archive is not the correct archive for the database, RESTORE VERIFY COMPATIBLE reports the error “Database has changed since last restore.” (SQLCODE -1012008, SQLSTATE QUA08) or “This restore cannot immediately follow the previous restore.” (SQLCODE -1012009, SQLSTATE QUA09).

Verification progress reporting

The RESTORE VERIFY process verifies every stripe specified in the command. As the verification process checks the stripes and their corresponding files, it reports progress in terms of the number of IQ blocks verified. After every 5000 IQ blocks verified, a message is displayed in the server log file with the percentage of the number of IQ blocks completed:

```
5000/100000 (5%) Blocks verified
```

A message is displayed in the server log file when 100% of the IQ blocks are verified.

The messages “VERIFY RESTORE Started,” the number of IQ blocks to be verified, and “VERIFY RESTORE successfully completed” are also written in the server log, when the verify action starts and completes, respectively. For example:

```
I. 11/17 06:45:24. VERIFY RESTORE Started
I. 11/17 06:45:24. Total number of IQ blocks to be
verified: 764
I. 11/17 06:45:24. Total number of IQ blocks verified:
764/764 ( 100 % )
I. 11/17 06:45:24. VERIFY RESTORE Successfully Complete
```

Verification error reporting

If the verification process finds errors after which it can continue, the process continues checking the archive and logs information for the errors detected. The errors for which the verification can continue are:

Header of block to be restored appears to be corrupted. (SQLCODE -10120111, SQLSTATE QUA11)

Media data appears corrupted (bad checksum). (SQLCODE -1012012, SQLSTATE QUA12)

Media meta data appears corrupted (boundary record). (SQLCODE -1012013, SQLSTATE QUA13)

Media meta data appear corrupted (multiple begin boundary records). (SQLCODE -1012014, SQLSTATE QUA14

If any of these errors are found and the verification process is able to continue to the end of the archive, the error “The verification of the provided archive has failed. Please check the server log for details of the errors thrown during verify.” is reported.

If any error pertaining to RESTORE is found other than the four errors above, the error that occurred is reported and the verification process stops.

Permissions

Must have DBA or backup authority.

See also

- “RESTORE statement” in Chapter 6, “SQL Statements” in *Sybase IQ Reference Manual*
- “BACKUP statement” in Chapter 6, “SQL Statements” in *Sybase IQ Reference Manual*
- Chapter 14, “Data Backup, Recovery, and Archiving” in *Sybase IQ System Administration Guide*

4.6.3 Unsupported procedures and views [CR 460819, CR 442581]

The following list contains Adaptive Server Enterprise stored procedures that Sybase IQ no longer supports:

- sp_addserver
- sp_configure
- sp_estspace
- sp_help
- sp_helpuser
- sp_who

The following list contains the catalog procedures that Sybase IQ no longer supports:

- sp_column_privileges
- sp_databases
- sp_datatype_info
- sp_server_info

Obsolete procedures and views are removed from databases created prior to Sybase IQ 12.7 ESD #2 when you run the ALTER DATABASE UPGRADE command using Sybase IQ 12.7 ESD #2 or higher.

If you do not run ALTER DATABASE UPGRADE, the procedures remain in the database.

The following list contains the Adaptive Server Enterprise views that Sybase IQ no longer supports:

- sysalternates
- sysaudits
- sysauditoptions
- sysconstraints
- syscharsets
- sysconfigures
- syscurconfigs
- sysdatabases
- sysdepends
- sysdevices
- sysengines
- syskeys
- syslanguages
- syslocks
- syslogs
- sysloginroles
- sysmessages
- sysprocedures
- sysprocesses
- sysprotects
- sysreferences
- sysremotelogins
- sysroles

- syssegments
- syssservers
- sysssrvroles
- systhresholds
- sysusages

4.6.4 DBSpace Header utility (iqheader) [CR 427630]

The dbspace management enhancements added in Sybase IQ 12.6 allowed users to view an IQ server dbspace configuration from a connected user. At times, an administrator might want to determine which server, if any, is using a particular LUN/device/file as a dbspace in order to analyze disk usage or to configure a multiplex query server.

In ESD #1, the iqheader utility reports the configuration of an arbitrary device regardless of whether it is currently in use by an IQ server.

The user interface is a standalone console application called iqheader (header.exe on Windows). The iqheader tool searches the device for a IQ dbspace header and reports the header information in a user-readable format.

Note LUN is a Logical Unit Number and is used to identify SCSI devices so the host can address and access the data on each disk drive in an array.

Syntax **iqheader** [[*dbspace_path*]

Parameters The iqheader application takes a single parameter, which is the device to be checked.

Usage **iqheader usage** When invoked with no parameters, a usage summary is reported and a non-zero status is returned:

```
>iqheader
Usage: iqheader [dbspace_path]
```

iqheader error When the specified target is not an IQ dbspace, an error message is reported and a non-zero status is returned:

```
>iqheader /dev/null
Not an IQ file: Error 0
```

Operating system errors When the specified target is unreadable or any file operation fails due to an error returned from the operating system, the native operating system error displays and a non-zero status is returned.

```
>iqheader /dev/rdsk/c1t32d0s1          <
Open Failed: No such file or directory

>iqheader /dev/rdsk/c1t3d0s1          <
Open Failed: Permission denied
```

iqheader output When a valid IQ dbspace is specified, iqheader prints the dbspace configuration to the console and returns a 0 exit status. Table 3 describes the fields in the iqheader output.

Table 3: iqheader output

Field	Description
Full Path	Full path after symbolic link resolution
Version	dbspace file format version
File ID	Unique number assigned to each dbspace
Create Time	Time of dbspace creation
RW Mode	Current Readwrite mode: RW, RO, RW, N/A (Upgraded)
Last RW Mode	Last dbspace mode
Size (MB)	DBSpace size in Megabytes
Reserve (MB)	DBSpace reserve size in Megabytes
Block Size	Size of IQ block in Bytes
Page Size	Size of IQ page in Bytes
First Block	First IQ block number mapped to dbspace
Block Count	Number of IQ blocks that map to actual disk blocks
Reserve Blocks	Number of IQ blocks that may be added to this dbspace
Last Real Block	Last IQ block number that maps to an actual disk block
Last Mapped Block	Last IQ block number mapped to dbspace
Online	Online Status (YES/NO)
Create ID	Commit ID in which dbspace was created
Alter ID	Last Commit ID in which dbspace was altered
DBID1	Location of 1st database identity
DBID2	Location of 2nd database identity
CREATE DBSPACE command	Template create dbspace command

Example

The following example shows output for the iqheader command when executed against an IQ dbspace on a Solaris raw partition.

```
>iqheader /dev/rdisk/c1t17d0s4
File Name: /dev/rdisk/c1t17d0s4

Full Path:
/devices/sbus@2,0/SUNW,socal@d,10000/sf@1,0/ssd@w21000
02037229873,0:e,raw

DBSpace Header Info
Version: 2
File ID: 16404
Create Time: 2006-04-03 09:32:00
RW Mode: RW
Last RW Mode: N/A
Size (MB): 1025
Reserve (MB): 0
Block Size: 4096
Page Size: 65536
First Block: 2611040
Block Count: 262575
Reserve Blocks: 0
Last Real Block: 2873614
Last Mapped Block: 3133247
Online: YES
Create ID: 90
Alter ID: 0
vDBID1: 0
DBID2: 0

CREATE DBSPACE command
create dbspace "[DBSPACENAME]" as '/dev/rdisk/c1t17d0s4'
iq [MAIN|TEMPORARY|LOCAL] store
server "[SERVER_NAME]" file id 16404 prefix 65536
finish 0 first block 2611040 block count 262575
```

See also

Chapter 5, “Working with Database Objects,” in the *Sybase IQ System Administration Guide*.

4.7 Data load, update, and extraction enhancements

This section contains new features and changes related to data load, update, and extraction.

4.7.1 INSERT...LOCATION supports 20MB large object (LOB) data [CR 453439]

Sybase IQ 12.7 ESD #3 increases the limit on the length of large object data (LOB) that can be retrieved from a remote database using INSERT...LOCATION from 32767 bytes to 20MB (20971520 bytes).

An Adaptive Server Enterprise IMAGE column inserted into a LONG BINARY column using the INSERT...LOCATION command is now silently right truncated at 20971520 bytes, rather than 32767 bytes.

An Adaptive Server Enterprise TEXT column inserted into a LONG VARCHAR column using the INSERT...LOCATION command is now silently right truncated at 20971520 bytes, rather than 32767 bytes.

See “Documentation updates and clarifications” for documentation changes related to this feature.

Users must be specifically licensed to use the Large Objects Management functionality. For details on the Large Objects Management option, see *Large Objects Management in Sybase IQ*.

4.7.2 Implicit conversion between BIT and CHAR/VARCHAR data types supported [CR 451862]

Sybase IQ 12.7 ESD #2 supports implicit conversion between BIT and CHAR and BIT and VARCHAR data types for arithmetic, comparison, INSERT, and UPDATE operations. For more information, see “Implicit conversion between BIT and CHAR/VARCHAR data types supported [CR 451862]” on page 142.

4.7.3 INSERT...LOCATION supports additional ASE data types [CR 442348]

12.7 ESD #2 supports inserting data from an Adaptive Server Enterprise database column of data type DATE, TIME, UNSIGNED SMALLINT, NCHAR, or NVARCHAR using the INSERT...LOCATION syntax of the INSERT statement.

Note The Adaptive Server Enterprise data type UNSIGNED SMALLINT is not supported by Sybase IQ INSERT...LOCATION on Linux 32-bit systems, but is supported by INSERT...LOCATION on all other Sybase IQ supported platforms.

4.7.4 LOAD TABLE supports BCP files as input [CR 437332]

Sybase IQ 12.7 ESD #2 supports using a BCP character file as input to the LOAD TABLE command. The BCP input file must be generated by the BCP OUT command with the -c option. The LOAD TABLE FORMAT BCP feature is an alternative to the iq_bcp utility and will perform better than iq_bcp in all cases.

Syntax

```
LOAD [ INTO ] TABLE [ owner.]table-name
... ( load-specification [, ...] )
... FROM { 'filename-string' | filename-variable } [, ...]
... [ CHECK CONSTRAINTS { ON | OFF } ]
... [ DEFAULTS { ON | OFF } ]
... QUOTES OFF
... ESCAPES OFF
... [ FORMAT { ascii | binary | bcp } ]
... [ DELIMITED BY 'string' ]
...
```

Parameters

load-specification:

```
{ column-name [ column-spec ]
| FILLER ( filler-type ) }
```

column-spec:

```
{ ASCII ( input-width )
| BINARY [ WITH NULL BYTE ]
| PREFIX { 1 | 2 | 4 }
| 'delimiter-string'
| DATE ( input-date-format )
| DATETIME ( input-datetime-format ) }
[ NULL ( { BLANKS | ZEROS | 'literal', ... } ) ]
| ENCRYPTED ( data-type 'key-string' [, 'algorithm-string' ] )
```

Example

The following LOAD TABLE statement is now supported with the FORMAT BCP load option:

```
LOAD TABLE t1 (c1, c2, c3)
FROM 'bcp_file.bcp'
FORMAT BCP
...
```

Usage

FORMAT Sybase IQ supports ASCII and binary input fields. The format is usually defined by the *column-spec* described above. If you omit that definition for a column, by default Sybase IQ uses the format defined by this option. Input lines are assumed to have ascii (the default) or binary fields, one row per line, with values separated by the column delimiter character.

Sybase IQ also accepts data from BCP character files as input to the LOAD TABLE command.

- The BCP data file loaded into Sybase IQ tables using the LOAD TABLE FORMAT BCP statement must be exported (BCP OUT) in cross-platform file format using the -c option. Files generated by the iq_bcp utility (using the -c option) can be used as input to the LOAD TABLE FORMAT BCP command.
- For FORMAT BCP, the default column delimiter for the LOAD TABLE statement is <tab> and the default row terminator is <newline>.
- For FORMAT BCP, the last column in a row must be terminated by the row terminator, not by the column delimiter. If the column delimiter is present before the row terminator, then the column delimiter is treated as a part of the data.
- Data for columns that are not the last column in the load specification must be delimited by the column delimiter only. If a row terminator is encountered before a column delimiter for a column that is not the last column, then the row terminator is treated as a part of the column data.
- Column delimiter can be specified via the DELIMITED BY clause. For FORMAT BCP, the delimiter must be less than or equal to 10 characters in length. An error is returned, if the delimiter length is more than 10.
- For FORMAT BCP, the load specification may contain only column names, NULL, and ENCRYPTED. An error is returned, if any other option is specified in the load specification.

For example, the following LOAD TABLE load specifications are valid:

```
LOAD TABLE x( c1, c2 null(blanks), c3 )
FROM 'bcp_file.bcp'
FORMAT BCP
...
```

```
LOAD TABLE x( c1 encrypted(bigint,'KEY-ONE','aes'),
c2, c3 )
FROM 'bcp_file.bcp'
FORMAT BCP
...
```

For information on the LOAD TABLE ENCRYPTED clause, see *Encrypted Columns in Sybase IQ*.

Error messages

Sybase IQ no longer returns an error message when `FORMAT BCP` is specified as a `LOAD TABLE` option.

- If the specified load format is not `ASCII`, `BINARY`, or `BCP`, Sybase IQ now returns the message "Only ASCII, BCP and BINARY are supported LOAD formats."
- If the `LOAD TABLE` column specification contains anything other than column name, `NULL`, or `ENCRYPTED`, then Sybase IQ returns the new error message "Invalid load specification for LOAD ... FORMAT BCP."
- If the column delimiter or row terminator size for the `FORMAT BCP` load is greater than 10 characters, then Sybase IQ returns the message "Delimiter '%2' must be 1 to %3 characters in length." (where %3 equals 10).

Messages corresponding to error or warning conditions which can occur for `FORMAT BCP` as well as `FORMAT ASCII` are the same for both formats.

4.7.5 LOAD TABLE supports loading default values [CR 396843]

Sybase IQ 12.7 ESD #4 supports loading a default value into a column as specified in the `LOAD TABLE` statement. You can now load a default value into a column, even if the column does not have a default value defined in the table schema. This feature carries the performance benefits of the `LOAD TABLE` statement and provides the user with more flexibility at load time.

The `LOAD TABLE` syntax includes a new `DEFAULT` clause in the column specification.

Syntax

```
LOAD [ INTO ] TABLE [ owner.]table-name
... ( load-specification [, ...] )
... FROM { 'filename-string' | filename-variable } [, ...]
... [ CHECK CONSTRAINTS { ON | OFF } ]
... [ DEFAULTS { ON | OFF } ]
... QUOTES OFF
... ESCAPES OFF
... [ FORMAT { ascii | binary } ]
... [ DELIMITED BY 'string' ]
... [ STRIP { ON | OFF } ]
.
.
.
```

Parameters

load-specification:

```
{ column-name [ column-spec ]
| FILLER ( filler-type ) }
```

column-spec:

```
{ ASCII ( input-width )
| BINARY [ WITH NULL BYTE ]
| PREFIX { 1 | 2 | 4 }
| 'delimiter-string'
| DATE ( input-date-format )
| DATETIME ( input-datetime-format ) }
[ NULL ( { BLANKS | ZEROS | 'literal', ... } ) ]
| ENCRYPTED ( data-type 'key-string' [, 'algorithm-string' ] )
| DEFAULT default-value
```

Examples

The following LOAD TABLE statements are now supported with the DEFAULT load clause:

```
LOAD TABLE t1 (c1 DEFAULT '12345 ', c2, c3, filler(1))
FROM 'LoadConst04.dat'
DEFAULTS ON
STRIP OFF
QUOTES OFF
ESCAPES OFF
DELIMITED BY ',';

LOAD TABLE t1 (c1, c2, c3 DEFAULT '10')
FROM 'bcp_file.bcp'
DEFAULTS ON
FORMAT BCP
QUOTES OFF
ESCAPES OFF;
```

Usage

DEFAULT Use the DEFAULT value specifier in the column specification to specify a load default column value. The LOAD TABLE DEFAULTS clause must be set to ON in order to use the default value specified in the LOAD TABLE statement. If the DEFAULTS clause is OFF, the specified load default value is not used and a NULL value is inserted into the column instead.

The LOAD TABLE DEFAULT value specification must contain at least one column that needs to be loaded from the file specified in the LOAD TABLE command. Otherwise, an error is reported and the load is not performed.

The specified load default value must conform to the supported default values for columns and default value restrictions as described in the section “Using column defaults” in Chapter 9, “Ensuring Data Integrity” of the *Sybase IQ System Administration Guide*. The LOAD TABLE DEFAULT value specification does not support AUTOINCREMENT, IDENTITY, or GLOBAL AUTOINCREMENT as a load default value.

The LOAD TABLE DEFAULT value specification must be of the same character set as that of the database.

Encryption of the default value is not supported for the load default values specified in the LOAD TABLE DEFAULT value specification.

A constraint violation caused by evaluation of the specified load default value is counted for each row that is inserted in the table.

Error messages

Sybase IQ no longer returns an error message when FORMAT BCP is specified as a LOAD TABLE clause.

- If the specified load format is not ASCII, BINARY, or BCP, Sybase IQ now returns the message “Only ASCII, BCP and BINARY are supported LOAD formats.”
- If the LOAD TABLE column specification contains anything other than column name, NULL, or ENCRYPTED, then Sybase IQ returns the new error message “Invalid load specification for LOAD ... FORMAT BCP.”
- If the column delimiter or row terminator size for the FORMAT BCP load is greater than 10 characters, then Sybase IQ returns the message “Delimiter ‘%2’ must be 1 to %3 characters in length.” (where %3 equals 10).

Messages corresponding to error or warning conditions which can occur for FORMAT BCP as well as FORMAT ASCII are the same for both formats.

- If the load default value specified is AUTOINCREMENT, IDENTITY, or GLOBAL AUTOINCREMENT, the error “Default value %2 cannot be used as a LOAD default value. %1” is reported.
- If the LOAD TABLE specification does not contain any columns that need to be loaded from the file specified, the error “The LOAD statement must contain at least one column to be loaded from input file.” is reported and the LOAD TABLE statement rolls back.

4.8 Query enhancements, optimization, and changes

This section contains new features and changes related to queries and optimization.

4.8.1 OS_FILE_CACHE_BUFFERING option behavior change [CR 487762]

In versions prior to Sybase IQ 12.7 ESD #9, the database option `OS_FILE_CACHE_BUFFERING` controls the file system buffering of both IQ Main dbspaces and IQ Temporary dbspaces. As of Sybase IQ 12.7 ESD #9, `OS_FILE_CACHE_BUFFERING` controls the file system buffering of IQ Main dbspaces only. The new database option `OS_FILE_CACHE_BUFFERING_TEMPDB` controls the file system buffering of IQ Temporary dbspaces.

See “Controlling file system buffering [CR 487762]” on page 85 for a description of the database option `OS_FILE_CACHE_BUFFERING_TEMPDB`.

4.8.2 Workload monitoring [CR 472513]

Indexes are often created to provide optimization metadata and to enforce uniqueness and primary/foreign key relationships. Once an index is created, however, DBAs face the challenge of quantifying benefits that the index provides.

Tables are often created in the IQ Main Store for the temporary storage of data that must be accessed by multiple connections or over a long period. These tables might be forgotten while they continue to use valuable disk space. Moreover, the number of these tables in a data warehouses is too large and the workloads are too complex to manually analyze usage.

Thus, unused indexes and tables waste disk space, increase backup time, and degrade DML performance.

In ESD #3, Sybase IQ offers tools for collecting and analyzing statistics for a defined workload. DBAs can quickly determine which database objects are being referenced by queries and should be kept. Unused tables, columns, and indexes can be dropped to reduce wasted space, improve DML performance, and decrease backup time.

Workload monitoring is implemented using stored procedures, which control the collection and report detailed usage of table, column, and, index information. These procedures complement INDEX_ADVISOR functionality, which generates messages suggesting additional column indexes that may improve performance of one or more queries. Once recommended indexes have been added, their usage can be tracked to determine if they are worth keeping.

For more information, see “INDEX_ADVISOR option” and “sp_iqindexadvice procedure,” in the *Sybase IQ Reference Manual*.

sp_iqworkmon procedure

Function	Controls collection of workload monitor usage information and reports monitoring collection status.
Syntax	<pre>sp_iqworkmon ['<i>action</i>'] [, '<i>mode</i>']</pre> action = 'start' , 'stop' , 'status' , 'reset' mode = 'index' , 'table' , 'column' , 'all' For example: <pre>sp_iqworkmon 'start' , 'all'</pre> If one argument is specified, it can only be <i>action</i>. For example: <pre>sp_iqworkmon 'stop'</pre>
Permissions	DBA authority required. Users without DBA authority must be granted EXECUTE permission in order to run the stored procedure.
Usage	action Specifies the control action to apply. A value of <i>start</i> starts monitoring for the specified mode immediately. A value of <i>stop</i> stops monitoring immediately. A value of <i>reset</i> clears all collected usage information. A value of <i>status</i> (the default) displays the current status without changing state. mode Specifies the type of monitoring to control. The INDEX, TABLE, and COLUMN keywords individually control monitoring of index usage, table usage, and column usage respectively. The default ALL keyword controls monitoring of all usage monitoring features simultaneously. A result set always displays when sp_iqworkmon is executed. If a specific mode is specified (such as index), only the row for that mode displays. Usage is collected only for SQL statements containing a FROM clause; for example, SELECT, UPDATE, and DELETE.

Table 4: *sp_iqworkmon* columns

Column name	Description
MonMode	Table, index, or column
Status	Started or stopped
Rowcount	Current number of rows collected

Example

The following table illustrates sample output from the *sp_iqworkmon* procedure.

MonMode	Status	Rowcount
index	started	15
table	started	10
column	started	31

sp_iqcolumnuse procedure

Function Reports detailed usage information for columns accessed by the workload.

Syntax **sp_iqcolumnuse**

Permissions DBA authority required. Users without DBA authority must be granted EXECUTE permission in order to run the stored procedure.

Description Columns from tables created in SYSTEM are not reported.

Table 5: *sp_iqcolumnuse* columns

Column name	Description
TableName	Table name
ColumnName	Column name
Owner	Username of column owner
UID**	Column Unique Identifier
LastDT	Date/time of last access
NRef	Number of query references

**UID is a number assigned by the system that uniquely identifies the instance of the column (where instance is defined when an object is created).

Example The following table illustrates sample output from the `sp_iqcolumnuse` procedure.

TableName	ColumnName	Owner	UID	LastDT	NRef
orders	o_orderdate	DBA	151	20070917 22:41:22..	13
orders	o_shippriority	DBA	154	20070917 22:41:22..	13
lineitem	l_orderkey	DBA	186	20070917 22:41:22..	13
lineitem	l_extendedp..	DBA	191	20070917 22:41:22..	13
lineitem	l_discount	DBA	192	20070917 22:41:22..	13
lineitem	l_shipdate	DBA	196	20070917 22:41:22..	13
#tmp1	expression	DBA	10000000001218	20070917 22:57:36..	1
#tmp1	expression	DBA	10000000001222	20070917 22:41:58..	1
...					

Note The long numbers in the example above are temporary IDs.

sp_iqindexuse procedure

Function Reports detailed usage information for secondary (non-FP) indexes accessed by the workload.

Syntax **sp_iqindexuse**

Permissions DBA authority required. Users without DBA authority must be granted EXECUTE permission in order to run the stored procedure.

Description Each secondary index accessed by the workload displays a row. Indexes that have not been accessed are not displayed. Index usage is broken down by optimizer, constraint, and query usage.

Indexes from tables created in SYSTEM are not reported.

Table 6: *sp_iqindexuse* columns

Column name	Description
IndexName	Index name
TableName	Table name
Owner	Username of index owner
UID**	Index Unique Identifier
Type	Index type
LastDT	Date/time of last access
NOpt	Number of metadata/uniqueness accesses
NQry	Number of Query accesses
NConstraint	Number of accesses for Unique or RI checks

**UID is a number assigned by the system that uniquely identifies the instance of the index (where instance is defined when an object is created).

Example

The following table illustrates sample output from the *sp_iqindexuse* procedure.

IndexName	TableName	Owner	UID	Type	LastDT	NOpt	NQry	NConstraint
n_nationkey_hg	nation	DBA	29	HG	20070917 22:08:06~	12	0	12
n_regionkey_hg	nation	DBA	31	HG	20070917 22:08:06~	12	0	0
r_regionkey_hg	region	DBA	47	HG	20070917 22:08:06~	12	0	12
s_suppkey_hg	supplier	DBA	64	HG	20070917 22:08:06~	12	0	12
p_partkey_hg	part	DBA	87	HG	20070917 22:08:06~	6	0	6
s_suppkey_hg	supplier	DBA	64	HG	20070917 22:08:06~	12	0	12
...								

sp_iqtableuse procedure

Function	Reports detailed usage information for tables accessed by the workload.
Syntax	sp_iqtableuse
Permissions	DBA authority required. Users without DBA authority must be granted EXECUTE permission in order to run the stored procedure.
Description	Tables created in SYSTEM are not reported.

Table 7: *sp_iqtableuse* columns

Column name	Description
TableName	Table name
Owner	Username of table owner
UID**	Table Unique Identifier
LastDT	Date/time of last access
NRef	Number of query references

**UID is a number assigned by the system that uniquely identifies the instance of the table (where instance is defined when an object is created).

sp_iqunusedcolumn procedure

Function	Reports IQ columns that were not referenced by the workload.
Syntax	sp_iqunusedcolumn
Permissions	DBA authority required. Users without DBA authority must be granted EXECUTE permission in order to run the stored procedure.
Description	Columns from tables created in SYSTEM or local temporary tables are not reported.

Table 8: *sp_iqunusedcolumn* columns

Column name	Description
TableName	Table name
ColumnName	Column name
Owner	Username of column owner

Example	The following table illustrates sample output from the sp_iqunusedcolumn procedure.
---------	---

TableName	ColumnName	Owner
sales_order	id	DBA
sales_order	cust_id	DBA
sales_order	order_date	DBA
sales_order	fin_code_id	DBA
sales_order	region	DBA
sales_order	sales_rep	DBA
sales_order_items	id	DBA
sales_order_items	line_id	DBA
sales_order_items	prod_id	DBA
sales_order_items	quantity	DBA
sales_order_items	ship_date	DBA
contact	id	DBA
contact	last_name	DBA
contact	first_name	DBA
contact	title	DBA
contact	street	DBA
contact	city	DBA
contact	state	DBA
contact	zip	DBA
contact	phone	DBA
contact	fax	DBA
customer	id	DBA
...		

sp_iqunusedindex procedure

Function	Reports IQ secondary (non-FP) indexes that were not referenced by the workload.
Syntax	sp_iqunusedindex
Permissions	DBA authority required. Users without DBA authority must be granted EXECUTE permission in order to run the stored procedure.
Description	Indexes from tables created in SYSTEM or local temporary tables are not reported.

Table 9: *sp_iqunusedindex* columns

Column name	Description
IndexName	Index name
TableName	Table name
Owner	Username of index owner
IndexType	Index type

Example

The following table illustrates sample output from the *sp_iqunusedindex* procedure.

IndexName	TableName	Owner	IndexType
ASIQ_IDX_T450_I7_HG	sales_order	DBA	HG
ASIQ_IDX_T450_C6_HG	sales_order	DBA	HG
ASIQ_IDX_T450_C4_HG	sales_order	DBA	HG
ASIQ_IDX_T450_C2_HG	sales_order	DBA	HG
ASIQ_IDX_T451_I6_HG	sales_order_items	DBA	HG
ASIQ_IDX_T451_C3_HG	sales_order_items	DBA	HG
ASIQ_IDX_T451_C1_HG	sales_order_items	DBA	HG
ASIQ_IDX_T452_I11_HG	contact	DBA	HG
ASIQ_IDX_T453_I10_HG	customer	DBA	HG
ASIQ_IDX_T454_I4_HG	fin_code	DBA	HG
ASIQ_IDX_T455_I5_HG	fin_data	DBA	HG
ASIQ_IDX_T455_C3_HG	fin_data	DBA	HG
ASIQ_IDX_T456_I8_HG	product	DBA	HG
ASIQ_IDX_T457_I4_HG	department	DBA	HG
ASIQ_IDX_T457_C3_HG	department	DBA	HG
ASIQ_IDX_T458_I21_HG	employee	DBA	HG
ASIQ_IDX_T458_C5_HG	employee	DBA	HG
ASIQ_IDX_T459_I7_HG	alt_sales_order	DBA	HG
ASIQ_IDX_T460_I6_HG	alt_sales_order_items	DBA	HG

sp_iqunusedtable procedure

Function	Reports IQ tables that were not referenced by the workload.
Syntax	sp_iqunusedtable
Permissions	DBA authority required. Users without DBA authority must be granted EXECUTE permission in order to run the stored procedure.
Description	Tables created in SYSTEM and local temporary tables are not reported.

Table 10: sp_iqunusedtable columns

Column name	Description
TableName	Table name
Owner	Username of table owner

Example

The following table illustrates sample output from the sp_iqunusedtable procedure.

TableName	Owner
fin_code	DBA
contact	DBA
employee	DBA
empl	DBA
alt_sales_order_items	DBA
sales_order	DBA
fin_data	DBA
department	DBA
alt_sales_order	DBA
sales_order_items	DBA
product	DBA
iq_dummy	DBA
customer	DBA
sale	DBA

sp_iqdbspace procedure enhancement

Block type reporting is enhanced to report temporary storage used by the workload monitor. Additional block types are defined in the following table.

Identifier	Block Type
T	Table use
U	Index use
N	Column use

4.8.3 Results returned in a case insensitive database [CR 472049]

In some cases, an optimization in Sybase IQ 12.7 causes the results of a comparison returned in a case insensitive database to be different from the results returned in previous releases of IQ. You cannot expect case sensitive results in a database that is case insensitive (CASE IGNORE). This has always been true for Sybase IQ.

Character strings inserted into tables are always stored in the case they are entered, regardless of whether the database is case sensitive or not. If the string Value is inserted into a character data type column, the string is always stored in the database with an uppercase V and the remainder of the letters lowercase. SELECT statements return the string as Value. If the database is not case sensitive, however, all comparisons make Value the same as value, VALUE, and so on. The IQ server may return results in any combination of lowercase and uppercase, so you cannot expect case sensitive results in a database that is case insensitive.

For example, given the following table and data:

```
CREATE TABLE tb (id int NOT NULL,  
                  string VARCHAR(30) NOT NULL);  
INSERT INTO tb VALUES (1, 'ONE');  
SELECT * FROM tb WHERE string = 'oNe';
```

The result of the SELECT can be 'oNe' (as specified in the WHERE clause) and not necessarily 'ONE' (as stored in the database).

Similarly, the result of

```
SELECT * FROM tb WHERE string = 'One';
```

can be 'One' and the result of

```
SELECT * FROM tb WHERE string = 'ONE';
```

can be 'ONE'.

4.8.4 READ ONLY cursors now see data at time of OPEN [CR 466135]

As of Sybase IQ 12.7 ESD #4, applications that use cursors declared FOR READ ONLY may notice a change in which version of the data is visible to a cursor. Read only cursors now see the version of table(s) on which the cursor is declared when the cursor is opened. Prior to this change, cursors declared FOR READ ONLY saw the version of table(s) at the time of the first FETCH.

For example,

```
CREATE TABLE t1 ( c1 INT );
INSERT t1 VALUES ( 1 );

DECLARE t1_cursor CURSOR FOR SELECT * FROM t1
FOR READ ONLY;
OPEN t1_cursor;

INSERT t1 VALUES ( 2 );

FETCH T1_CURSOR;
```

When the cursor is fetched, only one row may be fetched from the table. (There is only one row in the table, when the cursor is opened.)

Prior to this change, the cursor could fetch two rows from the table.

4.8.5 Behavior change for columns with null values [CR 435353]

An ESD #1 enhancement optimizes queries involving columns that have a significant number of null values. Such queries now run faster. To support this query optimization, the process of inserting or updating data in a table may take longer (compared with previous releases) in cases where a significant number of NULL values are being inserted into the table.

4.8.6 Support for Join_Preference hint on join equality conditions [CR 433686]

An ESD #1 enhancement allows users to specify a join algorithm preference that does not affect every join in the query.

Simple equality join predicates can be tagged with a predicate hint that allows a join preference to be specified for just that one join. If the same join has more than one join condition with a local join preference, and if those hints are not the same value, then all local preferences are ignored for that join. Local join preferences do not affect the join order chosen by the optimizer.

The following example requests a hash join:

```
AND (T.X = 10 * R.X, 'J:4')
```

For additional information, see “User-supplied condition hints” in Chapter 3, “SQL Language Elements” in the *Sybase IQ Reference Manual*.

4.8.7 WORD index and CONTAINS predicate support for character large object (CLOB) columns [CR 415547]

Sybase IQ 12.7 ESD #3 provides support for WORD (WD) indexes on LONG VARCHAR character large object (CLOB) columns. A WORD index is also called a Containment index.

Sybase IQ 12.7 ESD #3 also includes support of the CONTAINS predicate on LONG VARCHAR (CLOB) columns with a WORD (WD) index. Formerly, the CONTAINS predicate was supported only on CHAR and VARCHAR columns with a WD index.

Users must be specifically licensed to use the Large Objects Management functionality. For details on the Large Objects Management option, see *Large Objects Management in Sybase IQ*.

WD index support

The following enhancements support the WORD (WD) index on LONG VARCHAR (CLOB) columns:

- Sybase Central Java allows you to create a WD index on columns of CHAR, VARCHAR, and LONG VARCHAR data types.
- The widest column supported by the WD index increased from 32767 bytes up to the maximum width for a LOB column. (The maximum length is equal to 4GB multiplied by the database page size.)

The maximum word length supported by Sybase IQ is still 255 bytes.

- All sp_iqcheckdb options for WD indexes over CHAR and VARCHAR columns are also supported for LONG VARCHAR (CLOB) columns, specifically, allocation, check, verify, and repair modes.
- The sp_iqrebuildindex stored procedure supports rebuilding a WD index over a LONG VARCHAR (CLOB) column.

Note that Chinese text or documents in a binary format still require ETL pre-processing to locate and transform the words into a form that can be parsed by the WD index.

See “Documentation updates and clarifications” on page 100 for documentation changes related to this feature.

CONTAINS predicate support

Using the CONTAINS predicate, you can now search for string constants of maximum length 255 characters on a LONG VARCHAR (CLOB) column, in addition to CHAR and VARCHAR columns.

Note that the CONTAINS predicate is not supported on LONG BINARY (BLOB) columns. If you attempt to search for a string in a LONG BINARY column using a CONTAINS predicate, an error is returned.

For more information on CONTAINS string searches, see “CONTAINS conditions” in Chapter 3, “SQL Language Elements” of the *Sybase IQ Reference Manual*.

See “Documentation updates and clarifications” on page 100 for documentation changes related to this feature.

4.8.8 New database option **WD_DELETE_METHOD** [CR 415547]

Sybase IQ 12.7 ESD #3 provides the ability to specify the algorithm used during a delete operation from a WD index using the new database option **WD_DELETE_METHOD**.

WD_DELETE_METHOD option

Function	Specifies the algorithm used during a delete in a WD index.
Allowed values	0 – 3
Default	0
Scope	DBA permissions are not required to set this option. Can be set temporary, for an individual connection, or for the PUBLIC group. Takes effect immediately.
Description	This option chooses the algorithm used during a delete operation in a WD index. When this option is not set or is set to 0, the delete method is selected by the cost model. The cost model considers the CPU related costs as well as I/O related costs in selecting the appropriate delete algorithm. The cost model takes into account: • Rows deleted • Index size • Width of index data type • Cardinality of index data • Available temporary cache • Machine related I/O and CPU characteristics • Available CPUs and threads

Allowed values for WD_DELETE_METHOD:

- 0: The delete method is selected by the cost model. Cost model only selects either mid or large method for deletion.
- 1: Forces small method for deletion. Small method is useful when the number of rows being deleted is a very small percentage of the total number of rows in the table.
- 2: Forces large method for deletion. Large method is useful when the number of rows being deleted is a high percentage of the total number of rows in the table.
- 3: Forces mid method for deletion. Mid method is a variation of the small method and is generally faster than the small method.

Example

The following statement forces the large method for deletion from a WD index:

```
SET TEMPORARY OPTION WD_DELETE_METHOD = 2
```

See also

For more details about these delete methods, see “New database option WD_DELETE_METHOD [CR 415547]” on page 173.

4.8.9 LIKE predicate for character large object (CLOB) columns [CR 415546]

Sybase IQ 12.7 ESD #3 includes support of the LIKE predicate on LONG VARCHAR (CLOB) columns. Formerly, LIKE predicates were supported only on CHAR and VARCHAR columns.

Using the LIKE predicate, you can now search for a pattern on a LONG VARCHAR column. All patterns of 126 characters or less are supported. Patterns of length greater than 254 characters are not supported. Some patterns of length between 127 and 254 characters are supported, depending on the contents of the pattern.

Note that LIKE predicates are not supported on LONG BINARY (BLOB) columns. If you attempt to search for a pattern in a LONG BINARY column using a LIKE predicate, the error “Invalid data type comparison in predicate...” is returned.

4.8.10 String search functions for large object columns [CR 415543]

Sybase IQ 12.7 ESD #3 provides string search capabilities on large objects. Sybase IQ now includes support of string search functions for LONG VARCHAR (CLOB) and LONG BINARY (BLOB) data. The string search functions CHARINDEX and LOCATE have been enhanced to allow a LONG VARCHAR or a LONG BINARY column as the argument for the string expression to be searched and now return a 64 bit value for the position of the string in the LONG VARCHAR or LONG BINARY column.

The PATINDEX function has been enhanced to allow a LONG VARCHAR column as the argument for the string expression to be searched and now returns a 64 bit value for the position of the string in the LONG VARCHAR column.

Users must be specifically licensed to use the Large Objects Management functionality. For details on the Large Objects Management option, see *Large Objects Management in Sybase IQ*.

See “Documentation updates and clarifications” on page 100 for documentation changes related to this feature.

The following sections describe the string search functions that support LONG VARCHAR and LONG BINARY data types.

CHARINDEX function

The CHARINDEX function returns a 64 bit signed integer containing the position of the first occurrence of the specified string in a LONG VARCHAR column. CHARINDEX returns a 32 bit signed integer position for CHAR and VARCHAR columns.

The CHARINDEX function also supports searching LONG BINARY columns.

Syntax:

CHARINDEX(*string-expression*, *long-varchar-column*)

Parameters:

string-expression The string for which you are searching. This string is limited to 255 bytes.

long-varchar-column The name of the LONG VARCHAR column.

Usage:

- All the positions or offsets, returned or specified, in the CHARINDEX function are always character offsets and may be different from the byte offset for multibyte data.

- If the LONG VARCHAR cell being searched contains more than one instance of the string, CHARINDEX returns only the position of the first instance.
- If the column does not contain the string, the CHARINDEX function returns zero (0).
- Searching for a string longer than 255 bytes returns NULL.
- Searching for a zero-length string returns 1.
- If any of the arguments is NULL, the result is NULL.

For a full description and an example of the CHARINDEX function, see “CHARINDEX function [String]” in Chapter 5, “SQL Functions” of the *Sybase IQ Reference Manual*.

LOCATE function

The LOCATE function returns a 64 bit signed integer containing the position of the specified string in a LONG VARCHAR column. LOCATE returns a 32 bit signed integer position for CHAR and VARCHAR columns.

The LOCATE function also supports searching LONG BINARY columns.

Syntax:

```
LOCATE( long-varchar-column, string-expression
        [, numeric-expression ] )
```

Parameters:

long-varchar-column The name of the LONG VARCHAR column to search.

string-expression The string for which you are searching. This string is limited to 255 bytes.

numeric-expression The character position or offset at which to begin the search in the string. The *numeric-expression* is a 64 bit signed integer for LONG VARCHAR and LONG BINARY columns and is a 32 bit signed integer for CHAR, VARCHAR, and BINARY columns. The first character is position 1. If the starting offset is negative, LOCATE returns the last matching string offset, rather than the first. A negative offset indicates how much of the end of the string to exclude from the search. The number of characters excluded is calculated as $(-1 * \text{offset}) - 1$.

Usage:

- All the positions or offsets, returned or specified, in the LOCATE function are always character offsets and may be different from the byte offset for multibyte data.

- If the LONG VARCHAR cell being searched contains more than one instance of the string:
 - If *numeric-expression* is specified, LOCATE starts the search at that offset in the string.
 - If *numeric-expression* is not specified, LOCATE returns only the position of the first instance.
- If the column does not contain the string, the LOCATE function returns zero (0).
- Searching for a string longer than 255 bytes returns NULL.
- Searching for a zero-length string returns 1.
- If any of the arguments is NULL, the result is NULL.

For a full description and examples of the LOCATE function, see “LOCATE function [String]” in Chapter 5, “SQL Functions” of the *Sybase IQ Reference Manual*.

PATINDEX function

The PATINDEX function returns a 64 bit unsigned integer containing the position of the first occurrence of the specified pattern in a LONG VARCHAR column. PATINDEX returns a 32 bit unsigned integer position for CHAR and VARCHAR columns.

The PATINDEX function does not support searching LONG BINARY columns.

Syntax:

PATINDEX('%pattern%', *long-varchar-column*)

Parameters:

pattern The pattern for which you are searching. This string is limited to 126 bytes. If the leading percent wildcard is omitted, PATINDEX returns one (1) if the pattern occurs at the beginning of the column value, and zero (0) if the pattern does not occur at the beginning of the column value. Similarly, if the trailing percent wildcard is omitted, the pattern should occur at the end of the column value. The pattern uses the same wildcards as the LIKE comparison.

long-varchar-column The name of the LONG VARCHAR column.

Usage:

- All the positions or offsets, returned or specified, in the PATINDEX function are always character offsets and may be different from the byte offset for multibyte data.

- If the LONG VARCHAR cell being searched contains more than one instance of the string pattern, PATINDEX returns only the position of the first instance.
- If the column does not contain the pattern, the PATINDEX function returns zero (0).
- Searching for a pattern longer than 126 bytes returns NULL.
- Searching for a zero-length pattern returns 1.
- If any of the arguments is NULL, the result is zero (0).

For a full description and examples of the PATINDEX function, see “PATINDEX function [String]” in Chapter 5, “SQL Functions” of the *Sybase IQ Reference Manual*.

For more information on LIKE comparisons, see “LIKE conditions” in Chapter 3, “SQL Language Elements” of the *Sybase IQ Reference Manual*.

For information on a known issue related to the PATINDEX function, see “PATINDEX returns 0 for zero-length search string [CR 475209]” on page 94.

4.8.11 Using query plans with Interactive SQL for Java [338070]

In Sybase IQ 12.7 ESD #3, authorized users can display query plans in the Java-based Interactive SQL (dbisql) plan window. You can also save and print query plans from dbisql instead of accessing the *.iqmsg* file or query plan files on the server.

Sybase IQ supports SQL functions GRAPHICAL_PLAN and HTML_PLAN, which return IQ query plans in XML and HTML format, respectively, as a string result set. The behavior of these functions is controlled by database options QUERY_PLAN_TEXT_ACCESS and QUERY_PLAN_TEXT_CACHING.

Users can view query plans from the dbisql plan window in the following ways:

- Execute the query and open the plan window. Depending on the plan type you selected from the Plan option (Tools > Options > Plan), the appropriate plan displays in the plan window.

The IQ query plan displays only if the GRAPHICAL_PLAN option is selected. Other plans return the error message, “Plan type is not supported.”

- Enter the query in the SQL statements window and select from the menu SQL > Get Plan. Depending on the plan type you selected from the Plan option (Tools > Options > Plan), the appropriate plan displays in the plan window.

The IQ query plan displays only if the GRAPHICAL_PLAN option is selected. Other plans return the error message, “Plan type is not supported.”

- Use the SQL functions, GRAPHICAL_PLAN and HTML_PLAN, to return the query plan as a string result.

To access query plans, use the SQL functions, GRAPHICAL_PLAN and HTML_PLAN, for the following queries: SELECT, UPDATE, DELETE, INSERT SELECT, and SELECT INTO.

To save the query plans from Interactive SQL, use GRAPHICAL_PLAN or HTML_PLAN to retrieve the query plan and save the output to a file using the OUTPUT statement.

See “GRAPHICAL_PLAN function [String]” on page 68 and “HTML_PLAN function [String]” on page 71 for details.

To view saved plans, select File > Open from the Interactive SQL client menu and navigate to the directory where you saved your plan. You can also print plans displayed on the plan window by selecting File > Print.

GRAPHICAL_PLAN function [String]

Function	Returns the graphical query plan to Interactive SQL in XML format string.
Syntax	GRAPHICAL_PLAN (<i>string-expression</i>)
Parameters	string-expression SQL statement for which the plan is to be generated. It is generally a SELECT statement, but it can be an UPDATE or DELETE, INSERT SELECT, and SELECT INTO statement.

If the user does not provide an argument to the GRAPHICAL_PLAN function, the query plan is returned to the user from the cache. If there is no query plan in the cache, then the message “Plan not available” is returned.

The behavior of GRAPHICAL_PLAN function is controlled by database options QUERY_PLAN_TEXT_ACCESS and QUERY_PLAN_TEXT_CACHING. If QUERY_PLAN_TEXT_ACCESS is OFF (the default), the following message displays in the error dialog box:

```
Plan not available. The database option
QUERY_PLAN_TEXT_ACCESS is OFF
```

If a user needs access to the plan, the DBA must set option `QUERY_PLAN_TEXT_ACCESS ON` for that user.

If `QUERY_PLAN_TEXT_ACCESS` is `ON`, and the query plan for the string expression is available in the cache maintained on the server, the query plan is returned to the user from the cache.

If the query plan is not available in the cache and the user is authorized to view plans on the client, then a query plan with optimizer estimates (query plan with `NOEXEC` option `ON`) is generated and displays on the dbisql client plan window.

Note Sybase IQ does not support `NOEXEC` plan generation for `SELECT`, `UPDATE`, `DELETE`, `INSERT SELECT`, and `SELECT INTO` queries.

When a user requests a query plan that has not yet been executed, the query plan will not be available in the cache. Instead, a query plan with optimizer estimates is returned without `QUERY_PLAN_AFTER_RUN` statistics.

Query plans for stored procedures are not accessible using the `GRAPHICAL_PLAN` function. The error message, “The plan for this statement cannot be retrieved” is returned.

Users can view the query plan for cursors opened for IQ queries. A cursor is declared and opened using `DECLARE CURSOR` and `OPEN CURSOR` commands. Use the following query to obtain the query plan for the most recently opened cursor:

```
SELECT GRAPHICAL_PLAN ( );
```

With the `QUERY_PLAN_AFTER_RUN` option `OFF`, the plan can be displayed after `OPEN CURSOR` or `CLOSE CURSOR`. However, if `QUERY_PLAN_AFTER_RUN` is `ON`, then `CLOSE CURSOR` must be executed before the user requests the plan.

When Interactive SQL users select the option Short plan, Long plan, or Show UltraLite Plan from the Plan tab under Tools > Options, the following message displays on the plan window.

```
Plan not available. Plan type is not supported for  
Sybase IQ queries.
```

The user must change the plan option to Graphical plan before requesting the plan for an IQ query. UltraLite plans are not supported.

When Interactive SQL users request plans for UPDATE, DELETE, SELECT INTO, and INSERT SELECT queries, the NOEXEC plan is not supported. Set AUTO_REFETCH option OFF to view the plan for the last executed UPDATE, DELETE, SELECT INTO, and INSERT SELECT queries. To access the query plan, explicitly execute the query first and then request the plan.

Examples

The following SQL example passes a SELECT statement as a string parameter and returns the plan for executing the query. It saves the plan in the file *gplan.xml*.

Note If you use the OUTPUT statement's HEXADECIMAL clause set to ASIS to get formatted plan output, the values of characters are written without any escaping — even if the value contains control characters. ASIS is useful for text that contains formatting characters such as tabs or carriage returns.

```
SELECT GRAPHICAL_PLAN ('SELECT * FROM t1');
OUTPUT to 'c:\gplan.xml' hexadecimal ASIS quote '';
```

The following SQL example returns the query plan from the cache, if available.

```
SELECT GRAPHICAL_PLAN ( );
```

Standards and compatibility

- **SQL92** Vendor extension.
- **SQL99** SQL/foundation feature outside of core SQL.
- **Sybase** Not supported by Adaptive Server Enterprise.

See also

“HTML_PLAN function [String]” on page 71

“QUERY_PLAN_TEXT_ACCESS option” on page 73

“QUERY_PLAN_TEXT_CACHING option” on page 73

“AUTO_REFETCH option [DBISQL],” “NOEXEC option,”

“QUERY_PLAN_AFTER_RUN option,” and “OUTPUT statement” in the *Sybase IQ Reference Manual*

PLAN function [Miscellaneous], EXPLANATION function, [Miscellaneous], GRAPHICAL_ULPLAN function [Miscellaneous], LONG_ULPLAN function [Miscellaneous], and SHORT_ULPLAN function [Miscellaneous] in *Adaptive Server Anywhere SQL Reference Manual*.

HTML_PLAN function [String]

Function Returns query plans in HTML format string.

Syntax **HTML_PLAN** (*string-expression*)

Parameters **string-expression** SQL statement for which the plan is to be generated. It is primarily a SELECT statement but can be an UPDATE or DELETE statement.

If the user does not provide an argument to the HTML_PLAN function, the query plan is returned to the user from the cache. If there is no query plan in the cache, then message “No plan available” is returned.

The behavior of the HTML_PLAN function is controlled by database options QUERY_PLAN_TEXT_ACCESS and QUERY_PLAN_TEXT_CACHING. If QUERY_PLAN_TEXT_ACCESS is OFF (the default), the following message displays in the error dialog box:

```
Plan not available. The database option
QUERY_PLAN_TEXT_ACCESS is OFF
```

If QUERY_PLAN_TEXT_ACCESS is ON, and the query plan for the string expression is available in the cache maintained on the server, the query plan is returned to the user from the cache.

The HTML_PLAN function can be used to return query plans to Interactive SQL using SELECT, UPDATE, DELETE, INSERT SELECT, and SELECT INTO.

Users can view the query plan for cursors opened for IQ queries. A cursor can be declared and opened using DECLARE CURSOR and OPEN CURSOR commands. Use the following query to obtain the query plan for the most recently opened cursor:

```
SELECT HTML_PLAN ( );
```

With QUERY_PLAN_AFTER_RUN option OFF, the plan can be displayed after OPEN CURSOR or CLOSE CURSOR. However, if QUERY_PLAN_AFTER_RUN is ON, then CLOSE CURSOR must be executed before the user requests the plan.

When the Interactive SQL user selects the option Short plan, Long plan, or Show Ultralite Plan from the Plan tab under Tools > Options, the following message displays on the plan window.

```
Plan not available. Plan type is not supported for
Sybase IQ queries.
```

The user must change the plan option to Graphical plan before requesting the plan for an IQ query. UltraLite plans are not supported.

When Interactive SQL users request plans for UPDATE, DELETE, SELECT INTO, and INSERT SELECT queries, the NOEXEC plan is not supported. Set AUTO_REFETCH option OFF to view the plan for the last executed UPDATE, DELETE, SELECT INTO, and INSERT SELECT queries. To access the query plan, explicitly execute the query first and then request the plan.

When the user requests an HTML_PLAN for a SQL Anywhere query or for an OMNI/CIS decomposed query, the following message is returned:

```
No plan. HTML_PLAN function is not supported for this
type of statement or database.
```

Example

The following SQL example passes a SELECT statement as a string parameter and returns the HTML plan for executing the query. It saves the plan in the file *hplan.html*.

```
SELECT HTML_PLAN ('SELECT * FROM t1');OUTPUT to
'c:\hplan.html';
```

The following SQL example returns the HTML query plan from the cache, if available.

```
SELECT HTML_PLAN ( );
```

Standards and compatibility

- **SQL92** Vendor extension.
- **SQL99** SQL/foundation feature outside of core SQL.
- **Sybase** Not supported by Adaptive Server Enterprise.

See also

“GRAPHICAL_PLAN function [String]” on page 68

“QUERY_PLAN_TEXT_ACCESS option” on page 73

“QUERY_PLAN_TEXT_CACHING option” on page 73

“AUTO_REFECT option [DBISQL],” “NOEXEC option,”

“QUERY_PLAN_AFTER_RUN option,” and “OUTPUT statement” in the *Sybase IQ Reference Manual*

PLAN function [Miscellaneous], EXPLANATION function, [Miscellaneous], GRAPHICAL_ULPLAN function [Miscellaneous], LONG_ULPLAN function [Miscellaneous], and SHORT_ULPLAN function [Miscellaneous] in *Adaptive Server Anywhere SQL Reference Manual*.

QUERY_PLAN_TEXT_ACCESS option

Function	Enables or prevents users from accessing query plans from the Interactive SQL (dbisql) client or from using SQL functions to get plans.
Allowed values	ON, OFF
Default	OFF
Scope	DBA permissions are required to modify this option. Can be set temporary for an individual connection or for the PUBLIC group. Takes effect immediately.
Description	When QUERY_PLAN_TEXT_ACCESS option is ON, users can view, save, and print query plans from the dbisql client. When the option is OFF, query plans are not cached, and other query plan-related database options have no effect on the query plan display from the dbisql client. The following error message displays: <pre>No plan available. The database option QUERY_PLAN_TEXT_ACCESS is OFF.</pre>
See also	“GRAPHICAL_PLAN function [String]” on page 68“ “HTML_PLAN function [String]” on page 71 “QUERY_PLAN_TEXT_CACHING option” on page 73 “QUERY_DETAIL option,” “QUERY_PLAN_AS_HTML option,” “QUERY_PLAN_AFTER_RUN option,” and “OUTPUT statement” in the <i>Sybase IQ Reference Manual</i>. PLAN function [Miscellaneous], EXPLANATION function, [Miscellaneous], GRAPHICAL_ULPLAN function [Miscellaneous], LONG_ULPLAN function [Miscellaneous], and SHORT_ULPLAN function [Miscellaneous] in <i>Adaptive Server Anywhere SQL Reference Manual</i>.

QUERY_PLAN_TEXT_CACHING option

Function	Allow users to specify whether or not Sybase IQ generates and caches IQ plans for queries executed by the user.
Allowed values	ON, OFF
Default	OFF
Scope	DBA permissions are not required to modify this option. Can be set temporary for an individual connection. Takes effect immediately.

Description	IQ query plans vary in size and can become very large for complex queries. Caching plans for display on the dbisql client can have high resource requirements. The QUERY_PLAN_TEXT_CACHING option gives users a mechanism to control resources for caching plans. If the user turns this option OFF, the query plan is not cached for that user connection. If QUERY_PLAN_TEXT_ACCESS is turned OFF for a user, the query plan is not cached for the connections from that user, no matter how QUERY_PLAN_TEXT_CACHING is set.
See also	“GRAPHICAL_PLAN function [String]” on page 68“ “HTML_PLAN function [String]” on page 71 “QUERY_PLAN_TEXT_ACCESS option” on page 73 “QUERY_DETAIL option,” “QUERY_PLAN_AS_HTML option,” “QUERY_PLAN_AFTER_RUN option,” and “OUTPUT statement” in the <i>Sybase IQ Reference Manual</i>. PLAN function [Miscellaneous], EXPLANATION function, [Miscellaneous], GRAPHICAL_ULPLAN function [Miscellaneous], LONG_ULPLAN function [Miscellaneous], and SHORT_ULPLAN function [Miscellaneous] in <i>Adaptive Server Anywhere SQL Reference Manual</i>.

4.9 Security enhancements

This section contains new features and changes related to security.

4.9.1 Server support of TDS password encryption [CR 484702]

As of Sybase IQ 12.7 ESD #4, the Sybase IQ server supports TDS password encryption. The IQ server accepts a connection with an encrypted password sent by the client. In order to use this feature, you must use Open Client 15.0 ESD #7 or later.

4.9.2 Login management changes in Sybase Central [CR 442651, CR 448480]

The following Sybase Central changes for ESD #2 affect login management. (For an overview of this functionality, see Chapter 2 of the *Sybase IQ Reference Manual* and “Managing User IDs and Permissions” in Chapter 12 of the *Sybase IQ System Administration Guide*.)

For an IQ database, the Allowed to connect check box in the User Property Sheet, General tab is renamed to Enable Password. It continues to function in the General tab as before: checking the Enable Password checkbooks sets the password to non-null.

If login management is disabled, the Login Management tab is displayed but not enabled, while viewing the user/group properties.

If login management is enabled:

- Create Group and Create User wizards contain an additional page that allows the user to lock the new user/group being created out of the current database. The page also lets you lock out the user or group on specific servers in a multiplex.
- The Login Management tab in the user/group properties sheet is enabled based on two factors:
 - Passwords must be enabled by checking the “Enable Password” checkbooks in the General tab.
 - Login management is enabled via the Database Property sheet Login Management tab.

4.9.3 Sample code for password verification function [CR 440216]

Added for Sybase IQ 12.7 ESD #1, this example defines a number of procedures and functions which together implement advanced password rules, including requiring certain types of characters in the password, disallowing password reuse, and expiring passwords. These procedures and functions are called by the server using the VERIFY_PASSWORD_FUNCTION option and the LOGIN_PROCEDURE option when a User ID is created or a password is changed, and when a user connects. The application can call the procedure specified by the POST_LOGIN_PROCEDURE option to report that the password should be changed before it expires.

For more information on the POST_LOGIN_PROCEDURE option, see “POST_LOGIN_PROCEDURE option” on page 84.

For more information on the VERIFY_PASSWORD_FUNCTION and the LOGIN_PROCEDURE options, see Chapter 2, “Database Options” in the *Sybase IQ Reference Manual*.

Sample code

```

-- only DBA should have permissions on this table

CREATE TABLE DBA.t_pwd_history(
pk INT DEFAULT AUTOINCREMENT PRIMARY KEY,
change_date TIMESTAMP DEFAULT CURRENT TIMESTAMP,
    -- when pwd set
grace_date DATE, -- a day after password expires to
    -- allow user to log in
user_name CHAR(128), -- the user whose password is set
pwd_hash CHAR(32) ); -- hash of password value to detect
    -- duplicate passwords

-- called on every GRANT CONNECT TO ... IDENTIFIED BY...
-- statement to verify that the password conforms to
-- the password rules

CREATE FUNCTION DBA.f_verify_pwd( uid VARCHAR(128),
                                new_pwd VARCHAR(255) )
RETURNS VARCHAR(255)
BEGIN
-- a table with one row per character in new_pwd
DECLARE local temporary table pwd_chars(
pos INT PRIMARY KEY, -- index of c in new_pwd
c CHAR( 1 CHAR ) ); -- character
-- new_pwd with non-alpha characters removed
DECLARE pwd_alpha_only CHAR(255);
DECLARE num_lower_chars INT;

-- enforce minimum length (can also be done with
-- min_password_length option)
IF length( new_pwd ) < 6 THEN
RETURN 'password must be at least 6 characters long';
END IF;

-- break new_pwd into one row per character
INSERT INTO pwd_chars SELECT row_num, substr(
new_pwd, row_num, 1 )
FROM dbo.RowGenerator
WHERE row_num <= length( new_pwd );

-- copy of new_pwd containing alpha-only characters
SELECT LIST( c, '' ORDER BY pos ) INTO pwd_alpha_only
FROM pwd_chars
WHERE c BETWEEN 'a' AND 'z' OR c BETWEEN 'A' AND 'Z';

```

```
-- number of lower case characters IN new_pwd
SELECT COUNT(*) INTO num_lower_chars
FROM pwd_chars
WHERE CAST( c AS BINARY ) BETWEEN 'a' AND 'z';

-- enforce rules based on characters contained in
-- new_pwd
IF ( SELECT count(*) FROM pwd_chars
WHERE c BETWEEN '0' AND '9' ) < 1 THEN
RETURN 'password must contain at least one numeric
digit';
ELSEIF length( pwd_alpha_only ) < 2 THEN
RETURN 'password must contain at least two letters';
ELSEIF num_lower_chars = 0
OR length( pwd_alpha_only ) - num_lower_chars = 0
THEN
RETURN 'password must contain both upper- and
lowercase characters';
END IF;

-- not the same as any user name
-- (this could be modified to check against
-- a disallowed words table)
IF EXISTS( SELECT * FROM SYS.SYSUSER
WHERE lower( user_name ) IN ( lower( pwd_alpha_only ),
lower( new_pwd
) ) ) THEN
RETURN 'password or only alphabetic characters in
password ' || 'must not match any user name';
END IF;

-- not the same as any previous password for this user
IF EXISTS( SELECT * FROM t_pwd_history
WHERE user_name = uid
AND pwd_hash = hash( uid || new_pwd, 'md5'
) ) THEN
RETURN 'previous passwords cannot be reused';
END IF;

-- save the new password
INSERT INTO t_pwd_history( user_name, pwd_hash )
VALUES( uid, hash( uid || new_pwd, 'md5' ) );
RETURN( NULL );
END;

ALTER FUNCTION DBA.f_verify_pwd SET HIDDEN;
```

```
GRANT EXECUTE ON DBA.f_verify_pwd TO PUBLIC;
SET OPTION PUBLIC.verify_password_function =
'DBA.f_verify_pwd';

-- called on every connection to check for password
-- expiry

CREATE PROCEDURE DBA.p_login_check()
BEGIN
DECLARE uid CHAR(128);
DECLARE INVALID_LOGON EXCEPTION FOR SQLSTATE '28000';
DECLARE last_pwd_change DATE;
DECLARE grace_date DATE;
DECLARE is_dba CHAR;
DECLARE msg_str CHAR(255);
SET uid = connection_property( 'Userid' );
IF ( EXISTS( SELECT * FROM t_pwd_history
WHERE user_name = uid
) ) THEN
SELECT FIRST t.change_date, t.grace_date
INTO last_pwd_change, grace_date
FROM t_pwd_history t WHERE t.user_name = uid
ORDER BY t.change_date DESC;
END IF;

IF last_pwd_change IS NULL THEN
-- no password change in t_pwd_history, so create one
-- now
INSERT INTO t_pwd_history( user_name, pwd_hash )
VALUES( uid, 'unknown' );
COMMIT WORK;
ELSE
IF EXISTS( SELECT * FROM SYS.SYSUSERAUTHORITY a,
SYS.SYSUSER u
WHERE u.user_name = uid AND u.user_id =
a.user_id AND
a.auth = 'DBA' ) THEN
SET is_dba = 'Y';
ELSE
SET is_dba = 'N';
END IF;

-- remove any locks on t_pwd_history and
-- SYSUSERAUTHORITY
ROLLBACK WORK;
```

```

-- check if last password change was over five
-- months ago
IF CURRENT DATE > dateadd( month, 5, last_pwd_change )
THEN
-- never expire DBA accounts so that the database
-- does not get locked out by all users
IF CURRENT DATE < dateadd( month, 6, last_pwd_change
) OR
is_dba = 'Y' OR
( grace_date IS NOT NULL AND grace_date = CURRENT
DATE ) THEN
SET msg_str = 'The password for user ' || uid ||
' expires on ' ||
CAST( dateadd( month, 6, last_pwd_change )
AS DATE ) || '. Please change it before it expires.';
MESSAGE msg_str;

-- The post_login_procedure option is set to
-- p_post_login_check, which will cause a dialog to be
-- displayed notifying the user their password will
-- expire soon. dbisql and dbisqlc will display this
-- dialog, and user applications can call the
-- post_login_procedure and display this dialog.
-- May want to use xp_send_mail to notify user and/or
-- administrator.

ELSEIF grace_date IS NULL THEN

-- Allow one grace login day. The first login on
-- the grace day fails to ensure the user knows
-- their password has expired.

UPDATE t_pwd_history t SET t.grace_date = CURRENT DATE
WHERE t.grace_date IS NULL AND t.user_name = uid;
COMMIT WORK;
SET msg_str = 'The password for user ' || uid || '
has expired, but future logins will ' ||
'be allowed today only so that the password ' ||
'can be changed.';
MESSAGE msg_str;
RAISERROR 28000 msg_str;
RETURN;
ELSE
SET msg_str = 'The password for user ' || uid ||
' has expired and must be reset by your DBA.';
MESSAGE msg_str;

```

```
-- may want to use xp_send_mail to notify administrator
RAISERROR 28000 msg_str;
RETURN;
END IF;
END IF;
END IF;

CALL sp_login_environment;
END;

GRANT EXECUTE ON DBA.p_login_check TO PUBLIC;
SET OPTION PUBLIC.login_procedure =
'DBA.p_login_check';

-- called by dbisql, dbisqlc and some user applications
-- on every successful connection to check for warnings
-- which should be displayed

CREATE PROCEDURE DBA.p_post_login_check()
RESULT( warning_text VARCHAR(255), warning_action INT )
BEGIN
DECLARE uid CHAR(128);
DECLARE last_pwd_change DATE;
DECLARE warning_text CHAR(255);
DECLARE warning_action INT;
SET uid = connection_property( 'Userid' );
SELECT FIRST t.change_date
INTO last_pwd_change
FROM t_pwd_history t WHERE t.user_name = uid
ORDER BY t.change_date DESC;
IF CURRENT DATE > dateadd( month, 5, last_pwd_change )
THEN
SET warning_text = 'Your password expires on ' ||
CAST( dateadd( month, 6, last_pwd_change )
AS DATE ) || '. Please change it before it expires.';
SET warning_action = 1;
ELSE
-- There is no warning
SET warning_text = NULL;
SET warning_action = 0;
END IF;
-- return the warning (if any) through this result set
SELECT warning_text, warning_action;
END;
```

```
GRANT EXECUTE ON DBA.p_post_login_check TO PUBLIC;
SET OPTION PUBLIC.post_login_procedure =
'DBA.p_post_login_check';
```

4.9.4 sp_iqmpxcreatewriterlogin procedure [CR 421148]

Function	An ESD #1 enhancement includes a stored procedure that adds an external login for DBA to write server.
Syntax1	call sp_iqmpxcreatewriterlogin ('remote_user', 'remote_password')
Syntax2	sp_iqmpxcreatewriterlogin 'remote_user', 'remote_password'
Syntax3	sp_iqmpxcreatewriterlogin remote_user, remote_password
Permissions	DBA authority required.
Usage	remote_user User's login name on the write server, which acts as the external login of DBA from the query server. remote_password User's password on the write server, which acts as the external login of DBA from the query server.
Description	In order to propagate passwords from a query server to a write server and subsequently other nodes, Sybase IQ calls an RPC on the write server from the query server. The RPC uses user 'DBA' account to establish a connection to the write server. If you modify the user 'DBA' password on the query server or set the 'DBA' password to NULL, the RPC can no longer connect to the write server. You must execute sp_iqmpxcreatewriterlogin on the reader so that it can use an alternate account for the remote connection to the write server. The created external login is local to each query server and is deleted after the Multiplex synchronization unless a post synchronization script ('DBA.sp_mpxcfg_' + @@servername) is defined and called to restore it.
Errors	The following errors might occur. Causes are listed after each error. "Permission denied: you do not have permission to execute the procedure "sp_iqmpxcreatewriterlogin" Cause: A user without DBA privilege tried to execute the sp_iqmpxcreatewriterlogin procedure. RAISERROR executed: The stored procedure is not called from a query server. Sybase error code=17823, SQLState="WW012"

Cause: This procedure can only be called from a query server.

```
"RAISERROR executed: Userid "iqremoter" is not found or
could not be set as a login for write server. Sybase
error code=17824, SQLState="WW012"
```

Causes: The passed in value to procedure parameter *remote_password* is null, the remote user does not exist or not have the DBA privilege, or the DBA's password is null.

```
"ASA Error -110: Item 'DBA' already exists Sybase error
code = 12006, SQLState = "52010"
```

Cause: The external login already exists for DBA to writer server. You can drop it first by calling *sp_iqmpxdropwriterlogin*, and then call the *sp_iqmpxcreatewriterlogin* procedure.

4.9.5 *sp_iqmpxdropwriterlogin* procedure [CR 421148]

Function	An ESD #1 enhancement includes a stored procedure that drops existing external login for DBA to writer server.
Syntax1	call <i>sp_iqmpxdropwriterlogin</i>
Syntax2	<i>sp_iqmpxdropwriterlogin</i>
Permissions	DBA authority required.
Description	Used to drop an existing external login for DBA to write server. To change the external login, call <i>sp_iqmpxcreatewriterlogin</i> before calling the <i>sp_iqmpxcreatewriterlogin</i> procedure.
Errors	The following errors might occur. Causes are listed after each error. <pre>"Permission denied: you do not have permission to execute the procedure "sp_iqmpxdropwriterlogin"</pre> Cause: A user without DBA privilege tried to execute the <i>sp_iqmpxdropwriterlogin</i> procedure. <pre>"RAISERROR executed: The stored procedure is not called from a query server. Sybase error code=17823, SQLState="WW012"</pre> Cause: This procedure can only be called from a query server. <pre>"ASA Error -712: External login for server 'SYBIQ_MpxRemoteWriter' could not be found Sybase Error code=11221, SQLState="W0013"</pre> Cause: The external login does not exist or has already been dropped.

4.9.6 Password expiration warning [CR 420080]

The new password expiration warning feature in Sybase IQ 12.7 ESD #1 provides a simple method of notifying users when a password is about to expire. A warning message is delivered to the client application after the user successfully logs in. The new system stored procedure, `sp_iq_process_post_login`, and the new database option, `POST_LOGIN_PROCEDURE`, provide this enhancement to the current IQ Login Management facility.

Users who log in using `dbisql`, `dbisqlc`, or the Sybase Central Plug-in for Sybase IQ now see a password expiration warning, when appropriate, if Sybase IQ Login Management is enabled.

sp_iq_process_post_login procedure

Function	Sends a warning message to the client application for a user whose password is about to expire.
Syntax	sp_iq_process_post_login
Permissions	None.
Usage	If Sybase IQ Login Management is enabled and a user logs in using <code>dbisql</code>, <code>dbisqlc</code>, or the Sybase Central Plug-in for Sybase IQ, the <code>sp_iq_process_post_login</code> stored procedure runs immediately after a successful login. This procedure runs automatically, as the database option <code>POST_LOGIN_PROCEDURE</code> is set equal to <code>sp_iq_process_post_login</code> by default. Setting the <code>POST_LOGIN_PROCEDURE</code> option requires DBA authority. The password expiration warning is sent to the client application for a user whose password expires in 10 days or less. The password expiration warning procedure applies only to users added by the <code>sp_iqaddlogin</code> procedure.
See also	“<code>POST_LOGIN_PROCEDURE</code> option” on page 84 “<code>sp_iq_process_login</code> procedure” in Chapter 10, “System Procedures” in the <i>Sybase IQ Reference Manual</i> “<code>LOGIN_PROCEDURE</code> option” in Chapter 2, “Database Options” in the <i>Sybase IQ Reference Manual</i> “Login procedures and IQ User Administration” in Chapter 3, “Sybase IQ Connections” in the <i>Sybase IQ System Administration Guide</i>

	“Managing IQ user accounts and connections” in Chapter 12, “Managing User IDs and Permissions” in the <i>Sybase IQ System Administration Guide</i>
Description	After a user successfully logs in, the client application calls the stored procedure specified by the database option <code>POST_LOGIN_PROCEDURE</code>. The default setting of the <code>POST_LOGIN_PROCEDURE</code> option is <code>DBA.sp_iq_process_post_login</code>. When Sybase IQ Login Management is enabled, <code>sp_iq_process_post_login</code> returns a password expiration warning message to the client application, if the password of the user is about to expire. When Sybase IQ Login Management is disabled, the password expiration warning is not issued. The <code>sp_iq_process_post_login</code> procedure supports the client applications <code>dbisql</code>, <code>dbisqlc</code>, and the IQ plug-in for Sybase Central. <code>dbisql</code>, <code>dbisqlc</code>, and the IQ plug-in for Sybase Central display the warning “Your password will expire in n days! Please change your password as soon as possible.” In addition, <code>dbisql</code> and the IQ plug-in for Sybase Central also display a Reset Password dialog with the warning message, a request that the user change the password, and text boxes to supply and confirm a new password.

POST_LOGIN_PROCEDURE option

Function	Specifies a login procedure whose result set contains messages that are displayed by the client application immediately after a user successfully logs in.
Allowed values	String
Default	<code>DBA.sp_iq_process_post_login</code>
Scope	Can be set for an individual connection or the PUBLIC group. Requires DBA permissions to set the option. Takes effect immediately.
Description	The default post login procedure, <code>sp_iq_process_post_login</code>, executes immediately after a user successfully logs in. If you have DBA authority, you can customize the post login actions by creating a new procedure and setting <code>POST_LOGIN_PROCEDURE</code> to call the new procedure. <i>Do not</i> edit <code>sp_iq_process_post_login</code>. The customized post login procedure must be created in every database you use. The post login procedure supports the client applications <code>dbisql</code>, <code>dbisqlc</code>, and the IQ plug-in for Sybase Central.
See also	“Sample code for password verification function [CR 440216]” on page 75 “<code>sp_iq_process_post_login</code> procedure” on page 83

“sp_iq_process_login procedure” in Chapter 10, “System Procedures” in the *Sybase IQ Reference Manual*

“LOGIN_PROCEDURE option” in Chapter 2, “Database Options” in the *Sybase IQ Reference Manual*

“Login procedures and IQ User Administration” in Chapter 3, “Sybase IQ Connections” in the *Sybase IQ System Administration Guide*

“Managing IQ user accounts and connections” in Chapter 12, “Managing User IDs and Permissions” in the *Sybase IQ System Administration Guide*

4.10 Performance enhancements

This section contains new features and changes related to performance enhancements.

4.10.1 Controlling file system buffering [CR 487762]

As of Sybase IQ 12.7 ESD #9, you can control file system buffering on Linux, Linux IBM, and IBM AIX platforms using the database options `OS_FILE_CACHE_BUFFERING` and `OS_FILE_CACHE_BUFFERING_TEMPDB`. This direct I/O functionality is supported in Linux kernel versions 2.6.x. Sun Solaris and Windows platforms already support direct I/O.

To enable direct I/O on Linux kernel version 2.6 and AIX, also set the environment variable `IQ_USE_DIRECTIO` to 1. Direct I/O is disabled by default in Linux kernel version 2.6 and AIX. `IQ_USE_DIRECTIO` has no effect on Sun Solaris and Windows.

Note Sybase IQ does not support direct I/O on Linux kernel version 2.4. If you set the `IQ_USE_DIRECTIO` environment variable on Linux kernel version 2.4, the Sybase IQ server does not start. The error “Error: Invalid Block I/O argument, maybe <pathname> is a directory, or it exceeds maximum file size limit for the platform, or trying to use Direct IO on unsupported OS” is reported.

In versions prior to Sybase IQ 12.7 ESD #9, the database option `OS_FILE_CACHE_BUFFERING` controls the file system buffering of both IQ Main dbspaces and IQ Temporary dbspaces. As of Sybase IQ 12.7 ESD #9, `OS_FILE_CACHE_BUFFERING` controls the file system buffering of IQ Main dbspaces only. The new database option `OS_FILE_CACHE_BUFFERING_TEMPDB` controls the file system buffering of IQ Temporary dbspaces.

To enable direct I/O for IQ Main dbspaces, set `OS_FILE_CACHE_BUFFERING` to OFF. To enable direct I/O for IQ Temporary dbspaces, set `OS_FILE_CACHE_BUFFERING_TEMPDB` to OFF. The default value of both options is OFF.

`OS_FILE_CACHE_BUFFERING_TEMPDB` option

Function	Controls the use of file system buffering for IQ Temporary dbspaces.
Allowed values	ON, OFF
Default	OFF
Scope	Can be set for the PUBLIC group only. You must shut down and restart the database for the change to take effect. DBA permissions are required to set this option.
Description	Setting <code>OS_FILE_CACHE_BUFFERING_TEMPDB</code> to OFF prevents file system buffering for IQ Temporary Store files. Turning off file system buffering saves a data copy from the file system buffer cache to the main IQ buffer cache. Usually this reduces paging caused by competition for memory between the IQ buffer manager and the file system buffer of the operating system. When <code>OS_FILE_CACHE_BUFFERING_TEMPDB</code> reduces paging, this option improves performance; however, if the IQ page size for the database is less than the block size of the file system (typically only in testing situations), performance decreases, especially during multiuser operation. Experiment with this option to determine the best setting for different conditions. You must restart the database for the new setting to take effect. This option has no effect on HP-UX and HP-UXi. In Linux, direct I/O is supported in kernel versions 2.6.x. <code>OS_FILE_CACHE_BUFFERING</code> controls file system buffering for IQ Main Store files.

See also

“OS_FILE_CACHE_BUFFERING option” in Chapter 2, “Database Options” in the *Sybase IQ Reference Manual*

Chapter 5, “Managing System Resources” in the *Sybase IQ Performance and Tuning Guide*

4.11 Sybase IQ documentation on CD

In Sybase IQ 12.7, a SyBooks CD replaces the Technical Library CD. The SyBooks browser software runs on Windows and Linux platforms. Users with non-Linux UNIX platforms must use Acrobat Reader to open PDF files on the SyBooks CD.

The Sybase Product Manuals Web site is an online version of the SyBooks CD that you can access using a standard Web browser. To access the Sybase Product Manuals Web site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

5. Product compatibilities

This section describes the compatibility of Sybase IQ with other products. Certification with client and partner products is an ongoing process. See “Sybase certifications on the Web” on page 178 for instructions on accessing the most current online certification information.

5.1 Sybase products

The following Sybase products are certified with Sybase IQ 12.7:

- Adaptive Server Anywhere
- Adaptive Server Enterprise
- InfoMaker
- jConnect for JDBC
- Open Client
- PowerBuilder
- PowerDesigner
- Replication Server

Refer to “Sybase certifications on the Web” on page 178 for instructions on accessing the latest online certification information.

For information about product interoperability and known issues, see:

- *Sybase IQ Installation and Configuration Guide* section “Updating SDK” in Chapter 2, “Installing Sybase IQ”, and section “Running client applications” in Chapter 5, “Configuring Sybase IQ.”
- *New Features in Sybase IQ 12.7* sections “Integer data type support changes,” and “SYBASE_OCS behavior changes.”
- *Sybase IQ Reference Manual* section “SYBASE_OCS environment variable” in Chapter 1, “File Locations and Installation Settings.”

5.2 EMC Symmetrix and Hitachi certification

EMC Symmetrix and Hitachi are certified with Sybase IQ and are supported on all platforms. For details about the EMC Symmetrix certification, refer to the white paper “Configuring and Performance Tuning EMC Symmetrix for Sybase IQ Multiplex,” available on the Sybase Web site.

6. Restrictions

Read this section! Your system may produce unexpected results if you ignore the restrictions and other instructions listed below.

The information in this section applies to version 12.7 servers and databases that have been upgraded from any earlier version of Sybase IQ, unless specified otherwise.

6.1 Do not use server startup switch -f [CR 610313]

Table 1-1 in Chapter 1, “Running the Database Server,” contains an error. Unless instructed by Sybase Technical Support, never use the start_asiq startup switch -f for IQ databases. The switch will cause database corruption, because the catalog database is not synchronized with the IQ store if the server did not shut down cleanly. Documentation for this switch will be removed from future releases. See also “Restrictions on forced recovery [CR 610313]” on page 166.

6.2 Limit on number of backup devices [CR 543249]

When using the BACKUP command of Sybase IQ, users can parallelize the operation to multiple devices by specifying multiple TO clauses. Backups created using a backup statement with more than 36 TO clauses create backup archives that may be unreliable. While these archives appear to be correctly written, they may contain invalid headers and cannot be restored. The Sybase IQ BACKUP command provides no diagnostics or alerts that indicate that the archives have been written incorrectly. There is no current method to test the validity of the archives without performing a restore using the archives.

Future versions of Sybase IQ will enforce a limit of 36 devices in a BACKUP command.

This alert affects all versions of Sybase IQ up to and including Sybase IQ 12.7 ESD #5.

This alert also affects all versions of Sybase Risk Analytics Platform and Sybase RAP - The Trading Edition™.

6.2.1 Corrective actions for existing backup archives

Sybase strongly recommends that customers who have created backup archives to more than 36 devices test the archives by restoring them to a test area. Do not attempt to restore the archive to a production database, unless this test has completed successfully. If the test restore fails (a "Tapes out of order" error is returned), create new backup archives using 36 or fewer TO clauses.

6.2.2 Corrective actions for future backups

Users must create future backups using 36 or fewer TO clauses in a BACKUP command, a limit that will be enforced in future versions of Sybase IQ.

Keep backup commands small. Do not go to extremes with the number of devices, as this will cause I/O and hardware contention with more devices. As a practical guideline, use roughly 1 device per core on the machine to saturate CPU usage. Use up to 2 devices per core on faster systems.

6.3 ODBC option SQL_QUERY_TIMEOUT unsupported [CR 475631]

The ODBC driver used by Sybase IQ does not support the SQL_QUERY_TIMEOUT option

6.4 Dbspaces on Network Attached Storage unsupported [CR 477968]

Sybase IQ does not support creating dbspaces on NAS (Network Attached Storage) devices.

6.5 String comparisons on encrypted text [CR 436993]

If data is case insensitive, or uses a collation other than ISO_BINENG, you must decrypt ciphertext columns in order to perform string comparisons.

When performing comparisons on strings, the distinction between equal and identical strings is important for many collations and depends on CASE option of CREATE DATABASE. In a database set to CASE RESPECT and the ISO_BINENG collation, the defaults for Sybase IQ, equality, and identity questions are resolved the same way.

Identical strings are always equal, but the reverse may not be true. Strings are only identical if they are represented using the same byte values. When data is case insensitive or uses a collation where multiple characters must be treated as equal, the distinction between equality and identity is significant. ISO1LATIN1 is such a collation.

For example, the strings 'ABC' and 'abc' in a case insensitive database are not identical but are equal. In a case sensitive database, they are neither identical nor equal.

The ciphertext created by the Sybase encryption functions preserves identity but not equality. In other words, the ciphertext for 'ABC' and 'abc' will never be equal.

6.6 Sybase Central Java startup script [CR 430515]

On UNIX and Linux platforms, always start Sybase Central Java using the startup script `scjview` (default location `$SYBASE/shared/sybccentral43`).

If you install Sybase IQ on the same UNIX or Linux system as other Sybase products that use Sybase Central, the Sybase Central directory contains two startup scripts for Sybase Central Java:

- `scjview` – installed by Sybase IQ.
- `scjview.sh` – installed by Adaptive Server Enterprise, Replication Server, and other Sybase 15.0 products.

Do not use *scjview.sh* with the IQ plug-in. Using *scjview.sh* returns errors on startup and when connecting to IQ databases because libraries required by the IQ plug-in are not loaded.

6.7 Operations restriction on HP systems

On HP systems, Sybase IQ writes server messages to */var/adm/syslog/syslog.log* (the system log file) by default. This could quickly fill the file system. To avoid the problem, use the *-s* option on the *start_asiq* command to set it to a location other than “user”. For example:

```
start_asiq -s local2
```

For a complete list of *start_asiq* parameters, see the *Sybase IQ Reference Manual*.

6.8 Dbospace management and file placement

When you allocate file system files for dbspaces (System, IQ Main or IQ Temporary), do not place the files on a file system that is shared over a local area network. Violating this rule can lead to poor I/O performance and other problems, including overloading the local area network. On UNIX and Linux platforms, avoid Network File System (NFS) mounted file systems. On Windows, do not place dbospace files on network drives owned by another server.

Dbospace management should be performed by a single database administrator on a single connection to avoid conflicts.

6.9 Dbospace naming in RESTORE command [CR 561366]

If the dbospace name contains a file extension such as *.iq*, *.iqtmp*, or *.iqloc*, you must enclose it in double quotation marks when specifying the name in a RESTORE command RENAME clause.

For example:

```
RENAME local1 TO '/work/local1_res.iqloc.iqloc'  
DBSPACENAME "local1_res.iqloc"
```

or

```
RENAME "test_prod2.iq" TO '/test/test_prod7.iq'
```

6.10 Query restrictions

By default, Sybase IQ cursors are scrollable, meaning that Sybase IQ keeps all the query results in a buffer so that you can scroll backwards. If the query returns very large numbers (millions) of rows of output, you can improve performance by issuing the following command before running the query:

```
SET TEMPORARY OPTION Force_No_Scroll_Cursors = 'ON'
```

Some front-end applications make use of scrolling cursor operations, however, and require this option to be set "OFF".

If scrolling cursors are never used in your application, you should make this option a permanent public option. It will reduce memory usage and make a modest improvement in query performance.

6.11 User-defined functions and query results

In very few cases, differences in semantics between Adaptive Server Anywhere and Sybase IQ can produce different results if a query is issued in a user-defined function or functions in selects that have no FROM clause. For example, Sybase IQ treats the CHAR and VARCHAR data types as distinct and different, and Adaptive Server Anywhere treats CHAR data as if it were VARCHAR.

6.12 Interactive SQL restrictions

Sybase recommends that you use the Java edition of DBISQL, rather than DBISQLC or isql, unless instructed otherwise for specific tasks. Note that although dbisqlc is supported, dbisqlc does not contain all the features of the Java version dbisql.

- Do not use the Import option in the DBISQL Data menu (or Command > Options > Input Format in DBISQLC). This option is not supported for use with IQ databases. Use the LOAD TABLE statement or the INSERT statement to load data into IQ tables.
- If you run DBISQL (Interactive SQL Java) with the -q (quiet mode) option, and the data extraction commands (primarily setting the option TEMP_EXTRACT_NAME1 to an output file) are in a command file, you must first set and make permanent the DBISQL option "Show multiple result sets." If you do not set this option, the output file is not created.

To set the “Show multiple result sets” option, click Tools > Options in the DBISQL window, then check the box “Show multiple result sets” and click “Make permanent.”

- In DBISQL and DBISQLC, if you click the Help button or the Online Books button, the error “Cannot open the file” or “Cannot open help” is returned. Online help is not available for DBISQL and DBISQLC. [CR 370040]
- The query plan in the DBISQL Plan tab is an Adaptive Server Anywhere style query plan. For a Sybase IQ query plan, refer to the IQ *.iqmsg* file.
- If you set the terminal type on UNIX and Linux systems to “dumb” or “unknown” and then start DBISQLC, Sybase IQ returns an error. For example:

```
% setenv TERM dumb
% DBISQLC
Error at line 1
Unable to initialize screen routines
```

To avoid problems, run DBISQL (Interactive SQL Java) instead, or use an xterm window to run DBISQLC on UNIX and Linux systems. For example, you can start an xterm window with a scroll bar as follows:

```
% xterm -sb
```

6.13 SQL command line length is restricted

A current restriction limits the SQL command line length to the Catalog page size (not IQ page size) of the connected database. Because the default Catalog page size is 4KB, this is only a problem when the connection is to a database such as *utility_db*, which has a page size of 1024. This restriction may cause RESTORE commands that reference numerous dbspaces to fail.

To avoid the problem, make sure the length of SQL command lines is less than the Catalog page size.

Alternatively, start the engine with `-gp 32768` to increase Catalog page size.

7. Known problems

For a description of known issues in Sybase IQ 12.7, see the following sections. If there is a workaround for a problem, it is provided. Change request numbers appear in brackets, as in [CR 235449].

See also “Restrictions” on page 88 for more information.

“Documentation updates and clarifications” on page 100 contains details that were not documented in time for this release.

7.1 Data manipulation

This section reports known issues with data manipulation.

7.1.1 ROUND function truncates [CR 461126]

This issue is fixed in Sybase IQ 12.7 ESD #3.

The ROUND function truncates rather than rounds.

7.2 Sybase IQ operations

This section reports workarounds for issues with Sybase IQ operations.

7.2.1 PATINDEX returns 0 for zero-length search string [CR 475209]

The PATINDEX function should return 1 for a zero-length pattern. PATINDEX, however, returns 0.

7.2.2 Creating server entries with DSEDIT [CR 446092, CR 446419]

A problem occurs when creating a server entry using dsedit from the OCS 15 install via an xterm window. Clicking OK fails to dismiss the Server Add panel. If you click the X in the upper-right corner to close the xterm window, the system reports a Segmentation fault and core dumps.

If you try to add the IQ server while it is running, dsedit disallows saving the server entry because it detects that the port is in use. To create the entry, you must first shut down the server.

If the Sybase IQ server is running on the port you specify, dsedit displays a message indicating that port is in use. Note that this message may display behind the other dsedit windows. Click OK to dismiss the message and dsedit adds the server.

7.2.3 CIS returns “No Suitable Driver” error [CR 436095]

This issue is fixed in Sybase IQ 12.7 ESD #1.

When you use Component Integration Services (CIS) in certain geographic regions, connection attempts return the error No Suitable Driver. Java Development Kits (JDKs) used with Sybase IQ 12.7 support a specific set of time zone codes.

JDKs 1.3 and prior provide limited support for time zones and daylight savings time. If you encounter the No Suitable Driver error, use a TZ variable with GMT+nnn to avoid interactions between named time zones (for example, Asia/Seoul) and daylight savings time.

- For databases using default JDK 1.1.8:

Substitute GMT+9 for unsupported time zone KST:

```
setenv TZ GMT+9
```

- For databases using JDK 1.3:

Substitute GMT+9 for named time zone Asia/Seoul:

```
setenv TZ GMT+9
```

After you set the time zone environment variable to a supported setting and start the server, CIS works as expected. To ensure that the correct setting is always used, you can set the time zone in the start_asiq script.

See section “TZ environment variable” in Chapter 1, “File Locations and Installation Settings,” in the *Sybase IQ Reference Manual* for the supported time zone codes.

7.3 Sybase IQ queries

This section reports workarounds for query issues.

7.3.1 Query with derived table returns error [CR 427346]

This issue is fixed in Sybase IQ 12.7 ESD #2.

A query with a derived table that executed without error in Sybase IQ 12.5 returns an `item already exists` error in Sybase IQ 12.7 when run against a Sybase IQ server or with a SQL Anywhere server using a remote proxy table.

The error occurs in a query like the following where the derived table uses a column alias for a constant ('2' as `import_item`) and the outer query block has a column by the same name (`import_item`):

```
SELECT TB.dept_item, TB.import_item
INTO #t_Test_table1
FROM (
  SELECT '2' AS dept_item,
         '2' AS import_item
  FROM Test_Table A
) TB
GROUP by TB.dept_item, TB.import_item
```

The workaround is to specify column aliases, for example:

```
SELECT TB.dept_item AS dept_item, TB.import_item
AS import_item
```

7.4 Sybase IQ ETL

This section reports issues with Sybase IQ ETL Server and Sybase IQ ETL Development software. See “Documentation updates and clarifications” on page 100 for ETL documentation issues.

7.4.1 Required minimum screen resolution of 1024x768 pixels [CR 481841]

The Content Explorer dialog is too large to display when the screen resolution is set to 800x600 pixels. The workaround is to set the screen resolution to 1024x768 pixels.

7.4.2 IQ BulkLoad fails if Schema field is selected [CR 482941]

When you transfer data to IQ, the IQ BulkLoad component fails if you select the Schema field during the IQ BulkLoad configuration. The workaround is to avoid selecting the Schema field when configuring the IQ BulkLoad component.

7.4.3 File Content display in Text Data Provider [CR 456431]

When entering the configuration window of a text provider component the file content is read with the current (default is ISO-8859-1) encoding setting. It is not reread when a different encoding is selected. The preview section of the configuration window displays data correctly (according to the current settings) and the functionality of the component is not affected.

To read and display the file content with a specific encoding, do *one* of the following and then reopen the configuration window:

- Select the encoding in the configuration window and click Save.
- Select the encoding in the Property section.

7.4.4 Initialization file settings for Sybase IQ ETL Server on UNIX [CR 456080]

On UNIX installations of Sybase IQ ETL Server, initialization (*.ini*) file settings are not read by the GridNode application unless you have a *GridNode.ini* in place.

❖ Specifying settings for the GridNode application

- 1 Make a copy of the *Default.ini* file named *GridNode.ini* in the *etc* subdirectory of your installation path.
- 2 Specify general settings (to be shared by all applications) in *Default.ini*.
- 3 Specify exclusive settings for the GridNode application in *GridNode.ini*.
- 4 Disable (remove or comment) the keys for general settings in *GridNode.ini*.

7.4.5 Project execution might hang when database runs out of space [CR 455779]

When a database involved in a data transfer runs out of space, the execution of a project might hang. This happens because the database driver does not return with an error message but waits until space is available to continue. Ensure that sufficient space is available before running a project or job.

7.4.6 Special characters in Sybase IQ ETL installation paths [CR 454526]

Special characters in the installation path of Sybase IQ ETL Development or Server may lead to problems in accessing databases.

For example, connecting to an Oracle database client returns the following error when Sybase IQ ETL Server is installed in a path containing parentheses, such as *C:\Program Files (x86)\Sybase*:

```
ORA-12154: TNS:could not resolve the connect
identifier specified
```

Sybase recommends selecting installation paths that contain only alphanumeric characters. To avoid parentheses issues, remove “(x86),” which is added by the Windows 2003 EE 64-bit operating system, from the default target directory during installation.

7.4.7 Recommended settings for IQ database options [CR 447096, CR 447097]

The IQ database option `FORCE_NO_SCROLL_CURSORS` should be at the default setting (OFF) or a fetch error may occur while moving data via ETL.

To adjust database option settings, use the `SET OPTION` command in `dbisql` or, in Sybase Central, right-click the database and choose Set Options from the submenu.

8. Documentation for this version

All documentation for version 12.7 is included with your product shipment on CD:

- The Sybase IQ Getting Started CD contains release bulletins and installation guides in PDF format, and may also contain other documents or updated information not included on the SyBooks CD. It is included with your software. The Getting Started Sybase IQ 12.7 Extended Enterprise Edition CD also includes *Sybase IQ ETL User's Guide*, part number DC00608-01-0421-01.

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 Sybase IQ Getting Started CD also contains release bulletins and installation guides in an HTML format 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.

- The Sybase IQ 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 may 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.

Note The SyBooks browser software runs on Windows and Linux platforms. Users with non-Linux UNIX platforms must use Acrobat Reader to open PDF files on the SyBooks CD.

Table 11 lists documentation on the Sybase IQ 12.7 Sybooks CD.

Table 11: Sybase IQ Sybooks contents

Part Number	Book Title
DC00171-01-1270-01	<i>New Features in Sybase IQ 12.7</i>
DC38159-01-1270-01	<i>Introduction to Sybase IQ</i>
DC00170-01-1270-01	<i>Sybase IQ System Administration Guide</i>
DC00169-01-1270-01	<i>Sybase IQ Performance and Tuning Guide</i>
DC38151-01-1270-01	<i>Sybase IQ Reference Manual</i>
DC00462-01-1270-01	<i>Sybase IQ Error Messages Guide</i>
DC36584-01-1270-01	<i>Sybase IQ Troubleshooting and Recovery Guide</i>
DC00168-01-1270-01	<i>Sybase IQ Utility Guide</i>
DC00172-01-1270-01	<i>Large Objects Management in Sybase IQ</i>
DC00467-01-1270-01	<i>Encrypted Columns in Sybase IQ</i>

A print copy edition of the Sybase IQ documentation set, 48200-01-1270-01, can be ordered separately.

You may also need to refer to the documentation for Adaptive Server Anywhere:

- *Adaptive Server Anywhere Programming Guide* – Intended for application developers writing programs that directly access the ODBC, Embedded SQL™, or Open Client™ interfaces, this book describes how to develop applications for Adaptive Server Anywhere.

- *Adaptive Server Anywhere Database Administration Guide* – Intended for all users, this book covers material related to running, managing, and configuring databases and database servers.
- *Adaptive Server Anywhere SQL Reference Manual* – Intended for all users, this book provides a complete reference for the SQL language used by Adaptive Server Anywhere. It also describes the Adaptive Server Anywhere system tables and procedures.

You can also refer to the Adaptive Server Anywhere documentation in the SQL Anywhere Studio 9.0.1 collection on the Sybase Product Manuals Web site. To access this site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

9. Documentation updates and clarifications

This section contains information omitted from documentation and new information that needs emphasis.

See *New Features in Sybase IQ 12.7* for a description of changes to the Sybase IQ documentation set in version 12.7.

9.1 Customer issues

The *CustomerIssues.txt* file lists customer issues resolved in Sybase IQ 12.7. Issues resolved in a subsequent software update (ESD or EBF) are documented in the cover letter for that update. *CustomerIssues.txt* is installed in the following directories:

- `$SYBASE/ASIQ-12_7` on UNIX
- `%SYBASE%/ASIQ-12_7` on Windows

For more information, see “Sybase EBFs and software maintenance” on page 179.

9.2 Message and error logging

This section contains changes to error messages and log files.

9.2.1 System log file shows “ODBC seriously misbehaves” [CR 447093]

The message “ODBC seriously misbehaves” may appear in the system log file. This message is generated when the Sybase IQ ETL component DB Full Load specifies a large block size. Reducing the block size removes the printing of this message and could also increase performance.

9.2.2 Named pipes permissions [CR 447092]

When sending data to IQ via a named pipe, the error “Permission denied” or “File cannot be found” may result when executing a job. Make sure that both the Sybase IQ server and the Sybase IQ ETL Server have access to and read/write permissions on the named pipe.

9.2.3 ETL logs UTC (Coordinated Universal Time) [CR 446767]

The GridNode process logs its timestamp in the *system.log* in UTC (Coordinated Universal Time), not the host machine time.

9.2.4 New and changed messages for LOCK TABLE [CR 444921, 472791]

The following changes to Sybase IQ error messages support the 12.7 ESD #3 LOCK TABLE feature described in “LOCK TABLE support [CR 480880, CR 472791, CR 463196, 444921]” on page 30.

The LOCK TABLE command returns errors as follows.

- When the specified table does not exist, you see the message:

```
Table '%1' not found (SQLCODE -141, SQLSTATE 42W33)
```

‘%1’ is the name of the first user-specified table-name that does not exist.

- If you specify an invalid time interval on the WAIT option, you see the ASA error message:

```
Cannot convert %1 to a %2 (SQLCODE -157, SQL STATE 53018)
```

In the preceding messages, ‘%1’ is the invalid time-string and ‘%2’ is “timestamp.”

The LOCK TABLE for WRITE mode returns the following errors:

- When the table name appears more than once in the table-list, you see the message:

```
Item '%1' already exists (Constant
SQL_NAME_NOT_UNIQUE,
SQLCODE -110, SQLSTATE 52010)
```

'%1' is the name of the first user-specified table-name that is duplicate.

- When the table-list contains both IQ and SA tables, you see the ASA error message:

```
Table '%1' not found (SQLCODE 141, SQLSTATE 42W33)
```

'%1' is the name of the first SA table in the list.

- When lock(s) cannot be acquired during the specified time interval, you see the message:

```
Unable to acquire table locks in specified time
(LOCK_TIMEOUT, SQLCODE -1175L, SQLSTATE 42WC4)
```

- If the command is interrupted, you see the ASA error message:

```
Statement interrupted by user (SQLCODE -299,
SQLSTATE 57014)
```

- If any table is not an IQ base table, you see the following ASA error messages.

For temporary tables, join virtual tables, or Adaptive Server Anywhere tables, you see the ASA error:

```
Table '%1' not found (SQLCODE 141, SQLSTATE 42W33).
```

%1 is the name of the first non-IQ base table in the list.

- When attempting to lock a read-only table, you see this ASA error:

```
Permission denied: you do not have permission
to use the %1 (SQLCODE -121, SQLSTATE 42501)
```

In the preceding statement, %1 is the LOCK TABLE WRITE statement.

Examples of a locking a read-only table include, but are not limited to, locking a table owned by the multiplex write server on a query server or locking a table when the database is read-only.

- When the WITH HOLD option is specified, you see the error message:

```
Syntax error near '%1', %2 (SQLCODE -131, SQLSTATE
42W04)
```

- If the WAIT option is not specified, you see the message:

Syntax error near '%1', %2 (SQLCODE -131, SQLSTATE 42W04)

- If the user does not have at least one of the required permissions necessary to either insert, load, update, delete or truncate on all the table(s), you see the ASA error:

Permission denied: you do not have permission to use the %1 (SQLCODE -121, SQLSTATE 42501)

In the preceding statement, %1 is the LOCK TABLE WRITE statement.

The LOCK TABLE READ/EXCLUSIVE statement returns the following error.

- If the statement fails, even with WAIT option specified, you see the ASA error:

-210, User '%1' has the row in '%2' locked.

'%1' is the name of the user who owns the lock. '%2' is the user-specified table name.

DML update commands return the same error when a write-write conflict occurs.

9.2.5 New and changed messages for LOAD TABLE FORMAT BCP [CR 437332]

The following changes to Sybase IQ error messages support the 12.7 ESD #2 LOAD TABLE FORMAT BCP feature described in "LOAD TABLE supports BCP files as input [CR 437332]" on page 45.

- The message "Only ASCII and BINARY are supported LOAD formats." (SQLCODE -1000188L, SQLSTATE QDB88, Sybase error code 20848) is now "Only ASCII, BCP and BINARY are supported LOAD formats."
- The message "Delimiter '%2' must be 1 to 4 characters in length." (SQLCODE -1013054L, SQLSTATE QCA54, Sybase error code 20624) is now "Delimiter '%2' must be 1 to %3 characters in length.", where %3 is the maximum allowable length of the column delimiter or row terminator.
- A new error message has been added: "Invalid load specification for LOAD ... FORMAT BCP." (SQLCODE -1000296L, SQLSTATE QDC78, Sybase error code 21124).

9.3 Interactive SQL (DBISQL) online help

Online help is not available for the utilities Interactive SQL (DBISQLC) and Interactive SQL Java (DBISQL).

9.4 New Features in Sybase IQ 12.7

This section contains updates to *New Features in Sybase IQ 12.7*.

9.4.1 Perl interface (DBD::ASAny) not supported by Sybase IQ [CR 443439]

Under “Miscellaneous enhancements,” the paragraph stating “The new DBD::ASAny driver for the Perl DBI module allows you to access and modify Adaptive Server Anywhere databases from Perl scripts,” is incorrect.

Perl interface (DBD:ASAny) is not supported by Sybase IQ.

9.4.2 Encryption function name errors [CR 441789]

In Chapter 1, the “Sybase IQ Encrypted Column Option” section contains errors. The ENCRYPT and DECRYPT functions should read AES_ENCRYPT and AES_DECRYPT, respectively. These functions are listed correctly in *Encrypted Columns in Sybase IQ*.

9.4.3 Superscript display error [CR 437840]

In Chapter 1, “Limits changed,” number of rows per table appears as 248 in the PDF version, when it should appear as 2^{48} (2 to the 48th power). This display issue did not affect the HTML version on Infocenter or Chapter 2, “Physical Limitations” in the *Sybase IQ Reference Manual*.

9.4.4 Setting DBISQL options TEMPORARY [CR 435977]

Replace the section “DBISQL options can no longer be set TEMPORARY” in the section “Data load, update, and extraction changes” in Chapter 1, “New Features in Sybase IQ 12.7” with the following section:

DBISQL options can no longer be set TEMPORARY in BEGIN...END
Options that control how DBISQL interacts with the database can no longer be set TEMPORARY between the BEGIN and END keywords of a compound statement. For a list of affected options, see “DBISQL options” in Chapter 2, “Database Options,” in the *Sybase IQ Reference Manual*.

9.4.5 ASAODBC support on 64-bit platforms [CR 430313, CR 441021]

In Chapter 1, the “Startup and connection changes” section incorrectly states that ASAODBC is the only server class supported for 64-bit platforms. This is not true. ASAODBC is the only *ODBC* server class supported for 64-bit platforms. Other non-ODBC server classes are supported on 64-bit platforms. In addition to the error, this support was not a new feature for 12.7.

9.4.6 Errors in New Features in Sybase IQ 12.7

Certain sections in *New Features in Sybase IQ 12.7* are incorrect and should be disregarded.

In Chapter 1, “New Features in Sybase IQ 12.7,” disregard:

- Query timeout value added to plug-in preferences
- Validate Database wizard added
- Sun Solaris error initializing raw device
- SNMP Agent

In Chapter 2, “New Features in Sybase IQ 12.6 ESDs,” disregard:

- New start_asiq utility switches
- Domain names are case insensitive

9.5 User’s Guide for Sybase IQ ETL

This section contains updates to the *Sybase IQ ETL User’s Guide*.

9.5.1 Accessing the Sybase ETL Demo Repository

Sybase IQ ETL includes a demo repository with predefined projects and jobs designed to familiarize users with the product. The *Sybase IQ ETL User’s Guide* refers to that demo repository in Chapter 3, “Getting Started”, and Chapter 6, “Components”.

Starting Sybase IQ ETL Development displays the Repository window.

The default login data to access the demo repository is:

Connection: Repository
Client: transformer
Client user:TRANSFORMER
Password: transformer

These values are preset on initial login. On subsequent logins, you might need to select or enter this information.

❖ **Setting up a user account on the demo repository**

- 1 Enter a new Client user name and a password. (Note: do not change the Client.)
- 2 Activate the Register new user and the Show all objects option.
- 3 Click Logon.
- 4 Confirm your password and click OK.

Note If you enter a different Client or do not activate Show all objects you will not get access to the predefined transformation objects.

The default login information to the demo repository is automatically set up during Sybase IQ ETL Development installation. Use the following information to correct or restore the settings.

❖ **Configuring the ODBC User Data Source**

- 1 Configure the ODBC User Data Source:
 - Driver: Microsoft Access
 - Name: DEMO_Repository
 - Database: <datadir>\Demodata\IQETLDEMO_REP.MDB
- 2 Set up the repository connection in the Repository window:
 - a Select Repository from the Connection list and choose *one* of the following:
 - Edit
 - Add and type a name for the connection
 - b Select dbodbc from the Interface list.
 - c Select DEMO_Repository from the Host list.
 - d Click Save.

<datadir> is either the installation or the user data directory, see *Sybase IQ Installation and Configuration Guide* for details.

See also “Administering the repository” in Chapter 2 of the *Sybase IQ ETL User’s Guide* for details.

The projects in the demo repository require these additional ODBC User Data Sources:

Driver: Microsoft Access
Name: ETLDEMO_DWH; Database: DEMO_DWH.MDB
Name: ETLDEMO_GER; Database: DEMO_GER.MDB
Name: ETLDEMO_US; Database: DEMO_US.MDB

The database files are located in `<datadir>\Demodata\` where `<datadir>` is either the installation or the user data directory, see *Sybase IQ Installation and Configuration Guide* for details.

9.5.2 Sybase IQ ETL Demonstrations

Sybase IQ ETL provides flash demonstrations for nearly all components and some general functionality.

There are three ways to view the demos:

- Select Demonstration from a component popup menu in the Component Store.
- Choose either an Introduction or a Component Demonstration from the Help menu.
- View a demonstration directly from the training subdirectory of the installation path.

Note These demonstrations show general functionality of Sybase ETL products. Some of the features might not be available in your product edition.

9.5.3 Unicode support for Sybase IQ ETL [CR 447098]

The level of Unicode support for Sybase IQ ETL is:

- Unicode characters can be extracted, transformed, and loaded.
- Unicode characters in component properties are not supported. This includes:
 - File or directory names
 - Metadata like table or attribute names
 - Connection settings like database, schema, user, or password
- Unicode characters in ETL object names or descriptions are not supported.

9.6 System Administration Guide

This section contains updates to the *Sybase IQ System Administration Guide*.

9.6.1 Permissions to create external names

The first sentence in “Creating external logins” should change from “Only the login-name and the DBA account can add or modify an external login” to “Only the DBA account can add or modify an external login.”

9.6.2 Login auditing capability

The following information will be added to the *Sybase IQ System Administration Guide*.

Sybase IQ 12.7 provides the ability to log connection details for a given user in the transaction log file, then use your own scripts to extract information such as who logged in, how long they were logged in, and what operations were performed.

For details, see:

- “Auditing database activity” in Chapter 1, “Keeping Your Data Secure” in the *SQL Anywhere Studio Security Guide*. (Note that Sybase IQ does not support the -sf or -sk server option, and uses the sample database asiqdemo.db.)
- “The transaction log file” in Chapter 5, “Managing System Resources” in the *Sybase IQ Performance and Tuning Guide*.
- “AUDITING option [database]” in *Sybase IQ Reference Manual*.
- “Connection parameters” in Chapter 4, “Connection and Communication Parameters” in *Sybase IQ System Administration Guide*.
- “Log Translation utility (dbtran)” in Chapter 3, “Database Administration Utility” in the *Sybase IQ Utility Guide*.
- “sa_audit_string system procedure” in Chapter 8, “System Procedures” in *Sybase IQ Reference Manual*.

Auditing example

This example shows how the auditing feature records attempts at unauthorized information.

- 1 As database administrator, turn on the auditing option in the database.

You can do this from Sybase Central as follows:

- Connect using the Adaptive Server IQ Demo data source. This connects you as a DBA user.
- Select the asiqdemo database and from the File menu, choose Options.
- Select Auditing from the list of options, and type the value On in the Public Setting box. Click Set Permanent Now to set the option and Close to exit.

Alternatively, you can use Interactive SQL, as follows:

- Connect to the sample database from Interactive SQL as the DBA user.
- Execute the following statement:

```
SET OPTION PUBLIC.AUDITING = 'ON'
```

- 2 To enable auditing, you must also specify the type of events to audit, as follows. In Interactive SQL, enter:

```
sa_enable_auditing_type('Connect')
```

- 3 Add a user to the sample database, named AuditedUser, with password AuditedUser. You can do this from Sybase Central. Alternatively, you can use Interactive SQL and enter the following statement:

```
GRANT CONNECT TO AuditedUser  
IDENTIFIED BY 'AuditedUser'
```

- 4 Use Interactive SQL to connect to the sample database as AuditedUser and attempt to access confidential information in the employee table with the following query:

```
SELECT emp_lname, salary  
FROM employee
```

This returns the error: Permission denied: you do not have permission to select from "employee".

- 5 At a command prompt, execute the following command:

```
dbtran -g asiqdemo.log asiqdemo.sql
```

or, to run against the server:

```
dbtran -c "DSN=Adaptive Server IQ Demo"
-g asiqdemo.sql
```

Either command produces a file named *asiqdemo.sql*, which contains the transaction log information and a set of comments holding audit information. The lines that indicate the unauthorized AuditedUser attempt to access the Employees table are included in the file as follows:

```
---AUDIT-1010-0000936326 -- 2007/06/06
10:34:06.164 Checking Select permission on
DBA.employee - Failed

---AUDIT-1010-0000936365 -- 2007/06/06
10:34:06.164 Checking Select permission on
DBA.employee(salary) - Failed
```

- 6 Restore the sample database to its original state so that other examples you try in this documentation give the expected results.

Connect as a DBA user, and perform the following operations:

- Revoke Connect privileges from the user ID AuditedUser.
- Set the PUBLIC.AUDITING option to OFF.

9.6.3 sp_mpxcfg_<servername> no longer used [CR 500307]

Table 5-8 incorrectly lists the sp_mpxcfg_<servername> procedure as sp_iqmpxcfg_<servername>.

Multiplex no longer relies on SQL Remote, so this procedure will be removed from the documentation, along with the others that use the SQL Remote functionality, in the next GA release.

9.6.4 Default index for DATE data type [CR 499073]

When a table is created with the DATE data type, a two-byte FP index is created on the DATE field, which is independent of the settings in database option MINIMIZE_STORAGE.

If you want to create a three-byte FP or flat-style FP index on the DATE field, use the following IQ UNIQUE values when creating the table:

- For a three-byte FP — IQ UNIQUE should be between 65537 and 16777216.
- For flat-style FP — IQ UNIQUE should be higher than 16777216.

9.6.5 Explicit table lock acquisition [CR 444921]

As of 12.7 ESD #3, the LOCK TABLE statement supports acquiring WRITE locks on a set of tables instead of a single table. See “LOCK TABLE support [CR 480880, CR 472791, CR 463196, 444921]” on page 30 and “New and changed messages for LOCK TABLE [CR 444921, 472791]” on page 101.

The following information augments the section “How Locking Works,” in Chapter 10, “Transactions and Versioning.”

You can reserve WRITE locks on a set of tables within a new transaction using the LOCK TABLE statement. LOCK TABLE commits the current transaction and allows transactions to enqueue until the locks are available. For syntax, see “LOCK TABLE” in Chapter 6, “SQL Statements,” in the *Sybase IQ Reference Manual*.

For example, suppose that you are executing a critical transaction and you want to prevent update operation failure. Reserve WRITE locks on all the corresponding tables in advance, using commands like the following:

```
CREATE PROCEDURE test_lock()
BEGIN
  lbl:
  LOOP
    LOCK TABLE t1, t2, t3 IN WRITE MODE WAIT "30:00:00";
    IF SQLCODE <> 0 THEN
      PRINT 'Lock not yet acquired';
    ELSE
      LEAVE lbl;
    END IF;
  END LOOP lbl;
  UPDATE t1 set c1=5;
  INSERT INTO t2 values(10,10);
  TRUNCATE table t3;
  COMMIT;
END;
```

In the case of deadlocks, the last LOCK TABLE statement that became blocked is rolled back and an error returns to that transaction about the form of deadlock that occurred.

9.6.6 Files located relative to Catalog Store [CR 469000]

The information in “Example 2 — Moving the Catalog Store” in the “Moving database files” section of the “The RESTORE statement” section under “Restoring your databases” in Chapter 14, “Data Backup, Recovery and Archiving” is not correct regarding the directory in which a file is restored.

- In the second bullet, the sentences “Files originally created relative to the IQ_SYSTEM_MAIN dbspace, which holds the main IQ Store file, are restored relative to the main IQ Store file.” will be replaced with the following:

“Files originally created relative to the Catalog Store are restored relative to the Catalog Store.”

- The sentence “Relative pathnames in the RENAME clause work as they do when you create a database or dbspace: the main IQ Store dbspace, Temporary Store dbspaces, and Message Log are restored relative to the location of *db_file* (the Catalog Store); user-created IQ Store dbspaces are restored relative to the directory that holds the main IQ dbspace.” will be replaced with the following:

Relative pathnames in the RENAME clause work as they do when you create a database or dbspace: the main IQ Store dbspace, Temporary Store dbspaces, and Message Log are restored relative to the location of *db_file* (the Catalog Store); user-created IQ Store dbspaces are restored relative to the directory that holds the Catalog Store.”

9.6.7 Creating global or base tables affects other users [CR 466160]

The following warning replaces the warning currently in “Using table and column constraints” in Chapter 9, Ensuring Data Integrity,” in the *Sybase IQ System Administration Guide*.

Warning! Altering or creating global or base tables can interfere with other users of the database. For large tables, ALTER or CREATE TABLE can be a time-consuming operation. CREATE TABLE processing delays execution of other IQ processes until the statement completes. Although you can execute ALTER TABLE statements while other connections are active, you cannot execute them while any other connection uses the table to be altered. ALTER TABLE processing excludes other requests referencing the table being offered while the statement processes.

9.6.8 Configuring Open Client [CR 461831]

The following information was omitted from “Open Client applications and Sybase IQ” in Chapter 16, “Accessing Remote Data”.

When connecting to Sybase IQ using Open Client or when using the INSERT...LOCATION syntax, you can set various Open Client configuration parameters via an Open Client runtime configuration (.cfg) file. For example, you may change the maximum default number of connections, which is controlled by the value of the CS_MAX_CONNECT option.

The application name for INSERT...LOCATION is Sybase IQ. (The space between the words is required.) This application name is set at the Open Client connection level, not at the Open Client context level. For details about using an Open Client runtime configuration file and the options available, see the Open Client *Client-Library C Reference Manual*.

To have the .cfg take effect, stop and restart the Sybase IQ server. You may also specify certain configuration parameters in the INSERT...LOCATION command line. Parameters set in INSERT...LOCATION are superseded by parameters set in the configuration file.

9.6.9 Unique table names on a connection [CR 455684]

In Sybase IQ 12.7 ESD #3, an attempt to create a base table or a global temporary table will fail, if a local temporary table of the same name exists on that connection, as the new table cannot be uniquely identified by *owner.table*.

The information in the section “Unique table names on a connection [CR 455684]” on page 29 will be added in Chapter 5, “Working with Database Objects,” to the “Types of tables” section in “Creating tables” in the “Working with tables” section.

9.6.10 LOAD TABLE statement requires column specifications [CR 453567]

The simple LOAD TABLE example in Chapter 7, “Moving Data In and Out of Databases” is incorrect, because it contains no column specifications. The following example is correct:

```
LOAD TABLE department
( dept_id, dept_name, dept_head_id )
FROM '/d1/MILL1/dept.txt'
```

9.6.11 INSERT...LOCATION supports 20MB large object (LOB) data [CR 453439]

Sybase IQ 12.7 ESD #3 increases the limit on the length of large object data (LOB) that can be retrieved from a remote database using INSERT...LOCATION from 32767 bytes to 20MB (20971520 bytes). In Chapter 7, “Moving Data In and Out of Databases,” the last sentence in the first paragraph in the section “Loading ASE text and images” should be changed so that the paragraph says:

“Sybase IQ does not support the Adaptive Server Enterprise data type TEXT, but you can execute INSERT...LOCATION from an ASE database column of data type TEXT. Also note that INSERT...LOCATION does not support the use of variables in the SELECT statement. Text and image data inserted is silently right truncated at 20971520 bytes.”

9.6.12 Implicit conversion between BIT and CHAR/VARCHAR data types supported [CR 451862]

Sybase IQ 12.7 ESD #2 supports implicit conversion between BIT and CHAR and BIT and VARCHAR data types for arithmetic, comparison, INSERT, and UPDATE operations.

In Chapter 7, “Moving Data In and Out of Databases,” the tables in the section “Data conversions in IQ” should be updated as follows:

- In Table 7-6: “IQ conversions for comparison operations,” the conversion in the rows From: bit To: ch (CHARACTER data type) and From: bit To: vc (VARCHAR data type) should be “I” (implicit) instead of “E” (explicit).
- In Table 7-6: “IQ conversions for comparison operations,” the conversion in the rows From: char To: bt (BIT data type) and From: varchar To: bt should be “I” (implicit) instead of “E” (explicit).
- In Table 7-7: “IQ conversions for arithmetic operations,” the conversion in the rows From: bit To: ch (CHARACTER data type) and From: bit To: vc (VARCHAR data type) should be “I” (implicit) instead of “E” (explicit).
- In Table 7-7: “IQ conversions for arithmetic operations,” the conversion in the rows From: char To: bt (BIT data type) and From: varchar To: bt should be “I” (implicit) instead of “E” (explicit).
- In Table 7-8: “IQ conversions for INSERT and UPDATE,” the conversion in the rows From: bit To: ch (CHARACTER data type) and From: bit To: vc (VARCHAR data type) should be “I” (implicit) instead of “E” (explicit).
- In Table 7-8: “IQ conversions for INSERT and UPDATE,” the conversion in the rows From: char To: bt (BIT data type) and From: varchar To: bt should be “I” (implicit) instead of “E” (explicit).

9.6.13 Enabling ODBC trace output [CR 442772]

The following was omitted from the section “Using ODBC Data Sources on UNIX” in Chapter 3, “Sybase IQ Connections.”

To create an ODBC trace file, you must set the environment variables `ASA_ODBCTRACE_VERBOSE` and `ASA_ODBCTRACE_FILE`, for example:

```
$ setenv ASA_ODBCTRACE_VERBOSE ALL
$ setenv ASA_ODBCTRACE_FILE
/bluesun/fiona/odbctrace.out
$ dbisql -c "dsn=J123456" -nogui
```

On UNIX and Linux systems, use the `libdbodbc9.so` driver and leave it up to the driver to choose the multi-threaded or non-threaded driver. Tracing capability exists in the switch (`libdbodbc9.so`), not in the individual drivers (`libdbodbc9_n.so` or `libdbodbc9_r.so`). If you change the driver to point to the `_r` version, you are taking the switch out of the call sequence, preventing the tracing.

9.6.14 LOAD TABLE FORMAT syntax correction [CR 445737]

In the section “Bulk loading data using the LOAD TABLE statement” in Chapter 7, “Moving Data In and Out of Databases,” the syntax for the `FORMAT` option of the `LOAD TABLE` statement should not contain single quotes. The correct syntax is:

```
LOAD TABLE [ owner.]table-name
[ ( load-specification, ... ) ]
FROM 'filename-string', ...
[ FORMAT { ascii | binary } ]
...[ DELIMITED BY string ]
...[ STRIP { ON | OFF } ]
...[ QUOTES { ON | OFF } ]
...[ ESCAPES { ON | OFF } ]
...[ ESCAPE CHARACTER character ]
...[ WITH CHECKPOINT ON | OFF ]
...[ load-options ]
```

9.6.15 Group permissions cannot be revoked for a specific user [CR 444928]

When users are assigned membership in a group, they inherit all the permissions on tables, views, and procedures associated with that group. If you do not want a specific user to access a particular table, view, or procedure, then do not make that user a member of a group that has permissions on that object. You cannot revoke permissions for a specific user within a group.

9.6.16 Minimum IQ page size corrected [CR 443231]

The section “Creating databases” in Chapter 1, “Overview of Sybase IQ System Administration” contains an error.

For very large databases, you need an IQ page size of 128KB or larger.

9.6.17 INSERT...LOCATION supports additional ASE data types [CR 442348]

12.7 ESD #2 supports inserting data from an Adaptive Server Enterprise database column of data type UNSIGNED SMALLINT, NCHAR, NVARCHAR, DATE, or TIME using the INSERT...LOCATION syntax of the INSERT statement.

Note The Adaptive Server Enterprise data type UNSIGNED SMALLINT is not supported by Sybase IQ INSERT...LOCATION on Linux 32-bit systems, but is supported by INSERT...LOCATION on all other Sybase IQ supported platforms.

The following additions document the support of these ASE data types with INSERT...LOCATION.

In Chapter 7, “Moving Data In and Out of Databases,” the third bullet under “Note the following” in the section “Unsupported Adaptive Server Enterprise data types” should now say:

- Sybase IQ does not support the Adaptive Server Enterprise data types TEXT or unsigned smallint, but you can insert data from an ASE database column of data type TEXT or unsigned smallint using the LOCATION syntax of the INSERT statement.

Add the following row to Table 7-14: Integer data types in Chapter 7, “Moving Data In and Out of Databases.”

Adaptive Server Enterprise data type	Sybase IQ data type	Notes
unsigned smallint	Not supported	Sybase IQ does not support the Adaptive Server Enterprise data type unsigned smallint, but you can insert data from an ASE database column of data type unsigned smallint using INSERT...LOCATION.

Add the following row to Table 7-16: Character data types in Chapter 7, “Moving Data In and Out of Databases.”

Adaptive Server Enterprise data type	Sybase IQ data type	Notes
nchar	Not supported	Sybase IQ does not support the Adaptive Server Enterprise data type nchar, but you can insert data from an ASE database column of data type nchar using INSERT...LOCATION to character type columns.
nvarchar	Not supported	Sybase IQ does not support the Adaptive Server Enterprise data type nvarchar, but you can insert data from an ASE database column of data type nvarchar using INSERT...LOCATION to character type columns.

Add the following rows to Table 7-18: DATE/TIME data types in Chapter 7, “Moving Data In and Out of Databases.”

Adaptive Server Enterprise data type	Sybase IQ data type	Notes
date	date	You can insert data from an ASE database column of data type date using INSERT...LOCATION.
time	time	The Sybase IQ data type is the Time of day, containing hour, minute, second, and fraction of a second. The fraction is stored to 6 decimal places. A time value requires 8 bytes of storage. The Adaptive Server Enterprise data type time is between 00:00:00:000 and 23:59:59:999. You can use either military time or 12AM for noon and 12PM for midnight. A time value must contain either a colon or the AM or PM signifier. AM or PM may be in either uppercase or lowercase. A time value requires 4 bytes of storage. You can insert data from an ASE database column of data type time using INSERT...LOCATION.

9.6.18 ASE data type XML not supported [CR 442348]

In the section “Unsupported Adaptive Server Enterprise data types” in Chapter 7, “Moving Data In and Out of Databases,” the list of Adaptive Server Enterprise data types that are not supported by Sybase IQ in this version should include xml data type.

9.6.19 TIME to DATE conversion unsupported for INSERT and UPDATE [CR 442348]

In Table 7-8: “IQ conversions for INSERT and UPDATE” in Chapter 7, “Moving Data In and Out of Databases,” the conversion in the row From: time To: dt (DATE data type) should be changed to “U” (unsupported) from “E” (explicit).

9.6.20 DbSpace creation and raw devices [CR 441162]

The following clarification was omitted from “Adding dbspaces” in Chapter 1, “Overview of Sybase IQ System Administration.”

In the Create DBSpace wizard, on the Path screen, the raw device option is selected by default. Sybase recommends that you create main stores on raw devices.

When you specify a raw device for a new dbspace, Sybase IQ determines its size automatically and allocates the whole device for use as an IQ store. This may have unpredictable results if the device is actually a file device and Sybase does not recommend this practice.

If you indicate that the device is not raw, then the wizard enables a field where the size may be specified. The wizard also verifies that the given path exists before moving to the next page.

9.6.21 ESCAPE_CHARACTER option for system use only [CR 440545]

The ESCAPE_CHARACTER option is reserved for system use only. Do not change the setting of this option.

In Chapter 5, “Working with Database Options,” the paragraph following Table 5-2 should not refer to the ESCAPE_CHARACTER option. The paragraph should say:

“On Windows systems, when you specify device names that include a backslash, you must double the backslash to keep the system from mistaking a backslash/letter combination for an escape sequence such as tab or newline command.”

9.6.22 LOAD TABLE supports BCP files as input [CR 437332]

Sybase IQ 12.7 ESD #2 supports using a BCP character file as input to the LOAD TABLE command. For syntax and usage updates to the LOAD TABLE information in Chapter 7, “Moving Data In and Out of Databases” of the *Sybase IQ System Administration Guide*, see “LOAD TABLE supports BCP files as input [CR 437332]” on page 45.

The LOAD TABLE FORMAT BCP feature in Sybase IQ 12.7 ESD #2 eliminates the need for special handling of dates in Adaptive Server Enterprise BCP OUT files. The section “Ensuring accurate dates” in the section “Matching Adaptive Server Enterprise data types” in Chapter 7, “Moving Data In and Out of Databases” can be ignored and has been removed from the chapter for future releases.

9.6.23 Replacing a write server with intact files [CR 424433]

In the section “Replacing write servers,” in Chapter 14, “Data Backup, Recovery, and Archiving,” Step 2 in “Replacing a write server with intact files” should read as follows:

Use FTP or some other system facility to copy the write server’s SYSTEM dbspace files (*dbname.db*, typically, is the only one) to the new write server’s directory.

You must copy *dbname.db* to the query server directory that will be promoted. This cannot be a new directory and must be the query server directory. The file cannot be copied to any other directory. If the write server’s database or log file has a name that is different from the query server’s database or log file, then the name must be changed when it is copied over. For example:

```
cp /writesvr_dir/w1.db /querysvr_dir/q1.db
cp /writesvr_dir/w1.log /querysvr_dir/q1.log
```

9.6.24 Clarification on using ODBC data sources on UNIX [CR 422223]

In Chapter 3, “Sybase IQ Connections,” the following correction applies to “Using ODBC data sources on UNIX.”

After the statement “References to ODBC functions are resolved at run time,” the section should read as follows:

To connect with ODBC data sources, the location of your *.odbc.ini* file must be referenced by the \$ODBCINI or the \$ODBCHOME variable. Sybase IQ searches the directories specified by the variables below in the following order:

- 1 *\$ODBCINI* – must contain the exact full pathname of the *.odbc.ini* file.
- 2 *\$ODBCHOME* – must be set to the directory that contains the *.odbc.ini* file.

9.6.25 WORD index and CONTAINS predicate support for character large object (CLOB) columns [CR 415547]

Sybase IQ 12.7 ESD #3 provides support for WORD (WD) indexes on LONG VARCHAR character large object (CLOB) columns and support of the CONTAINS predicate on LONG VARCHAR (CLOB) columns with a WD index.

The following documentation updates will be made to Chapter 6, “Using Sybase IQ Indexes” for this feature.

- In the section “Overview of indexes,” the bullet regarding the WD index in the section “Sybase IQ index types” will be changed to say:
“WD Used to index keywords by treating the contents of a CHAR, VARCHAR, or LONG VARCHAR column as a delimited list.”
- In the section “Choosing an index type,” the following row will be added to Table 6-2 “Query type/index” in the section “Types of queries”:

Type of query usage	Recommended index type
In a CONTAINS predicate	WD

- In the section “Choosing an index type,” the following phrase will be added to the third bullet regarding other factors in the section “Types of queries”:
LONG VARCHAR data can only be used in the default and WD index types;
- In the section “Choosing an index type,” the following row will be added to Table 6-4 “Indexes supported for data types” in the section “Data types in the index”:

Data type	Supported indexes	Unsupported indexes
long varchar	WD	CMP, HG, HNG, LF, DATE, TIME, DTTM

- Several changes will be made to the section “The Containment (WD) index type” in the section “Sybase IQ index types.”
 - The description of the Containment (WD) index will include the LONG VARCHAR data type:

“The Containment (WD) index allows you to store words from a column string of CHAR, VARCHAR, and LONG VARCHAR data.”
 - Three of the bullets in the “Recommended use” section will be modified as follows:
 - The column-name must identify a CHAR, VARCHAR, or LONG VARCHAR column in a base table.
 - The minimum permitted column width is 3 bytes and the maximum permitted column width is the maximum width for a LOB column. (The maximum length is equal to 4GB multiplied by the database page size.)
 - The sp_iqcheckdb (DBCC consistency checker) allocation, check, verify, and repair modes support the WD index on CHAR, VARCHAR, and LONG VARCHAR columns.
 - The third item in the Disadvantages column of Table 6-9 “WD advantages/disadvantages” will be changed to say:

“Can only use this index if data in your columns is CHAR, VARCHAR, or LONG VARCHAR.”

9.6.26 LOAD TABLE supports loading default values [CR 396843]

Sybase IQ 12.7 ESD #4 supports loading a default value into a column as specified in the LOAD TABLE statement. You can now load a default value into a column, even if the column does not have a default value defined in the table schema. The LOAD TABLE syntax includes a new DEFAULT value specification in the column specification.

The following bullets will be added to the list of column default value requirements in the section “Working with column default values” in “Using column defaults” in Chapter 9, “Ensuring Data Integrity” of the *Sybase IQ System Administration Guide*.

- The LOAD TABLE DEFAULTS clause must be set to ON in order to use the default value specified in the LOAD TABLE statement DEFAULT value specification. If the DEFAULTS clause is OFF, the specified load default value is not used and a NULL value is inserted into the column instead.

- The LOAD TABLE DEFAULT value specification must contain at least one column that needs to be loaded from the file specified in the LOAD TABLE command.
- The LOAD TABLE DEFAULT *default-value* must be of the same character set as that of the database and must conform to the supported default values for columns and default value restrictions. The LOAD TABLE DEFAULT clause does not support AUTOINCREMENT, IDENTITY, or GLOBAL AUTOINCREMENT as a load default value.
- Encryption of the default value is not supported for the load default values specified in the LOAD TABLE DEFAULT value specification.

For syntax and usage updates to the section “Bulk loading data using the LOAD TABLE statement” in Chapter 7, “Moving Data In and Out of Databases” of the *Sybase IQ System Administration Guide*, see “LOAD TABLE supports loading default values [CR 396843]” on page 47.

9.6.27 Bulk loading using the LOAD TABLE statement

In Chapter 7, “Moving Data In and Out of Databases,” the last sentence in “Using command files to load data” contains a cross reference error. It should read as follows:

To create a command file, follow the instructions in Chapter 2, “Using Interactive SQL (dbisql),” in the *Sybase IQ Utility Guide*.

9.6.28 JDBC and debugger documentation

Appendix B, “Data Access Using JDBC,” and Appendix C, “Debugging Logic in the Database,” have not been updated for version 12.7.

9.7 Reference Manual

This section contains updates to the *Sybase IQ Reference Manual*.

9.7.1 SET OPTION statement for groups [CR612762]

In the Usage section for the SET OPTION statement, the third paragraph should read as follows:

“Specifying either a user ID or the PUBLIC user ID determines whether the option is set for an individual user, a user group represented by userid, or the PUBLIC user ID (the user group to which all users are a member). If the option applies to group user ID, option settings are not inherited by members of the group. If no user group is specified, the option change is applied to the currently logged-on user ID that issued the SET OPTION statement.”

9.7.2 Web service examples return errors [CR 610258]

The Sybase IQ documentation refers to web service examples in the section “Creating web service client functions and procedures” in *Adaptive Server Anywhere Database Administration Guide*. When run with Sybase IQ, some of the code in these examples returns errors.

On page 255, the sample code to run with Sybase IQ should be:

```
CREATE FUNCTION get_webservices ()
RETURNS LONG VARCHAR
URL 'HTTP://localhost:81/get_webservices'
TYPE 'SOAP'
```

On page 256, the sample code to run with Sybase IQ should be:

```
SELECT *
FROM OPENXML( get_webservice(), '//row' )
WITH (Name char(128) 'service_name',
Secure char(1) 'secure_required' )
```

9.7.3 Setting LOGIN PROCEDURE database option [CR 581936]

Information was omitted from the description for the LOGIN_PROCEDURE option in “Alphabetical list of options,” in Chapter 2, “Database Options,” in the *Sybase IQ Reference Manual*.

After the sentence, “The customized login procedure must be created in every database you might use,” add the following:

When you set LOGIN_PROCEDURE to call a new procedure, the custom procedure needs to either call sp_login_environment or detect when a TDS connection occurs (see the default sp_login_environment code) and call sp_tsql_environment directly.

9.7.4 Database name size limit [CR 529339]

The database name size limit of 250 bytes was missing in the documentation. It has been added to Table 6-1 in the “Physical Limitations” chapter.

This revision will be available in the next release.

9.7.5 sp_iqstatus property name changes [CR 528489]

The sp_iqstatus stored procedure’s output has changed. “Number of Main DB Spaces” now reads “Number of Main DB Files” and “Number of Temp DB Spaces” now reads “Number of Temp DB Files.”

This correction will appear in the “System Procedures” chapter in the next release of the Sybase IQ documentation.

9.7.6 QUERY_TEMP_SPACE_LIMIT default change [CR 496907]

The default for the QUERY_TEMP_SPACE_LIMIT option changed from 2000 to 0.

This correction will appear in the “Database Options” chapter in the next release of the Sybase IQ documentation.

9.7.7 LIKE conditions clarification [CR 488337]

A section describing like conditions was edited for clarity. Revisions will be available in the next GA release. The corrected paragraph reads:

LIKE predicates, where the pattern starts with non-wildcard characters, may execute faster if an HG or LF index is available.

9.7.8 Maximum number of UNION branches per query updated [CR 487488]

The number of UNION branches per query in Table 8-1, Chapter 8, “Physical Limitations” should be updated to 512:

“512. If each branch has multiple tables in the FROM clause, the limit on tables per query reduces the number of UNION branches allowed.”

9.7.9 Correct physical limit for SQL statements [CR 486412]

Table 8-1 in Chapter 8, “Physical Limitations” contained an incorrect limit description for the maximum length of SQL statement.

The corrected content reads:

The maximum length of SQL statements is limited to the amount of memory available for the IQ catalog and to the size of the catalog stack.

If your SQL statements are long, increase the catalog stack size using -gss and increase catalog memory cache size using -c or the combination of -ch plus -cl.

When printing the SQL statement in error messages, the text is limited to the IQ catalog page size. To print long commands, you can start the server with an increased -gp setting, although the default of -gp 4096 should generally be used.

This correction will appear in the next release of the Sybase IQ documentation.

9.7.10 BINARY data types and trailing zeros [CR 484687]

The compatibility information regarding the treatment of trailing zeros with BINARY data types in Chapter 4, “SQL Data Types” is not correct. In the section “Compatibility,” the first sentence and Table 4-4 should be replaced with the following:

“The treatment of trailing zeros in binary data types is different in Sybase IQ, Adaptive Server Anywhere, and Adaptive Server Enterprise. Table 4-4 shows the differences.”

Table 4-4: Treatment of trailing zeros

Data type	Sybase IQ	SQL Anywhere	ASE
BINARY NOT NULL	padded	not padded	padded
BINARY NULL	padded	not padded	not padded
VARBINARY NOT NULL	truncated, not padded	truncated, not padded	truncated, not padded
VARBINARY NULL	truncated, not padded	truncated, not padded	truncated, not padded

The following clarification should be added to the section “Treatment of trailing zeros” in the “Notes” section:

When you select a BINARY value, you must specify the value with the padded zeros or use the CAST function. For example,

```
SELECT * FROM zeros WHERE bnot = 0x0123000000;
```

or

```
SELECT * FROM zeros WHERE bnot = CAST(0x0123 as
binary(5));
```

9.7.11 LOAD TABLE IGNORE CONSTRAINT CHECK correction [CR 484525]

The description of the LOAD TABLE IGNORE CONSTRAINT CHECK *limit* clause in the section “LOAD TABLE statement” in Chapter 6, “SQL Statements” should say:

“If *limit* specifies zero, the number of CHECK constraint violations to ignore is infinite.”

9.7.12 TEMP_EXTRACT_ESCAPE_QUOTES option [CR 483767]

Function	Specifies whether all quotes in fields containing quotes are escaped in the output of the data extraction facility for an ASCII extraction.
Allowed values	ON, OFF
Default	OFF
Scope	Can be set for an individual connection. Takes effect immediately.
Description	This option is ignored unless TEMP_EXTRACT_QUOTE is the default or set to the value of "" (double quotes), and TEMP_EXTRACT_BINARY is OFF, and either TEMP_EXTRACT_QUOTES or TEMP_EXTRACT_QUOTES_ALL is ON.
See also	“TEMP_EXTRACT_BINARY option,” “TEMP_EXTRACT_QUOTES option,” and “TEMP_EXTRACT_QUOTES_ALL option” in the <i>Sybase IQ Reference Manual</i> .

9.7.13 sa_eng_properties new results set table [CR 481961]

The next GA release of the *Sybase IQ Reference Manual* includes the following table, which makes sa_eng_properties syntax more useful.

Table 12: sa_eng_properties result sets

Column name	Data type	Description
PropNum	integer	The database server property number
PropName	varchar(255)	The database server property name
PropDescription	varchar(255)	The database server property description
Value	long varchar	The database server property value

9.7.14 Maximum number of users correction [CR 480409]

Corrected maximum number of users (connected and concurrent) in Table 8-1, chapter 8, “Physical Limitations” to read:

- 1000 on 64-bit platforms: AIX, HP, Linux, and Sun Solaris
- 200 on 32- and 64-bit platforms on Windows

9.7.15 Extended GROUP BY semantics not supported by IQ [CR 479578]

Sybase IQ does not support extended group by semantics. All references in the documentation have been removed.

9.7.16 NOEXEC option clarification [CR 477741]

Operations such as INSERT...VALUES, LOAD, and SYNCHRONIZE are not affected by the NOEXEC option because they do not include a query.

9.7.17 SELECT statements can operate on stored procedure result sets [CR 475130]

The following was omitted from the SELECT statement in Chapter “SQL Statements”:

In SELECT statements, a stored procedure call can appear anywhere a base table or view is allowed. Note that CIS functional compensation performance considerations apply. For example, a SELECT statement can also return a result set from a procedure. For syntax and an example, see the FROM clause documentation in the *SQL Anywhere Server SQL Reference*.

9.7.18 sp_iqsysmon syntax correction [CR 474858]

In Chapter 10, “System Procedures,” the sp_iqsysmon buffer manager report section abbreviation in the “Batch mode usage” section should be “bufman,” not “bufnam.” The correct sentence is:

“If you specify “mbufman tbufman” or “bufman”, both the Main and Temporary Store buffer managers are monitored.”

In Table 10-42 the abbreviation for Buffer manager should be “(m/t)bufman.”

9.7.19 Understanding `sp_iqindexfragmentation` results [CR 474601]

A request was made to clarify how users understand the results of the `sp_iqindexfragmentation` procedure when using the `GARRAY_FILL_FACTOR_PERCENT` and `GARRAY_PAGE_SPLIT_PAD_PERCENT` options. This section contains the edits, which will be available in the next GA release of the Sybase IQ documentation.

`sp_iqindexfragmentation` procedure

Function	Reports information about the percentage of page space taken up within the btrees, garrays, and bitmap structures in Sybase IQ indexes. For garrays, the fill percentage calculation does not take into account the reserved space within the garray groups, which is controlled by the <code>GARRAY_FILL_FACTOR_PERCENT</code> option.
Syntax	<pre>dbo.sp_iqindexfragmentation ('<i>target</i>') <i>target</i>: table <i>table-name</i> (index <i>index-name</i> (...))</pre>
Permissions	DBA authority required. Users without DBA authority must be granted EXECUTE permission in order to run the stored procedure.
Usage	table-name Target table <i>table-name</i> reports on all nondefault indexes in the named table. index-name Target index <i>index-name</i> reports on the named index within the specified table. You can specify multiple indexes within the table, but you must repeat the index keyword with each index specified.

Example

The following procedure reports the internal index fragmentation for nonunique HG index cidhg in table customers:

```
dbo.sp_iqindexfragmentation ('index customers.cidhg ')
```

Index	Index type	btree node pages	Fill factor percent
dba.customers.cidhg	HG	3	75
SQLCODE	0		
Fill Percent	btree pages	garray pages	bitmap pages
0 - 10%	0	0	0
11 - 20%	0	0	0
21 - 30%	0	0	0
31-40%	0	0	22
41 - 50%	0	0	0
51 - 60%	0	0	10
61 - 70%	2	0	120
71 - 80%	138	3	64
81 - 90%	24	122	14
91 - 100%	18	1	0

According to this output, of the 182 btree pages in nonunique HG index cidhg, 2 are between 61% and 70% full, 138 are 71% to 80% full, 24 are 81% - 90% full, and 18 are 91% - 100% full. Usage for garray and bitmap pages is reported in the same manner. All percentages are truncated to the nearest percentage point. HG indexes also display the value of option GARRAY_FILL_FACTOR_PERCENT. Those index types that use a btree also display the number of node (nonleaf) pages. These are HG, LF, WD, DATE, and DTTM.

If an error occurred during execution of the stored procedure for this index, the SQLCODE would be nonzero.

See also

GARRAY_FILL_FACTOR_PERCENT option and
GARRAY_PAGE_SPLIT_PAD_PERCENT option in the following sections.

GARRAY_FILL_FACTOR_PERCENT option**Function**

Specifies the percent of space on each HG garray pages to reserve for future incremental inserts into existing groups. The garray tries to pad out each group to include a pad of empty space set by the value. This space is used for rows added to existing index groups.

Allowed values

0 – 1000

Default	25
Scope	DBA permissions are not required to set this option. Can be set temporary for an individual connection or for the PUBLIC group. Takes effect immediately.
Description	An HG index can reserve some storage on a per-group basis (where group is defined as a group of rows with equivalent values). Reserving space consumes additional disk space but can help the performance of incremental inserts into the HG index. If you plan to do future incremental inserts into an HG index, and those new rows have values that are already present in the index, a nonzero value for this option might improve incremental insert performance. If you do not plan to incrementally update the index, you can reduce the values of this option to save disk space.

GARRAY_PAGE_SPLIT_PAD_PERCENT option

Function	Determines per-page fill factor during page splits on the garray and specifies the percent of space on each HG garray page to reserve for future incremental inserts. Splits of a garray page try to leave that percentage empty. This space is used for rows added to new index groups.
Allowed values	0 – 1000
Default	25
Scope	DBA permissions are not required to set this option. Can be set temporary for an individual connection or for the PUBLIC group. Takes effect immediately.
Description	An HG index can reserve storage at the page level that can be allocated to new groups when additional rows are inserted. Reserving space consumes additional disk space but can help the performance of incremental inserts into the HG index. If future plans include incremental inserts into an HG index, and the new rows do not have values that are already present in the index, a nonzero value for the GARRAY_PAGE_SPLIT_PAD_PERCENT option could improve incremental insert performance. If you do not plan to incrementally update the index, you can reduce the values of this option to save disk space.

9.7.20 Change to sp_iqlocks output [CR 472791]

For Sybase IQ ESD #3, make the following changes to the Description column of table 10-32, “sp_iqlocks columns:”

- Add to the list of character strings: EW – exclusive and write, SW – share and write.
- Change to the list of character strings: S – share and exclusive.
- Change the first sentence of the second paragraph to “All locks listed have either S, E, EW, or SW. They may also have P, A, or both.”
- Add the line “Sybase IQ obtains a share lock before a write lock. If a connection has exclusive lock, share lock is not displayed. For write locks, if a connection has all-exclusive, share, and write locks, EW is displayed, else SW.”

9.7.21 Results returned in a case insensitive database [CR 472049]

In some cases, an optimization in Sybase IQ 12.7 causes the results of a comparison returned in a case insensitive database to be different from the results returned in previous releases of IQ. You cannot expect case sensitive results in a database that is case insensitive (CASE IGNORE). This has always been true for Sybase IQ.

To clarify the results expected from a case insensitive database, the information in the section “Results returned in a case insensitive database [CR 472049]” on page 58 will be added to the description of the CASE clause in the “Usage” section of the “CREATE DATABASE statement” section in Chapter 6, “SQL Statements.”

9.7.22 Clarify NULL in STDDEV and VARIANCE functions [CR 470567]

In the “SQL Functions” chapter, the STDDEV and VARIANCE functions contain an incorrect parameter, NULL. The information on NULL has been moved to the Usage section and is reworded for clarity:

NULL Returns NULL for a one-element input set in Sybase IQ 12.7. In prior versions, NULL returned zero.

This correction was also made to STDDEV_SAMP() and VAR_SAMP() functions and will appear in the next release of the Sybase IQ documentation.

9.7.23 INSERT...LOCATION PACKETSIZE limitation for TDS [CR 470047]

When INSERT...LOCATION is transferring data between a Sybase IQ server and a remote Sybase IQ or Adaptive Server Anywhere server, the value of the INSERT...LOCATION TDS PACKETSIZE parameter is always equal to 512 bytes, even if you specify a different value for PACKETSIZE.

The following sentence will be added to the PACKETSIZE description in the “Usage” section of the “INSERT statement” section in Chapter 6, “SQL Statements”:

“When INSERT...LOCATION is transferring data between a Sybase IQ server and a remote Sybase IQ or Adaptive Server Anywhere server, the value of the INSERT...LOCATION TDS PACKETSIZE parameter is always equal to 512 bytes, even if you specify a different value for PACKETSIZE.”

9.7.24 SELECT ... INTO returns errors if global variable is used [CR 469969]

The ROW_COUNT option could produce inconsistent results when used in a query involving global variables, system functions or proxy tables. Such queries are partly executed using CIS (Component Integrated Services). In such cases, use SELECT TOP *n* instead of setting ROW_COUNT, or set the global variable to a local one and use that local variable in the query.

9.7.25 Page size clarification [CR 468276]

Added the following statement to the end of the PAGE SIZE section under the CREATE DATABASE statement: “A larger page size might also be needed to execute queries involving large numbers of tables or views.”

This clarification will appear in the next release of the Sybase IQ documentation set.

9.7.26 CASE statement syntax correction [CR 467835]

The CASE statement syntax should end with END, not END CASE, as documented.

CASE value-expression

...WHEN [constant | NULL] THEN statement-list ...

... [WHEN [constant | NULL] THEN statement-list] ...

...ELSE statement-list

... END

This correction will appear in the next release of the Sybase IQ document set.

9.7.27 DISK_STRIPING_PACKED default value error [CR 467624]

The description of DISK_STRIPING_PACKED states incorrectly that the default value is OFF, when the default is ON.

The correction will appear in the next release of the Sybase IQ document set.

9.7.28 Added missing ALL syntax [CR 466582]

Syntax for the ALL condition was missing in the section, “Search conditions,” topic “ALL or ANY conditions.” The following syntax will appear in the next release of the Sybase IQ document set:

The syntax for ALL conditions is:

expression compare **ALL** (*subquery*)

where *compare* is a comparison operator.

9.7.29 READ ONLY cursors now see data at time of OPEN [CR 466135]

As of Sybase IQ 12.7 ESD #4, applications that use cursors declared FOR READ ONLY may notice a change in which version of the data is visible to a cursor.

The following statement will be added to the sections “DECLARE CURSOR statement [ESQL][SP],” “FETCH statement [ESQL][SP],” and “OPEN statement [ESQL][SP]” in Chapter 6, “SQL Statements” for this behavior change:

“A cursor declared FOR READ ONLY sees the version of table(s) on which the cursor is declared when the cursor is opened, not the version of table(s) at the time of the first FETCH.”

9.7.30 LOCK TABLE access clarification [CR 463204]

The following information will be added to the “Usage” section of the “LOCK TABLE statement” section in Chapter 6, “SQL Statements”:

LOCK TABLE statements run on tables in the IQ main store on the write server do not affect access to those tables from connections on query servers. For example:

On a write server connection, issue the command:

```
LOCK TABLE wrt1 WITH HOLD IN EXCLUSIVE MODE
```

sp_iqlocks on the write server confirms that the table wrt1 has an exclusive (E) lock.

The result of sp_iqlocks run on a connection on a query server does not show the exclusive lock on table wrt1. The user on this connection can see updates to table wrt1 on the write server.

Other connections on the write server can see the exclusive lock on wrt1 and attempting to select from table wrt1 from another connection on the write server returns the error "User DBA has the row in wrt1 locked."

9.7.31 SELECT statements without a FROM clause [CR 462007]

When constant values are in SELECT statements that include UNION ALL views but omit a FROM clause, use iq_dummy to avoid errors.

If you omit the FROM clause, or if all tables in the query are in the SYSTEM dbspace, Adaptive Server Anywhere processes the query, instead of Sybase IQ, and might behave differently, especially with regard to syntactic and semantic restrictions and the effects of option settings.

If you have a query that does not require a FROM clause, you can force Sybase IQ to process the query by adding the clause "FROM iq_dummy," where iq_dummy is a one-row, one-column table that you create in your database.

9.7.32 LOCK TABLE support [CR 444921]

As of 12.7 ESD #3, the LOCK TABLE statement supports acquiring WRITE locks on a set of tables. The LOCK TABLE statement also allows enqueueing for READ, WRITE, and EXCLUSIVE locks for a specified time period. For updates to LOCK TABLE syntax and usage in Chapter 6, "SQL Statements," see "LOCK TABLE support [CR 480880, CR 472791, CR 463196, 444921]" on page 30.

9.7.33 Multiple dbspaces and sp_iqindexinfo [CR 471137]

The following sentence was omitted from the Function section for sp_iqindexinfo in Chapter 10, "System Procedures:"

If the object resides on several dbspaces, sp_iqindexinfo returns the space used in all dbspaces, as shown in the example.

9.7.34 Index limit correction [CR 471124]

Table 8-1 in Chapter 8, “Physical Limitations,” shows a limit of 32,767 indexes per table. The correct limit is 2^{32} (approximately 4,000,000) indexes per table.

Corrections to the *Sybase IQ Reference Manual* will be available in the next GA release.

9.7.35 Evaluation order of conditions with logical operator OR [CR 469968]

In the section “Conditions with logical operators” in Chapter 3, “SQL Language Elements,” the description of combining conditions using OR will be changed as follows:

Conditions are combined using OR as follows:

condition1 **OR** *condition2*

The combined condition is TRUE if either condition is TRUE, FALSE if both conditions are FALSE, and UNKNOWN otherwise. There is no guaranteed order as to which condition, *condition1* or *condition2*, will be evaluated first.

9.7.36 GARRAY_FILL_FACTOR_PERCENT clarifications [CR 469543]

The function and description of the GARRAY_FILL_FACTOR_PERCENT option have been clarified in Sybase IQ 12.7 ESD #3. Corrections to the *Sybase IQ Reference Manual* will be available in the next GA release.

Function	Specifies the percent of space on each HG garray pages to reserve for future incremental inserts into existing groups.
Allowed values	0 – 1000
Default	25
Scope	DBA permissions are not required to set this option. Can be set temporary for an individual connection or for the PUBLIC group. Takes effect immediately.
Description	An HG index can reserve storage on a per-group basis (where group is defined as a group of rows with identical values). Reserving space consumes additional disk space but can help the performance of incremental inserts into the HG index. If future plans include incremental inserts into an HG index, and the new rows have values that are already present in the index, a nonzero value for the GARRAY_FILL_FACTOR_PERCENT option may improve incremental insert performance.

If you do not plan to incrementally update the index, you can reduce this option to save disk space.

9.7.37 GARRAY_PAGE_SPLIT_PAD_PERCENT option [CR 469513]

This option was implemented in Sybase IQ 12.7 but not documented. Corrections to the *Sybase IQ Reference Manual* will be available in the next GA release.

Function	Specifies the percent of space on each HG garray pages to reserve for future incremental inserts. This space can be used for rows added to both new and existing index groups.
Allowed values	0 – 1000
Default	25
Scope	DBA permissions are not required to set this option. Can be set temporary for an individual connection or for the PUBLIC group. Takes effect immediately.
Description	An HG index can reserve storage at the page level that can be allocated to new or existing groups when additional rows are inserted. Reserving space consumes additional disk space but can help the performance of incremental inserts into the HG index. If future plans include incremental inserts into an HG index, and the new rows have values that are already present in the index, a nonzero value for the GARRAY_PAGE_SPLIT_PAD_PERCENT option may improve incremental insert performance. If you do not plan to incrementally update the index, you can reduce this option to save disk space.

9.7.38 LOAD TABLE syntax missing ON PARTIAL INPUT ROW [CR 469027]

In Chapter 6, “SQL Statements,” the “Syntax” section of the “LOAD TABLE statement” section is missing the ON PARTIAL INPUT ROW parameter. The following line will be added to the LOAD TABLE syntax:

```
[ ON PARTIAL INPUT ROW { ROLLBACK | CONTINUE } ]
```

9.7.39 Files located relative to Catalog Store [CR 469000]

The information in the “Usage” section of the “CREATE DBSPACE statement” section in Chapter 6, “SQL Statements” contains incorrect information regarding the directory in which a dbspace is created, if an explicit directory is not specified.

The sentences “A filename without an explicit directory is created in the same directory as the initial dbspace of that store. Any relative directory is relative to that initial dbspace.” will be replaced with the following:

“A filename without an explicit directory is created in the same directory as the Catalog Store of the database. Any relative directory is relative to the Catalog Store.”

9.7.40 TOP_NSORT_CUTOFF_PAGES option [468580]

This option, added in 12.6 ESD #8, was not documented.

Function	Sets the result size threshold for TOP N algorithm selection.
Allowed values	1 – 1000
Default	1
Description	The TOP_NSORT_CUTOFF_PAGES option sets the threshold, measured in pages, where evaluation of a query that contains both a TOP clause and ORDER BY clause switches algorithms from ordered list-based processing to sort-based processing. Ordered list processing performs better in cases where the TOP N value is smaller than the number of result rows. Sort-based processing performs better for large TOP N values. In some cases, increasing TOP_NSORT_CUTOFF_PAGES can improve performance by avoiding sort-based processing.
See also	“SELECT Statement” in <i>Sybase IQ Reference Manual</i> .

9.7.41 Clarify TEMP_EXTRACT_DIRECTORY option [CR 468430]

The TEMP_EXTRACT_DIRECTORY option controls not only whether a user is allowed to use the data extraction facility but the directory into which temp extract files are placed. It can also override a directory path specified in the TEMP_EXTRACT_NAME_n options.

If the TEMP_EXTRACT_DIRECTORY option is set to the string FORBIDDEN (case insensitive) for a user, then that user is not allowed to perform data extracts. An attempt by this user to use the data extraction facility results in an error: “You do not have permission to perform Extracts”.

If TEMP_EXTRACT_DIRECTORY is set to FORBIDDEN for the PUBLIC group, then no one can run data extraction.

If TEMP_EXTRACT_DIRECTORY is set to a valid directory path, temp extract files are placed in that directory, overriding a path specified in the TEMP_EXTRACT_NAME options.

If TEMP_EXTRACT_DIRECTORY is set to an invalid directory path, an error occurs: “Files does not exist File: <invalid path>”

If TEMP_EXTRACT_DIRECTORY is blank, then temp extract files are placed in directories according to their specification in TEMP_EXTRACT_NAME. If no path is specified as part of TEMP_EXTRACT_NAME, the extract files are by default placed in the server startup directory.

The Sybase IQ manuals will be updated to reflect these changes in the next GA release.

9.7.42 Incorrect value listed for ON_TSQL_ERROR [CR 467582]

An incorrect value was listed for the ON_TSQL_ERROR option in Chapter 2, “Database Options.” It should be CONTINUE, not CONDITION. This correction will appear in the next version of the documentation.

9.7.43 Determining version using sp_iqstatus output [CR 463635]

To display space that can be reclaimed by dropping connections, use sp_iqstatus and add the results from the two returned rows:

```
(DBA)> select * from sp_iqstatus() where name like
'%Versions:%'
Execution time: 6.25 seconds
Name                               Value
-----
Other Versions: 2 = 1968Mb
Active Txn Versions: 1 = C:2175Mb/D:2850Mb

(First 2 rows)
```

The above example output shows that 1 active write transaction created 2175MB and destroyed 2850 MB of data. The total data consumed in transactions and not yet released is 4818MB, or 1968MB + 2850MB = 4818MB.

If `sp_iqstatus` shows a high percentage of main blocks in use on a multiplex server, run `sp_iqversionuse` to find out which versions are being used and the amount of space that can be recovered by releasing versions. See `sp_iqversionuse` in the *Sybase IQ Reference Manual*.

You can also display version information by using `sp_iqtransaction` to see total amount of MainTableKBCreated and MainTableKBDropped over transactions. See `sp_iqtransaction` in the *Sybase IQ Reference Manual* for details.

9.7.44 Correction to ASA NUMERIC data type default precision [CR 461147]

The default precision of Adaptive Server Anywhere NUMERIC and DECIMAL data types is incorrect in the “Compatibility” section of the “Numeric data types” section in Chapter 4, “SQL Data Types”. The default precision of the Adaptive Server Anywhere NUMERIC and DECIMAL data types is 30, not 3.

9.7.45 IDENTITY_INSERT option errors [CR 456339, CR 440042, CR 439496, CR 418760]

In Table 2-1, the VALUES column for the option IDENTITY_INSERT contains an error. The value *string* is followed by a single quotation mark where none should appear.

The following changes apply to Chapter 2, “Database Options.”

In the section “IDENTITY_INSERT option,” the final example should read:

To turn the option off, specify the equals sign and an empty string:

```
SET TEMPORARY OPTION IDENTITY_INSERT = "
```

The following note was omitted from the Scope section.

Note If you set a user level option for the current option, the corresponding temporary option is also set. For details, see “Scope and duration of database options” in Chapter 2, “Database Options.”

The following note was omitted from the Example section.

To illustrate the effect of user level options on temporary options (see note above), if you are connected to the database as DBA, and issue:

```
SET OPTION IDENTITY_INSERT = 'customer'
```

the value for the option is set to customer for the user DBA and temporary for the current connection. Other users who subsequently connect to the database as DBA find their option value for IDENTITY_INSERT is customer also.

9.7.46 Using correct syntax when creating stored procedures [CR 456307]

There are two ways to create stored procedures: T-SQL and SQL92. BEGIN TRANSACTION, for example, is T-SQL specific when using CREATE PROCEDURE syntax. Do not mix syntax when creating stored procedures.

9.7.47 Unique table names on a connection [CR 455684]

In Sybase IQ 12.7 ESD #3, an attempt to create a base table or a global temporary table will fail, if a local temporary table of the same name exists on that connection, as the new table cannot be uniquely identified by *owner.table*.

The information in the section “Unique table names on a connection [CR 455684]” on page 29 will be added to the following chapters and sections:

- In Chapter 6, “SQL Statements,” the “Usage” section in “DECLARE LOCAL TEMPORARY TABLE statement”
- In Chapter 6, “SQL Statements,” the “Usage” section in “CREATE TABLE statement”

9.7.48 UNION ALL view performance impacted by SORT DESC [CR 454910]

Certain optimizations, such as pushing a DISTINCT operator into a UNION ALL view, are not applied when the ORDER BY is DESC because the optimization that evaluates DISTINCT below a UNION does not apply to DESC order. For example, the following query would impact performance:

```
SELECT DISTINCT state FROM testVU ORDER BY state DESC;
```

To work around this performance issue, queries should have the DISTINCT operator evaluated before the ORDER BY, where the sort order is ASC and the optimization can be applied:

```
SELECT c.state FROM (SELECT DISTINCT state
                     FROM testVUA) c
ORDER BY c.state DESC;
```

See “SELECT statement” in *Sybase IQ Reference Manual*.

9.7.49 INSERT...LOCATION supports 20MB large object (LOB) data [CR 453439]

Sybase IQ 12.7 ESD #3 increases the limit on the length of large object data (LOB) that can be retrieved from a remote database using INSERT...LOCATION from 32767 bytes to 20MB (20971520 bytes). In Chapter 6, “SQL Statements,” in the “Usage” section of the “INSERT statement” section, the last sentence in the paragraph describing inserting Adaptive Server Enterprise TEXT and IMAGE data should be changed, so that the paragraph says:

“Sybase IQ does not support the Adaptive Server Enterprise data type TEXT, but you can execute INSERT...LOCATION (Syntax 3) from both an IQ CHAR or VARCHAR column whose length is greater than 255 bytes, and from an ASE database column of data type TEXT. ASE TEXT and IMAGE columns can be inserted into columns of other Sybase IQ data types, if Sybase IQ supports the internal conversion. All data inserted is silently right truncated at 20971520 bytes.”

9.7.50 Approximate numeric data type clarification [CR 452284]

Replace the section “Approximate numeric data types” in Appendix A, “Compatibility with Other Databases” with the following section.

Approximate numeric data types

Adaptive Server Enterprise differs from Adaptive Server Anywhere and Sybase IQ in how the FLOAT(precision) data type is interpreted: that is, when to create a 4-byte data type and when to create an 8-byte data type. When a precision is not specified, the Adaptive Server Anywhere and Sybase IQ option FLOAT_AS_DOUBLE makes the FLOAT keyword behave like the Adaptive Server Enterprise FLOAT keyword.

For Adaptive Server Enterprise, the precision specified in FLOAT(precision) means decimal precision. For Adaptive Server Anywhere and Sybase IQ, precision is an integer expression that specifies the number of bits used in the mantissa of the floating point number.

For more information on the Sybase IQ FLOAT data type, see “Numeric data types” in Chapter 4, “SQL Data Types” of the *Sybase IQ Reference Manual*. For more information on the Sybase IQ FLOAT_AS_DOUBLE option, see “FLOAT_AS_DOUBLE option [TSQL]” in Chapter 2, “Database Options” of the *Sybase IQ Reference Manual*.

9.7.51 Implicit conversion between BIT and CHAR/VARCHAR data types supported [CR 451862]

Conversion between
BIT and
CHAR/VARCHAR
data types

Sybase IQ 12.7 ESD #2 supports implicit conversion between BIT and CHAR and BIT and VARCHAR data types for arithmetic, comparison, INSERT, and UPDATE operations. The following section will be added to the chapter “SQL Data Types” in the *Sybase IQ Reference Manual*.

Sybase IQ supports implicit conversion between BIT and CHAR and BIT and VARCHAR data types for comparison operators, arithmetic operations, and INSERT and UPDATE statements.

BIT to CHAR, BIT to VARCHAR, CHAR to BIT, and VARCHAR to BIT conversion examples These examples illustrate both implicit and explicit conversions between BIT and CHAR and BIT and VARCHAR data types.

Given the following tables and data:

```
CREATE TABLE tchar(c1 CHAR(9))
CREATE TABLE tvarchar(c2 VARCHAR(9))
CREATE TABLE tbar(c2 BIT)
CREATE TABLE tbit(c2 BIT)

INSERT tbar VALUES(1)
INSERT tbar VALUES(0)
```

Implicit conversion of BIT to CHAR / CHAR to BIT and implicit conversion of BIT to CHAR:

```
INSERT tchar SELECT c2 FROM tbar
SELECT c1, char_length(c1) FROM tbar

c1,char_length(tchar.c1)
-----
'1',9
'0',9
```

Implicit conversion of CHAR to BIT:

```

INSERT tbit SELECT c1 FROM tchar
SELECT c2 FROM tbit

c2
--
true
false

```

Implicit conversion of BIT to VARCHAR / VARCHAR to BIT and implicit conversion of BIT to VARCHAR:

```

INSERT tvvarchar SELECT c2 FROM tbar
SELECT c2, char_length(c2) FROM tvvarchar

c2,char_length(tvvarchar.c2)
-----
'1',1
'0',1

```

Implicit conversion of VARCHAR to BIT:

```

INSERT tbit SELECT c2 FROM tvvarchar
SELECT c2 FROM tbit

c2
--
true
false

```

Explicit conversion of BIT to CHAR / CHAR to BIT and explicit conversion of BIT to CHAR:

```

INSERT tchar SELECT CONVERT (CHAR(9), c2) FROM tbar
SELECT c1, char_length(c1) FROM tchar

c1,char_length(tchar.c1)
-----
'1',9
'0',9

```

Explicit conversion of CHAR to BIT:

```
INSERT tbit SELECT CONVERT (BIT, c1) FROM tchar
SELECT c2 FROM tbit

c2
--
true
false
```

Explicit conversion of BIT to VARCHAR / VARCHAR to BIT and explicit conversion of BIT to VARCHAR:

```
INSERT tvvarchar SELECT CONVERT(VARCHAR(9), c2)
FROM tbar
SELECT c2, char_length(c2) FROM tvvarchar

c2,char_length(tvvarchar.c2)
-----
'1',1
'0',1
```

Explicit conversion of VARCHAR to BIT:

```
INSERT tbit SELECT CONVERT (BIT, c2) FROM tvvarchar
SELECT c2 FROM tbit

c2
--
true
false
```

9.7.52 VARCHAR and trailing blanks correction [CR 451561, CR 450225]

Appendix A, “Compatibility with Other Sybase Databases” incorrectly states in the section “Character data types” that Adaptive Server Enterprise trims trailing blank spaces from VARCHAR values, but Sybase IQ does not. The correct statement is:

Adaptive Server Enterprise trims trailing blank spaces from VARCHAR values. Sybase IQ trims trailing blanks from VARCHAR values depending on the form of the data and the operation. For details, see “Character data types” in Chapter 4, “SQL Data Types.”

VARCHAR data and trailing blanks

The following information has been added to the section “Character data types” in Chapter 4, “SQL Data Types”:

Data inserted via INSERT, UPDATE, or LOAD TABLE can be in one of the following forms:

- Enclosed in quotes
- Not enclosed in quotes
- Binary

For a column of data type VARCHAR, trailing blanks within the data being inserted are handled as follows:

- 1 For data enclosed in quotes, trailing blanks are never trimmed.
- 2 Data not enclosed in quotes:
 - Trailing blanks are always trimmed on insert and update.
 - For a LOAD statement, you can use the STRIP ON/OFF LOAD option to specify whether to have the trailing blanks trimmed. The STRIP ON/OFF option applies only to variable-length non-binary data. For example, assume the following schema:

```
CREATE TABLE t( c1 VARCHAR(3) );
LOAD TABLE t( c1 ',' ) ... STRIP ON           // trailing blanks trimmed

LOAD TABLE t( c1 ',' ) ... STRIP OFF          // trailing blanks not trimmed

LOAD TABLE t( c1 ASCII(3) ) ... STRIP ON      // trailing blanks not trimmed
LOAD TABLE t( c1 ASCII(3) ) ... STRIP OFF     // trailing blanks trimmed

LOAD TABLE t( c1 BINARY ) ... STRIP ON        // trailing blanks trimmed
LOAD TABLE t( c1 BINARY ) ... STRIP OFF       // trailing blanks trimmed
```

- 3 For binary data, trailing blanks are always trimmed.

You should not depend on the existence of trailing blanks in VARCHAR columns, when you write your applications. If an application relies on trailing blanks, a CHAR column should be used instead of a VARCHAR column

9.7.53 sa_checkpoint_execute fails to copy .db and .log files [CR 450484]

When using `sa_checkpoint_execute` to copy *asiqdemo.** files to another directory, all files are copied except the *.db* and *.log* files. ASA error -910 is returned.

This error is not a product defect but a Windows limitation; the Windows copy command cannot copy catalog files while they are open by the database.

9.7.54 sa_checkpoint_execute revised example [CR 450192]

The example for the `sa_checkpoint_execute` system procedure has been revised as follows:

Assuming you have created a subdirectory named *backup*, the following statement issues a checkpoint, copies all of the *asiqdemo* database files to the backup subdirectory, and completes the checkpoint:

```
sa_checkpoint_execute 'cp asiqdemo.* backup/'
```

9.7.55 xp_cmdshell supported by Sybase IQ [CR 449977]

`xp_cmdshell` is a system extended procedure that allows a database server to execute external shell commands. This functionality was available in Sybase IQ 12.7 and prior versions.

For details, see *Adaptive Server Anywhere SQL Reference Manual* on the Sybase Infocenter at <http://infocenter/help/index.jsp>.

Note In Chapter 7 of *Adaptive Server Anywhere SQL Reference*, the section “System extended stored procedures” states that users must be granted EXECUTE permission or have DBA authority. Several of the subsequent procedures, however, show permissions listed as none, including `xp_cmdshell`. `xp_cmdshell` requires DBA authority.

9.7.56 TOP range modification [CR 448015]

In Chapter 6, “SQL Statements,” SELECT statement, the TOP range listed in “FIRST or TOP number-of-rows” has been modified. The range is 1–2147483647 (not 1-32767).

Also, removed a clause and sentence from the fourth paragraph in the same section: “except that the maximum number of rows returned for TOP is 32767. ROW_COUNT does not have an upper limit for the number of rows returned.”

9.7.57 Corrections for new stored procedures permissions [CR 447961]

Several of the new 12.7 stored procedures have incorrect information for permissions in Chapter 10, “System Procedures.” The correct permission information for each stored procedure is listed in the following table.

<i>Procedure name</i>	<i>Permissions</i>
sp_iqdatatype	DBA authority required. Users without DBA authority must be granted execute permission in order to run the stored procedure.
sp_iqevent	DBA authority required. Users without DBA authority must be granted execute permission in order to run the stored procedure.
sp_iqhelp	DBA authority required. Users without DBA authority must be granted execute permission in order to run the stored procedure.
sp_iqjoinindex	DBA authority required. Users without DBA authority must be granted execute permission in order to run the stored procedure.
sp_iqpkeys	DBA authority required. Users without DBA authority must be granted execute permission in order to run the stored procedure.
sp_iqprocedure	DBA authority required. Users without DBA authority must be granted execute permission in order to run the stored procedure.
sp_iqprocparm	DBA authority required. Users without DBA authority must be granted execute permission in order to run the stored procedure.
sp_iqshowpsex	DBA authority required. Users without DBA authority must be granted execute permission in order to run the stored procedure.
sp_iqsysmon	DBA authority required. Users without DBA authority must be granted execute permission in order to run the stored procedure.
sp_iqversionuse	DBA authority required. Users without DBA authority must be granted execute permission in order to run the stored procedure.
sp_iqwho	DBA authority required. Users without DBA authority must be granted execute permission in order to run the stored procedure.

9.7.58 ALTER SERVER CONNECTION CLOSE syntax [CR 447958, CR 429434]

The syntax for the ALTER SERVER CONNECTION CLOSE clause was omitted from the *Sybase IQ Reference Manual*. The following information should be added to the “ALTER SERVER statement” section in Chapter 6, “SQL Statements.”

- Add the following clause to the ALTER SERVER syntax:

```
[ CONNECTION CLOSE [ CURRENT | ALL | connection-id ] ]
```

- Add the following examples to the “Examples” section:

The following example closes all connections to the remote server named `rem_test`.

```
ALTER SERVER rem_test  
CONNECTION CLOSE ALL
```

The following example closes the connection to the remote server named `rem_test` that has the connection ID 142536.

```
ALTER SERVER rem_test  
CONNECTION CLOSE 142536
```

- Add the following information to the “Usage” section:

CONNECTION CLOSE clause When a user creates a connection to a remote server, the remote connection is not closed until the user disconnects from the local database. The CONNECTION CLOSE clause allows you to explicitly close connections to a remote server. You may find this useful when a remote connection becomes inactive or is no longer needed.

The following SQL statements are equivalent and close the current connection to the remote server:

```
ALTER SERVER server-name CONNECTION CLOSE  
  
ALTER SERVER server-name CONNECTION CLOSE CURRENT
```

You can close both ODBC and JDBC connections to a remote server using this syntax. You do not need DBA authority to execute either of these statements.

You can also disconnect a specific remote ODBC connection by specifying a connection ID, or disconnect all remote ODBC connections by specifying the ALL keyword. If you attempt to close a JDBC connection by specifying the connection ID or the ALL keyword, an error occurs. When the connection identified by *connection-id* is not the current local connection, the user must have DBA authority to be able to close the connection.

9.7.59 INFER_SUBQUERY_PREDICATES option default is ON [CR 447453]

The default value of the INFER_SUBQUERY_PREDICATES database option is ON in Sybase IQ 12.7. (This is a change from Sybase IQ version 12.6, in which the default value of this option is OFF.)

In Chapter 2, “Database Options,” the default value of INFER_SUBQUERY_PREDICATES in Table 2-1 should be changed to ON.

In Chapter 2, “Database Options,” the default value of INFER_SUBQUERY_PREDICATES in the section “INFER_SUBQUERY_PREDICATES option” should be changed to ON.

9.7.60 Column heading inconsistent when calling sp_iqrelocate using isql vs. dbisql [CR 446968]

In ESD #2, the following procedures were modified to make their SELECT clause column headings consistent with the result clause column headings:

```
sp_iqcontext
  ConnOrCursor      char(22),
  ConnHandle        unsigned bigint,
  Name              varchar(255),
  Userid            varchar(255),
  numIQCursors      unsigned int,
  IQthreads         unsigned int,
  TxnID             unsigned bigint,
  ConnOrCurCreateTime datetime,
  IQconnID          unsigned bigint,
  IQGovernPriority   char(14),
  CmdLine           varchar(4096)
```

```

sp_iqdbspaceinfo
    DbspaceName      char(128),
    Object            varchar(128),
    MinBlk            unsigned bigint,
    MaxBlk            unsigned bigint,
    ObjSize           varchar(5),
    DbospaceSize      varchar(5)

sp_iq_process_post_login
    warning_text      varchar(255),
    warning_action    int

sp_iqindexadvice
    Advice            varchar(1024),
    NInst             unsigned bigint,
    LastDT            datetime

sp_iqindexinfo
    Object            varchar(128),
    DbospaceName      char(128),
    ObjSize           varchar(5),
    DBSpPct           smallint,
    MinBlk            unsigned bigint,
    MaxBlk            unsigned bigint

sp_iqrelocate
    Object            varchar(128),
    "Blocks Relocated" bigint,
    Status            char(10)

sp_iqshowcompression
    Column            char(128),
    Compression       char(3)

sp_iqspaceinfo
    Name              varchar(128),
    NBlocks           unsigned bigint,
    DbospaceName      char(128)

sp_iqtransaction
    Name              varchar(256),
    Userid            varchar(256),
    TxnID             unsigned bigint,
    CmtID             unsigned bigint,
    VersionID         unsigned bigint,
    State             char(12),
    ConnHandle        unsigned bigint,
    IQConnID          unsigned bigint,
    MainTableKBCr     unsigned bigint,
    MainTableKBDr     unsigned bigint,

```

TempTableKBCr	unsigned bigint,
TempTableKBDr	unsigned bigint,
TempWorkspaceKB	unsigned bigint,
TxnCreateTime	char(26),
Dbremote	bit,
CursorCount	unsigned bigint,
SpCount	unsigned bigint,
SpNumber	unsigned bigint

9.7.61 Updatable cursor text error [CR 446647]

The following paragraph was incorrectly included in the DECLARE CURSOR statement [ESQL] [SP] and has been removed for the next release:

This information should be added to “Updatable cursor limitations” in the “Usage” section for the DECLARE CURSOR statement description in the “SQL Statements” chapter.

9.7.62 LOAD TABLE FORMAT syntax correction [CR 445737]

In the section “LOAD TABLE statement” in Chapter 6, “SQL Statements,” the syntax for the FORMAT option should not contain single quotes. The correct syntax is:

```
LOAD [ INTO ] TABLE [ owner.]table-name
...(load-specification [, ...])
...FROM { 'filename-string' | filename-variable } [, ...]
...[ CHECK CONSTRAINTS { ON | OFF } ]
...[ DEFAULTS { ON | OFF } ]
...QUOTES OFF
...ESCAPES OFF
...[ FORMAT { ascii | binary } ]
...
```

In the “Usage” section of the “LOAD TABLE statement” section, the fifth paragraph in the description of the QUOTES option should be replaced with the following paragraph, which does not have single quotes in the FORMAT option:

“The data extraction facility provides options for handling quotes (TEMP_EXTRACT_QUOTES, TEMP_EXTRACT_QUOTES_ALL, and TEMP_EXTRACT_QUOTE). If you plan to extract data to be loaded into an IQ table and the string fields contain column or row delimiter under default ASCII extraction, use the TEMP_EXTRACT_BINARY option for the extract and the FORMAT binary and QUOTES OFF options for LOAD TABLE.”

9.7.63 Description of TDS in error [CR 443108]

In Chapter 2, “Database Options,” the description for TDS is incorrect. The paragraph should read:

Connections to Sybase IQ can be made through the TDS (tabular data stream) protocol (Open Client and jConnect™ for JDBC™ connections) or through the Sybase IQ protocol (ODBC, Embedded SQL).

9.7.64 sp_iqclient_lookup procedure [CR 443012]

The sp_iqclient_lookup procedure was added in Sybase IQ 12.7 ESD #2.

Function	Allows a client application to determine the Sybase IQ user account responsible for a particular data stream, as observed in a network analyzer originating from a specific client IP address/port.
Syntax	sp_iqclient_lookup [<i>IPaddress</i>], [<i>Port</i>], [<i>UserID</i>]
Parameters	IPaddress Specifies the IP address of the originating client application Port Specifies the port number of the originating client application UserID Specifies the Sybase IQ user ID
Remarks	The sp_iqclient_lookup procedure takes the client IP address and port number and returns a single row containing Number (the connection ID), IPaddress, Port, and UserID.

```
1> sp_iqclient_lookup 158.76.235.71,3360
2> go
```

Number	IPaddress	Port	UserID
-----	-----	----	-----
15	158.76.235.71	3360	rdeniro

An optional third argument can be passed to select just the UserID. If no arguments are passed then the procedure returns all current logins with their IP addresses and port numbers. For example:

```
sp_iqclient_lookup
```

Number	IPaddress	Port	UserID
-----	-----	----	-----
11	162.66.131.36	2082	mbrando
21	162.66.100.233	1863	apacino
22	162.66.100.206	8080	jcaan
23	162.66.100.119	6901	rduvall
24	162.66.100.125	7001	dkeaton
25	162.66.100.124	6347	jcazale

```
(6 rows affected)
```

```
(return status = 0)
```

If a client application is not using TCP/IP or for internal connections, the address appears as 127.0.0.1.

Note This information is available for logged on users only. No historical login data is kept on the server for this purpose.

Permissions

DBA.

Side effects

The `sp_iqclient_lookup` stored procedure may impact server performance, which will vary from one installation to another. Finding the login name entails scanning through all current active connections on the server; therefore, the impact may be greater on servers with large numbers of connections. Furthermore this information cannot be cached as it is dynamic — sometimes highly dynamic. It is, therefore, a matter for the local system administrator to manage the use of this stored procedure, as well as monitor the effects on the server, just as for any other client application that uses server facilities.

Examples

Shows IP addresses for UserID jcazale:

```
sp_iqclient_lookup null, null, jcazale
```

Number	IPaddress	Port	UserID
-----	-----	----	-----
11	162.66.131.36	2082	jcazale
15	164.66.131.36	1078	jcazale

Shows IP addresses from client IP 162.66.131.36:

```
sp_iqclient_lookup '162.66.131.36'
```

Number	IPaddress	Port	UserID
-----	-----	----	-----
11	162.66.131.36	2082	jcazale
12	162.66.131.36	1078	jcaan

Note The result is empty when the user specifies an incorrect argument.

9.7.65 OFF not an option for CLOSE_ON_ENDTRANS [CR 444249]

In Chapter 6, “SQL Statements,” “Set statement [T-SQL],” OFF is not an option for CLOSE_ON_ENDTRANS and has been removed from table 6-14.

In the subsequent bulleted list, CLOSE_ON_ENDTRANS should read as follows:

SET CLOSE_ON_ENDTRANS { ON } When CLOSE_ON_ENDTRANS is set to ON (the default and only allowable value), cursors are closed at the end of a transaction. With the option set ON, CLOSE_ON_ENDTRANS provides Transact-SQL-compatible behavior.

9.7.66 ESCAPE_CHARACTER option for system use only [CR 440545]

The ESCAPE_CHARACTER option is reserved for system use only. Do not change the setting of this option.

Chapter 2, “Database Options” requires the following changes:

- Table 2-1: Remove the ESCAPE_CHARACTER option and its values.
- Table 2-3: Remove the value and default of the ESCAPE_CHARACTER option.
- Add the following sentence to the note that follows Table 2-3:

The ESCAPE_CHARACTER option is reserved for system use only. Do not change the setting of this option.

9.7.67 Unowned LVC cells message clarification [CR 439550]

Three “LVC cells” messages were added to the output of `sp_iqcheckdb` in Sybase IQ 12.6 ESD #2, but descriptions were not included in the manuals. An example message:

```
sp_iqcheckdb 'check index
EFG2JKL.ASIQ_IDX_T208_C504_FP'
-----
Index Statistics:
** Inconsistent Index:
abcd.EFG2JKL.ASIQ_IDX_T208_C504_FP ***** FP
Indexes Checked: 1
** Unowned LVC Cells: 212 *****
```

The `sp_iqcheckdb` messages include:

- Unowned LVC cells
- Duplicate LVC cell rows
- Unallocated LVC cell rows

These messages indicate inconsistencies with a VARCHAR or CLOB column. Unowned LVC cells represent a small amount of unusable disk space and can safely be ignored. Duplicate and Unallocated LVC cells are serious errors that can only be resolved by dropping the damaged columns.

To drop a damaged column, create a new column from a copy of the old column, then drop the original column and alter rename the new column to the old column.

Note LVC is a VARCHAR column with a width greater than 255. CLOB also uses LVC.

9.7.68 ALTER TABLE column-modification syntax [CR 438076]

In Chapter 6, “SQL Statements,” the ALTER TABLE statement “Parameters” section should contain the following syntax:

```
column-modification:
SET DEFAULT default-value
| DROP DEFAULT
| ADD [ CONSTRAINT column-constraint-name ] CHECK ( condition )
| { DELETE | DROP } CONSTRAINT column-constraint-name
| { DELETE | DROP } CHECK
| SET COMPUTE ( expression )
| DROP COMPUTE
```

9.7.69 LOAD TABLE supports BCP files as input [CR 437332]

Sybase IQ 12.7 ESD #2 supports using a BCP character file as input to the LOAD TABLE command. For syntax and usage updates to the “LOAD TABLE statement” section in Chapter 6, “SQL Statements” of the *Sybase IQ Reference Manual*, see “LOAD TABLE supports BCP files as input [CR 437332]” on page 45.

9.7.70 Setting DBISQL options TEMPORARY [CR 435977]

In Chapter 2, “Database Options,” Syntax 1 in the section “DBISQL options” should contain the keyword TEMPORARY:

```
SET [ TEMPORARY ] OPTION  
... [ userid. | PUBLIC. ]option-name = [ option-value ]
```

Also the following sentence belongs at the beginning of the “DBISQL options” Description section:

“Syntax 1 with the TEMPORARY keyword cannot be used between the BEGIN and END keywords of a compound statement.”

9.7.71 IQGovernPriority value of -1 [CR 435854]

In Chapter 10, “System Procedures,” the complete description of the IQGovernPriority column in Table 10-6 of the sp_iqcontext stored procedure is:

“A value that indicates the order in which the queries of a user are queued for execution. In the range of allowed values, 1 indicates high priority, 2 (the default) medium priority, and 3 low priority. A value of -1 indicates that IQGovernPriority does not apply to the operation. This value is set per user with the database option IQGOVERN_PRIORITY. For details, see “Setting query priority” in Chapter 3, “Optimizing Queries and Deletions” of the *Sybase IQ Performance and Tuning Guide*.”

9.7.72 NEWID() function in column default values [CR 435215]

In Chapter 5, “SQL Functions,” the last sentence in the “Usage” section of the section “NEWID function [Miscellaneous]” should be replaced with the sentence:

“You can use a value generated by the NEWID function as a column default value in a Sybase IQ table.”

This change is a Sybase IQ 12.7 ESD #1 enhancement.

9.7.73 INSERT ... LOCATION example correction [CR 432896]

In Chapter 6, “SQL Statements,” replace the third example in the “Examples” section of the “INSERT statement” section with the following example:

- Inserts data from the `l_shipdate` and `l_orderkey` columns of the `lineitem` table from the Sybase IQ database `asiqdet` on the remote server `detroit` into the corresponding columns of the `lineitem` table in the current database:

```
INSERT INTO lineitem
    (l_shipdate, l_orderkey)
LOCATION 'detroit.asiqdet'
PACKETSIZE 512
' SELECT l_shipdate, l_orderkey
FROM lineitem '
```

9.7.74 INSERT ... LOCATION statement clarification [CR 432896]

In Chapter 6, “SQL Statements,” replace the first three paragraphs of the Syntax 3 description in the “Usage” section of the “INSERT statement” section with the following four paragraphs:

“Syntax 3 INSERT...LOCATION is a variation of Syntax 2 that allows you to insert data from an Adaptive Server Enterprise or Sybase IQ database. The *servername.dbname* specified in the LOCATION clause identifies the remote server and database for the table in the FROM clause. To use Syntax 3, the Adaptive Server Enterprise or Sybase IQ remote server to which you are connecting must exist in the Sybase Open Client *interfaces* or *sql.ini* file on the local machine.

The following Open Client restrictions apply to queries using Syntax 3: You can insert a maximum of 2147483647 rows.

The SELECT statement can be delimited by either curly braces or straight single quotation marks. (Curly braces represent the start and end of an escape sequence in the ODBC standard, and might generate errors in the context of ODBC.)

The local Sybase IQ server connects to the server and database you specify in the LOCATION clause. The results from the queries on the remote tables are returned and the local server inserts the results in the current database. If you do not specify a server name in the LOCATION clause, Sybase IQ ignores any database name you specify, since the only choice is the current database on the local server.”

9.7.75 Message log wrapping enabled automatically [CR 431841]

The second paragraph of the description in “IQMSG_LENGTH_MB option” in Chapter 2, “Database Options” should read as follows:

Initially, message log wrapping is disabled (IQMSG_LENGTH_MB is 0); messages are always appended to the end of the file, and the file continues to grow. If the disk fills, Sybase IQ enables log wrapping automatically.

9.7.76 More information on LOAD_MEMORY_MB [CR 426783]

The amount of virtual memory used by load command is a function of the total number of bytes (as defined by the table schema) for all columns being loaded. For example, consider the following schema:

```
CREATE table xx(c1 integer, c2 varchar(300), c2 double)
```

The total number of bytes is 312:

integer	4 bytes
varchar(300)	300 bytes
double	8 bytes

total	312 bytes

Using this example, the load would require approximately 140MB:

```
312 * 45 * 10000 => 140MB
```

The amount of virtual memory used can become quite large if many columns (such as in a very wide table) are loaded at once. The wider the table, the more the load memory. The more users doing loads, the more heap/load memory is allocated outside IQ.

There are several courses of action you can take if you encounter the following error:

```
"All available virtual memory has been used ..."
```

You can set an upper limit on the amount of virtual memory a LOAD command can use by setting LOAD_MEMORY_MB to a non-zero value, with 2000MB the maximum allowed value.

You can also adjust BLOCK FACTOR or BLOCK SIZE LOAD command options. These command options default to 10000 and 500000, respectively, but you can set them to any number. Setting them lower forces the load to use less virtual memory.

You can also resort to loading a subset of the columns at a time, which is referred to as a partial-width load.

9.7.77 Subquery support in CREATE VIEW statement [CR 426715]

In Chapter 6, “SQL Statements,” the CREATE VIEW statement omits a subquery restriction. In the “Usage” section, the AS option should read as follows:

AS The SELECT statement on which the view is based must not contain an ORDER BY clause, a subquery in the SELECT list, or a TOP or FIRST qualification. It may have a GROUP BY clause and may be a UNION.

9.7.78 Result data type of a REPLACE function [CR 417507]

In Chapter 5, “SQL Functions,” the REPLACE function included a new Usage section as follows:

The result data type of a REPLACE function is a LONG VARCHAR. If you use REPLACE in a SELECT INTO statement, you must have a Large Objects Management option license or use CAST and set REPLACE to the correct data type and size.

There are two ways to work around this issue:

- Declare a local temporary table and then do an INSERT:

```
DECLARE local temporary table #mytable
(name_column char(10)) on commit preserve rows;
INSERT INTO #mytable SELECT REPLACE(name, '0', '1')
FROM dummy_table01;
```

- Use CAST:

```
SELECT CAST(replace(name, '0', '1') AS Char(10))
into #mytable from dummy_table01;
```

In ESD #1, this usage also applies to the CONVERT function and string concatenation operators.

The result data type of an AES_ENCRYPT function may be a LONG VARBINARY. If you use AES_ENCRYPT in a SELECT INTO statement, you must have a Large Objects Management option license or use CAST and set AES_ENCRYPT to the correct data type and size.

9.7.79 WORD index and CONTAINS predicate support for character large object (CLOB) columns [CR 415547]

Sybase IQ 12.7 ESD #3 provides support for WORD (WD) indexes on LONG VARCHAR character large object (CLOB) columns and support of the CONTAINS predicate on LONG VARCHAR (CLOB) columns with a WD index.

The following sentence will be added to the “Description” section of the “sp_iqrebuildindex procedure” section in Chapter 10, “System Procedures”:

“sp_iqrebuildindex will rebuild a WD index on a column of data type LONG VARCHAR (CLOB).”

The first sentence following the syntax in the “CONTAINS conditions” section in the “Search conditions” section in Chapter 3, “SQL Language Elements” will be changed to say:

“The *column-name* must be either a CHAR, VARCHAR, or LONG VARCHAR (CLOB) column in a base table and must have a WD index.”

The following sentence will be added to the “Data types” section under “Notes” in the “CREATE INDEX statement” section in Chapter 6, “SQL Statements”:

Only the default and WD index types can be created on LONG VARCHAR data.

9.7.80 New database option WD_DELETE_METHOD [CR 415547]

Sybase IQ 12.7 ESD #3 provides the ability to specify the algorithm used during a delete operation from a WD index using the new database option WD_DELETE_METHOD. The description of the new option will be added to Chapter 2, “Database Options.” See “New database option WD_DELETE_METHOD [CR 415547]” on page 62 for details.

9.7.81 Views and procedures not supported by sp_iqrename [CR 408049]

The sp_iqrename procedure does not support renaming views or procedures.

9.7.82 LOAD TABLE supports loading default values [CR 396843]

Sybase IQ 12.7 ESD #4 supports loading a default value into a column as specified in the LOAD TABLE statement. You can now load a default value into a column, even if the column does not have a default value defined in the table schema. The LOAD TABLE syntax includes a new DEFAULT value specification in the column specification.

For syntax and usage updates to the section “LOAD TABLE statement” in Chapter 6, “SQL Statements” of the *Sybase IQ Reference Manual*, see “LOAD TABLE supports loading default values [CR 396843]” on page 47.

9.7.83 ASE_BINARY_DISPLAY option setting incorrect

In Chapter 2, “Database Options,” the setting for the ASE_BINARY_DISPLAY option is incorrect. The default setting for the option is OFF.

9.7.84 ENABLE_THREAD_ALLOWANCE option setting incorrect

The default setting for the ENABLE_THREAD_ALLOWANCE option is incorrect in Chapter 2, “Database Options.” The default setting is ON.

9.7.85 INSERT ... LOCATION syntax error

An error appears in the INSERT statement in Chapter 6, “SQL Statements.”

Braces should appear in the example, as follows:

```
INSERT INTO foo
LOCATION 'ase.database'
QUOTED_IDENTIFIER ON {select "c1" from xxx};
```

9.7.86 UUIDTOSTR function syntax correction

In Chapter 5, “SQL Functions,” the syntax of the UUIDTOSTR function should be changed to:

UUIDTOSTR (*uuid-expression*)

9.7.87 NEWID, STRTOUID, and UUIDTOSTR are native IQ functions

In Appendix A, “Compatibility with Other Sybase Databases,” the last bullet in the section “SQL functions” should be:

- Adaptive Server Anywhere supports the NEWID, STRTOUID, and UUIDTOSTR functions; Adaptive Server Enterprise does not. These are native functions in Sybase IQ, so CIS functional compensation performance considerations do not apply.

9.7.88 IQ_USER_LOGIN_INFO_TABLE system table column descriptions

In Chapter 9, “System Tables,” the IQ_USER_LOGIN_INFO_TABLE system table column descriptions should be changed as follows:

- Remove the description of the column password_warning_days.
- Add the description:
password_created Date and time of password creation.

9.8 Introduction to Sybase IQ

This section contains updates to *Introduction to Sybase IQ*.

9.8.1 WORD index and CONTAINS predicate support for character large object (CLOB) columns [CR 415547]

Sybase IQ 12.7 ESD #3 provides support for WORD (WD) indexes on LONG VARCHAR character large object (CLOB) columns and support of the CONTAINS predicate on LONG VARCHAR (CLOB) columns with a WD index.

Table 5-1 “Sybase IQ column index types” in Chapter 5, “Indexing and Loading Data” will be modified, so that the Purpose column for the WD index says:

“Used to index keywords by treating the contents of a CHAR, VARCHAR, or LONG VARCHAR column as a delimited list.”

9.9 Installation and Configuration Guide

For updates about installation and migration, including database upgrades, see “Special installation and migration instructions” on page 11.

9.10 Large Objects Management in Sybase IQ

This section contains updates to *Large Objects Management in Sybase IQ*.

9.10.1 INSERT...LOCATION supports 20MB large object (LOB) data [CR 453439]

Sybase IQ 12.7 ESD #3 increases the limit on the length of large object data (LOB) that can be retrieved from a remote database using INSERT...LOCATION from 32767 bytes to 20MB (20971520 bytes).

In Chapter 2, “Binary Large Object (BLOB) data,” in the section “Large Object data types LONG BINARY and BLOB,” the paragraph that describes inserting Adaptive Server Enterprise IMAGE data should be changed to say:

“An Adaptive Server Enterprise IMAGE column can be inserted into a LONG BINARY column using the INSERT...LOCATION command. All IMAGE data inserted is silently right truncated at 20971520 bytes.”

In Chapter 3, “Character Large Object (CLOB) data,” in the section “Large Object data types LONG VARCHAR and CLOB,” the paragraph that describes inserting Adaptive Server Enterprise TEXT data should be changed to say:

“An Adaptive Server Enterprise TEXT column can be inserted into a LONG VARCHAR column using the INSERT...LOCATION command. All TEXT data inserted is silently right truncated at 20971520 bytes.”

9.10.2 LOAD TABLE FORMAT syntax correction [CR 445737]

In the section “Loading large object data” in Chapter 6, “Moving Large Object Data,” the syntax for the FORMAT option of the LOAD TABLE statement should not contain single quotes. The correct syntax is:

```
LOAD [ INTO ] TABLE [ owner.]table-name
... ( column-name load-column-specification [, ...] )
... FROM 'filename-string' [, ...]
...[ QUOTES { ON | OFF } ]
...ESCAPES OFF
...[ FORMAT { ascii | binary } ]
...[ DELIMITED BY 'string' ]
...
```

9.10.3 LOAD TABLE supports BCP files as input [CR 437332]

Sybase IQ 12.7 ESD #2 supports using a BCP character file as input to the LOAD TABLE command. For syntax and usage updates to the LOAD TABLE information in Chapter 6, “Moving Large Object Data” of the *Large Objects Management in Sybase IQ*, see “LOAD TABLE supports BCP files as input [CR 437332]” on page 45.

For LOAD TABLE FORMAT BCP, the load specification may contain only column names, NULL, and ENCRYPTED. This means that you cannot use secondary files when loading LONG BINARY and LONG VARCHAR columns using the LOAD TABLE FORMAT BCP option.

9.10.4 WORD index and CONTAINS predicate support for character large object (CLOB) columns [CR 415547]

Sybase IQ 12.7 ESD #3 provides support for WORD (WD) indexes on LONG VARCHAR character large object (CLOB) columns and support of the CONTAINS predicate on LONG VARCHAR (CLOB) columns with a WD index.

The paragraph describing index usage in the “Large Object data types LONG VARCHAR and CLOB” section in Chapter 3, “Character Large Object (CLOB) data” will be replaced with the following paragraph:

“A WORD (WD) index can be created on a LONG VARCHAR column. Other non-FP index types and join indexes cannot be constructed on a LONG VARCHAR column.”

9.10.5 LIKE predicate for character large object (CLOB) columns [CR 415546]

Sybase IQ 12.7 ESD #3 includes support of the LIKE predicate on LONG VARCHAR (CLOB) columns.

The following sentence will be added to the section “LONG VARCHAR columns in queries” in Chapter 3, “Character Large Object (CLOB) data”:

“You can use the LIKE predicate to search for a pattern on a LONG VARCHAR column. All patterns of 126 characters or less are supported. Patterns of length greater than 254 characters are not supported. Some patterns of length between 127 and 254 characters are supported, depending on the contents of the pattern.”

The following paragraph will be added to the section “LONG BINARY columns in queries” in Chapter 2, “Binary Large Object (BLOB) data”:

“LIKE predicates are not supported on LONG BINARY (BLOB) columns. If you attempt to search for a pattern in a LONG BINARY column using a LIKE predicate, the error “Invalid data type comparison in predicate...” is returned.”

The following error will be added to Chapter 8, “Error and Warning Messages”:

Error 1001013

Message text

Invalid data type comparison %1

Item	Value
SQLCode	-1001013L
Constant	EMSG_TYPECOMPAREERROR
SQLState	QFA13
ODBC 2 State	ERROR
ODBC 3 State	ERROR
Sybase Error Code	20522
Severity Code	14
Parameter 1	location of the exception

Probable cause

This error is reported if you attempt to search for a pattern in a LONG BINARY column using a LIKE predicate. LIKE predicates are not supported on LONG BINARY (BLOB) columns.

9.10.6 String search functions for large object columns [CR 415543]

Sybase IQ 12.7 ESD #3 provides string search capabilities on large objects. Sybase IQ now includes support of string search functions for LONG VARCHAR (CLOB) and LONG BINARY (BLOB) data. The string search functions CHARINDEX and LOCATE have been enhanced to allow a LONG VARCHAR or a LONG BINARY column as the argument for the string expression to be searched and now return a 64 bit value for the position of the string in the LONG VARCHAR or LONG BINARY column.

The PATINDEX function has been enhanced to allow a LONG VARCHAR column as the argument for the string expression to be searched and now returns a 64 bit value for the position of the string in the LONG VARCHAR column.

For the descriptions of the enhanced string search functions that will be added to Chapter 4, “Function Support” of *Large Objects Management in Sybase IQ*, see “String search functions for large object columns [CR 415543]” on page 64.

The first sentence in the section “LONG BINARY columns in queries” in Chapter 2, “Binary Large Object (BLOB) data” will be modified to say:

“In WHERE clauses of the SELECT statement, LONG BINARY columns can only be used in IS NULL and IS NOT NULL expressions, in addition to the BYTE_LENGTH64, BYTE_SUBSTR64, BYTE_SUBSTR, BIT_LENGTH, OCTET_LENGTH, CHARINDEX, and LOCATE functions.”

The first sentence in the section “LONG VARCHAR columns in queries” in Chapter 3, “Character Large Object (CLOB) data” will be modified to say:

“In WHERE clauses of the SELECT statement, LONG VARCHAR columns can only be used in IS NULL and IS NOT NULL expressions, in addition to the BIT_LENGTH, OCTET_LENGTH, CHAR_LENGTH, CHAR_LENGTH64, SUBSTRING64, SUBSTRING, CHARINDEX, LOCATE, and PATINDEX functions.”

9.10.7 LOAD TABLE supports loading default values [CR 396843]

Sybase IQ 12.7 ESD #4 supports loading a default value into a column as specified in the LOAD TABLE statement. You can now load a default value into a column, even if the column does not have a default value defined in the table schema. The LOAD TABLE syntax includes a new DEFAULT value specification in the column specification.

For syntax and usage updates to the LOAD TABLE information in Chapter 6, “Moving Large Object Data” of the *Large Objects Management in Sybase IQ*, see “LOAD TABLE supports loading default values [CR 396843]” on page 47.

9.11 Troubleshooting and Recovery Guide

This section contains updates to the *Sybase IQ Troubleshooting and Recovery Guide*.

9.11.1 Restrictions on forced recovery [CR 610313]

The first paragraph and note in “Starting servers in forced recovery mode,” in Chapter 2, “System Recovery and Database Repair,” in the *Sybase IQ Troubleshooting and Recovery Guide* are incorrect. Please disregard the following:

“If a server fails to start with an exception or an assert when opening a database, start the server with forced recovery. Forced recovery allows the server to start if the allocation map or checkpoint information is inconsistent. In this mode, options display information about inconsistencies. You can also specify options to repair such inconsistencies.

Use forced recovery only when normal database recovery fails to restore the database to a running state.”

The correct information is as follows:

“Follow this procedure only if you see s_buf or free list errors during failure recovery. If SQL Anywhere recovery errors occur, -iqfrec will not correct the problem.”

See also “Do not use server startup switch -f [CR 610313]” on page 88.

9.11.2 Insufficient threads [CR 478486]

The following should be appended to “Insufficient threads” in the section “Resource issues,” in Chapter 1, “Troubleshooting Hints,” in the *Sybase IQ Troubleshooting and Recovery Guide*.

If the server runs out of threads or if sufficient threads are not available to a connection during a restore, Sybase IQ may return the error `Ran out of threads. Start up server with more threads. (SQLCODE - 1012024)`. The RESTORE command will try to allocate a “team” of threads for the restore operation. Sybase IQ will try to allocate at least one thread per backup device plus two threads per CPU, plus one thread to the “team” for the restore. Make sure that enough threads have been allocated on a per connection and per team basis as well as to the server. See “MAX_IQ_THREADS_PER_TEAM option” and “MAX_IQ_THREADS_PER_CONNECTION option” in Chapter 2, “Database Options,” in the *Sybase IQ Reference Manual*.

9.11.3 Multiplex and single-node modes [CR 330391]

The following information replaces the section titled “Single-node mode” in the section “Recovering multiplex databases” in Chapter 2, “System Recovery and Database Repair.”

In multiplex mode, query servers continue to run while the write server issues DDL operations. The table version log (TLV log) stores information about DDL operations and communicates information about new table versions to the query servers. The write server preserves older table versions for as long as needed.

In single-node mode, no version logging takes place. Query servers shut down automatically and must be synchronized to restart.

The `-iqmpx_sn 1` server startup option starts the server in single-node mode. For syntax, see Chapter 1, “Running the Database Server” in *Sybase IQ Utility Guide*.

Warning! You must specify the override startup switch (`-iqmpx_ov 1`) and start in single-node mode (`-iqmpx_sn 1`) when starting a multiplex write server after any failure. Never use multiplex mode (the default) for recovery.

If a server runs out of space during a checkpoint operation, try restarting in single-node mode. For example, if a server ran out of space and was killed, it could precipitate a rare situation where normal recovery failed. Suppose that many versions were created at the server and they were all retained because the query server used an older version or there was a dbremote delay. If the write server had too many recovered transactions to checkpoint, it could run out of space during the recovery checkpoint. Subsequently, users could not connect or add dbspaces. Starting in single-node mode deallocates retained versions and allows checkpoints to succeed.

In order to resume multiplex operations, you must stop the write server, restart it in multiplex mode, and synchronize the query servers. For details, see “Synchronizing query servers” and “Using scripts to synchronize servers” in Chapter 5, “Working with Database Objects,” in the *Sybase IQ System Administration Guide*.

9.12 Utility Guide

This section contains updates to the *Sybase IQ Utility Guide*. These corrections will appear in the next release of the Sybase IQ documentation set.

9.12.1 Controlling stack size per thread [CR 481648]

The following should be added to Chapter 1, “Running the Database Server,” in the *Sybase IQ Utility Guide*.

-gss { *integer* | *integerK* | *integerM* } Sets the stack size for each SQL Anywhere thread in the server. Has no effect on Windows operating systems. The default stack size per SA thread is 64KB. The maximum size is 4MB.

9.12.2 Backing up the transaction log during database use [CR 477499]

The second example in “Backup utility (dbbackup)” in Chapter 3, “Database Administration Utilities,” in the *Sybase IQ Utility Guide* is incorrect. Please replace it with the following example.

The following example truncates the log of a live server, renames the old log prior to truncation, and creates a exact copy of the log in the target directory:

```
dbbackup -t -r -o dbbackup.out -c
"eng=taj_asiqdemo;dbn=asiqdemo;uid=DBA;pwd=SQL"
logbackup
```

9.12.3 Maximum Catalog page size clarification [CR 468276]

The description of the example for the @filename switch in the section on “Server command-line switches” states the maximum Catalog page size is set to 4MB. This is incorrect. The description should say the maximum Catalog page size is set to 4096 bytes.

In the description of the -gp parameter, the sentence “Sybase highly recommends that you set the -gp switch to 4096” has been changed to “Sybase highly recommends that you set the -gp switch to at least 4096.”

These corrections will appear in the next release of the Sybase IQ documentation set.

9.12.4 Correction to -iqlocalreplay server switch description [CR 464167]

In Chapter 1, “Running the Database Server,” the iqlocalreplay server switch description in the section “Starting the database server” contains errors. The last sentence should read, “IQ tables and user-defined data types must always be enabled for replay.”

The description of -iqlocalreplay none should read “No actions that affect grants, procedures or views are replayed.”

The following was omitted from the description of -iqlocalreplay userandpwd:

The userandpwd value must be used with the keyword all or grants to take effect. For example, set -iqlocalreplay "all+userandpwd" for the log to replay all DDL commands and actions for userandpwd. Set -iqlocalreplay "grants+userandpwd" for the log to replay GRANT and REVOKE commands and actions for userandpwd.

9.12.5 Correction to value of -iqro server switch [CR 461571]

In Chapter 1, “Running the Database Server,” the value of the -iqro server switch is incorrect in the section “Starting the database server.” The correct value is -iqro 1.

9.12.6 No master database in iq_bcp [CR 447965]

In Chapter 3, “Database Administration Utilities,” section “Bulk Copy utility (iq_bcp),” database_name in “Options” contains an error. It should read as follows:

database_name Is optional if the table being copied is in your default database. Otherwise, you must specify a database name.

The words, “or in master,” have been removed from the first sentence.

9.12.7 Using ping utility to test connection [CR 463766]

In the Chapter 3 section “Ping utility (dbping),” the following information should be appended to the description of the Load specified library (-l) option.

You can use dbping to verify connectivity with the ODBC Driver Manager on UNIX systems. Unlike dbisql or other tools, dbping allows you to explicitly test the components as you would with a third-party tool. For example:

```
dbping -m -c "dsn=dsnname" -l /full/path/to/libodbc.so
```

where *libodbc.so* is the third-party tool's ODBC driver.

Note This CR was formerly listed as 442804.

9.12.8 Ping utility is threaded [CR 442796]

Chapter 3, “Database Administration Utilities,” contains an error in the section “Ping utility (dbping).” The statement that you must use *dbping_r* in order to use a threaded connection library on UNIX is incorrect. The dbping utility is threaded by default, and *dbping_r* is not required or installed by Sybase IQ.

9.12.9 LOAD TABLE supports BCP files as input [CR 437332]

Add the following note to the description of the iq_bcp utility in the section “Bulk Copy utility (iq_bcp)” in Chapter 3, “Database Administration Utilities” of the *Sybase IQ Utility Guide*:

Note The LOAD TABLE FORMAT BCP feature is an alternative to the iq_bcp in utility and will perform better than iq_bcp in all cases. For more information on LOAD TABLE FORMAT BCP, see “LOAD TABLE statement” in Chapter 1, “SQL Statements” of *Reference: Statements and Options*.

9.12.10 Effect of start_asiq -ti on connection to local server using shared memory [CR 422369]

The start_asiq -ti server switch, which specifies the client idle time before shutdown, has no effect on connections to a local server using shared memory.

This information will be added to the -ti switch description in Chapter 1, “Running the Database Server” section “Starting the database server” under “Usage.”

9.12.11 Loading a NULL value into VARCHAR using iq_bcp

The following information will be added to the “Constraints” section of the “Bulk Copy utility (iq_bcp)” section in Chapter 3, “Database Administration Utilities”:

“If you load a NULL value into a VARCHAR column that does not accept NULL values, iq_bcp transforms the NULL into a “. This transformation is consistent with Adaptive Server Enterprise behavior.”

9.13 Performance and Tuning Guide

This section contains updates to the *Sybase IQ Performance and Tuning Guide*.

9.13.1 Correct BNF grammar in OLAP chapter [CR 487599]

Grammar rule 53 beneath “BNF grammar for OLAP function” did not describe the sort specification list. The corrected grammar reads:

```
<ORDER SPECIFICATION> ::= ORDER BY <SORT SPECIFICATION
LIST>
  <SORT SPECIFICATION LIST> ::= <SORT SPECIFICATION>
  [ { <COMMA> <SORT SPECIFICATION> } ... ]
  <SORT SPECIFICATION> ::= <SORT KEY>
  [ <ORDERING SPECIFICATION> ] [ <NULL ORDERING> ]
  <SORT KEY> ::= <VALUE EXPRESSION>
  <ORDERING SPECIFICATION> ::= ASC | DESC
  <NULL ORDERING> := NULLS FIRST | NULLS LAST
```

This information will be available in the next GA release.

9.13.2 Corrections to OLAP chapter [CR 479852, 447363]

Chapter 4 “Using OLAP,” contains documentation errors, which are corrected in the next version of the *Sybase IQ Performance and Tuning Guide*:

- Cube example 1 column headings have been corrected to match the example syntax: avg salary column head now reads AVERAGE.
- Cube example 2 column headings have been corrected to match the example syntax: state now reads c_state, sex now reads c_gender, dept_id now reads c_dept, count now reads count(*), and avg salary now reads AVERAGE.
- In “Windowing” under “Analytic functions,” the following syntax is repeated in two locations:

```
OLAP_FUNCTION() OVER (PARTITION BY col1, col2...)
```

The redundant paragraph and syntax will be removed in the next version.

- In “An OLAP window’s three essential parts,” the first line of the syntax contains a repeated word, WINDOW, in the window definition list. The second instance will be removed in the next version.
- In “Ascending and descending order for value-based frames” the syntax is incorrect and should be ordered by DESC instead of ASC. The corrected syntax should read as follows:

```
ORDER BY year DESC range BETWEEN CURRENT ROW and 1  
FOLLOWING
```

- In “Windowing aggregate functions,” “Windowing aggregate example 1” contains two errors. Column head Year should be Y, and the dates in the discussion output should match the syntax; 1993 should be 2000 and 1994 should be 2001.
- In “Windowing aggregate functions,” “Windowing aggregate example 2” contains incorrect column headings. Employee should say E_name and Salary should say Sal.

9.13.3 Incorrect function listed in OLAP chapter [CR 479239]

Table 4-3, incorrectly lists LENGTH (*string-expression*) where it should list LN (*numeric-expression*).

The corrections will appear in the next release of the Sybase IQ documentation.

9.13.4 Managing thread usage [CR 478486]

The last bullet in the section “Sybase IQ options for managing thread usage,” in Chapter 5, “Managing system resources,” in the *Sybase IQ Performance and Tuning Guide* should read as follows:

To set the maximum number of threads a single user will use, issue the command `SET OPTION MAX_IQ_THREADS_PER_CONNECTION`. Some operations try to allocate and use a “team” of threads. To set then number of threads available to a “team”, issue the command `SET OPTION MAX_IQ_THREADS_PER_TEAM`. These options can be used to control the amount of resources a particular operation consumes. For example, the DBA can set this option before issuing an `INSERT`, `LOAD`, `BACKUP`, or `RESTORE` command.

9.13.5 Restricting concurrent queries with `-iqgovern` [CR 432651]

In Chapter 5, “Managing System Resources,” the description for `-iqgovern` under “Restricting concurrent queries,” is incorrect. It has been corrected in the next version of the manual to read as follows:

The `-iqgovern` switch lets you specify the number of concurrent queries on a particular server. This is not the same as the number of connections, which is controlled by your license. By specifying the `-iqgovern` switch, you can help IQ optimize paging of buffer data out to disk, and avoid overcommitting memory. The default value of `-iqgovern` is $(2 \times \text{the number of CPUs}) + 10$. You may need to experiment to find an ideal value. For sites with large numbers of active connections, try setting `-iqgovern` slightly lower.

9.13.6 Estimating temporary cache size [CR 425940]

The following was omitted from “Managing buffer caches” in Chapter 5, “Managing System Resources” in the *Sybase IQ Performance and Tuning Guide*.

For guidelines about estimating temporary cache size, see the Hardware Sizing Guide for Sybase IQ 12.6 and 12.7 at <http://www.sybase.com/detail?id=1062015>. This technical white paper describes the relationship between hardware configuration and Sybase IQ memory management.

9.13.7 New database option `WD_DELETE_METHOD` [CR 415547]

The following section will be added to “Optimizing delete operations” in Chapter 3, “Optimizing Queries and Deletions” to supplement the information on the `WD_DELETE_METHOD` database option introduced in Sybase IQ 12.7 ESD #3.

For more information on the `WD_DELETE_METHOD` database option, see “New database option `WD_DELETE_METHOD` [CR 415547]” on page 62.

WD delete operations

Sybase IQ chooses one of three algorithms to process delete operations on columns with a WD (Word) index:

- **Small delete** Small delete for WD provides optimal performance when the rows deleted contain few distinct words, so that not many WD pages need to be visited. The WD small delete algorithm performs an ordered access to the WD. Worst case I/O is bounded by the number of index pages. Small delete incorporates the cost of sorting the words and record IDs in the records to delete.
- **Mid delete** Mid delete for WD is a variation of WD small delete, and is useful under the same conditions as small delete, that is, when the rows deleted contain few distinct words. Mid delete for WD sorts only words in the records to delete. This sort is parallel, with parallelism limited by the number of words and CPU threads available. For Word index, the mid delete method is generally faster than small delete.
- **Large delete** Large delete for WD provides optimal performance when the rows deleted contain a large number of distinct words, and therefore need to visit a large number of “groups” in the index. The large delete scans the WD in order, until all rows are deleted. Worst case I/O is bounded by the number of index pages. Large delete is parallel, but parallelism is limited by the internal structure of the index and the distribution of groups from which to delete.

WD delete costing

The WD delete cost model considers many factors including I/O costs, CPU costs, available resources, index metadata, and parallelism.

You can use the `WD_DELETE_METHOD` database option to control WD delete performance.

Using WD delete performance option

The value of the parameter specified with the `WD_DELETE_METHOD` option forces the use of the specified delete algorithm as follows:

- 0 = Mid or large delete as selected by the cost model
- 1 = Small delete
- 2 = Large delete
- 3 = Mid delete

9.14 IQ plug-in for Sybase Central

This section contains updates to the IQ plug-in for Sybase Central.

9.14.1 WORD index and CONTAINS predicate support for character large object (CLOB) columns [CR 415547]

Sybase IQ 12.7 ESD #3 provides support for WORD (WD) indexes on LONG VARCHAR character large object (CLOB) columns and support of the CONTAINS predicate on LONG VARCHAR (CLOB) columns with a WD index.

The table “Sybase IQ column index types” in Lesson 5 in the Sybase Central tutorial and in the “Managing indexes” topic will be modified, so that the Purpose column for the WD index says:

“Used to index keywords by treating the contents of a CHAR, VARCHAR, or LONG VARCHAR column as a delimited list.”

10. 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.

Before you contact
Technical Support

Technical Support needs information about your environment to resolve your problem. Before contacting Technical Support, run the getiqinfo script to collect as much information as possible automatically. You may also need to collect some information manually. Providing this information helps expedite the resolution of your problem.

In the following list, * indicates items collected by getiqinfo:

- Type of hardware, amount of memory, number of CPUs*
- Operating system and version (for example, Red Hat Linux 2.4.9-e.40smp)*
- Operating system patch level*
- Front-end tool used (for example, Brio Query)
- Connectivity protocol used (for example, ODBC, JDBC, TDS)
- Open Client version
- Configuration type (single user or multiuser)

- Message log file* (*very important*)
Named *dbname.iqmsg*, located by default in the directory where you started the database server
- Stack trace file for the date and time this problem occurred
Named *stktrc-YYYYMMDD-HHMMSS_#.iq*, located in the directory where you started the database server*
- Command or query that produced the error
- Query plan* (recorded in *.iqmsg* file)

Note The query plan is collected automatically by `getiqinfo`. If you collect information manually, you must type the following commands and then rerun the command that produced the error:

```
SET TEMPORARY OPTION Query_Plan = 'ON'
SET TEMPORARY OPTION Query_Detail = 'ON'
```

The plan will be in the message log file.

If you have performance problems, set the following option:

```
SET TEMPORARY OPTION Query_Plan_After_Run = 'ON'
```

This enables Technical Support to see which steps in the query processing used the time.

- Server logs
 - For UNIX and Linux, *ASIQ-12_7/logfiles/<servername>.00n.stderr* and *ASIQ-12_7/logfiles/<servername>.00n.srvlog**
 - On Windows platforms,
*ASIQ-12_7/logfiles\<servername>.00n.srvlog**
- Startup and connection option settings, from the configuration file (by default, *dbname.cfg*)*
- Database option settings and output from `sa_conn_properties`* (if the server is still running)
- Schema and indexes for the database
- Output from `sp_iqstatus` and `sp_iqcheckdb`
- On multiplex databases, you must execute `getiqinfo` on the write server and/or query servers, and the following information is also collected:

- `servername.out` on query servers
- `write_server_name/repDirs/logfiles/servername.dbrlog`

A checklist for recording this information for Technical Support is provided at the end of this release bulletin.

For more information on `getinfo`, see “Collecting diagnostic information using `getinfo`” in the *Sybase IQ Troubleshooting and Recovery Guide*.

Online support

You can find additional help from the Sybase online support Web site at <http://www.sybase.com/support/>. MySybase, a link at the top of the page, is a free service that allows you to search through closed support cases, the latest software bulletins, and resolved and known problems, using a view customized for your needs. You can even open a Technical Support case online.

11. 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 may also contain 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 may 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.

Note The SyBooks browser software runs on Windows and Linux platforms. Users with non-Linux UNIX platforms must use Acrobat Reader to open PDF files on the SyBooks CD.

- 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/>.

- Infocenter is an online version of SyBooks that you can view using a standard Web browser. To access the Infocenter Web site, go to Sybooks Online Help at <http://infocenter.sybase.com/help/index.jsp>.

11.1 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 Partner Certification Report.
- 3 In the Partner Certification Report filter select a product, platform, and timeframe and then click Go.
- 4 Click a Partner 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.

11.2 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, type 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.

11.3 IQ Newsgroup

Sybase newsgroups provide a means for users to exchange information over the Internet. The newsgroup for Sybase IQ is sybase.public.iq.

For information on subscribing to Sybase newsgroups, configuring your newsreader or Web browser, and for guidelines on posting, go to Newsgroups at <http://www.sybase.com/support/newsgroups>.

Note Sybase newsgroups are moving to a new, more fault-tolerant environment. If you already subscribe to Sybase newsgroups, you need to replace your existing newsgroup headers with the new ones using the steps that follow.

❖ **Replacing existing newsgroup information**

- 1 Remove the forums.sybase.com account from your newsreader.
- 2 Add the forums.sybase.com account again.
All headers will be marked as unread.
- 3 Resubscribe to your newsgroups, following the instructions at Newsgroups at <http://www.sybase.com/support/newsgroups>.
- 4 Download headers into those newsgroups.

11.4 IQ User's Group

The IQ User's Group provides a forum for any IQ users to exchange information on IQ. To have your name added to the IQ User's Group list, send e-mail to iqug@odscompanies.com.

12. 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.

Sybase IQ 12.7 and the HTML documentation have 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.

For Section 508 compliant core documentation, see your Sybooks CD or the Infocenter Web site at <http://infocenter.sybase.com/help/index.jsp>.

For Section 508 compliant installation guides and release bulletins, see the HTML files on your Getting Started CD.

The online help for this product is also provided as HTML-based JavaHelp, which you can navigate using a screen reader.

For information about accessibility support in the Sybase IQ plug-in for Sybase Central, see “Using accessibility features” in the *Introduction to Sybase IQ*. The online help for this product, which you can navigate using a screen reader, also describes accessibility features, including Sybase Central keyboard shortcuts and using this product without a mouse.

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, and Section 508 compliance statements for Sybase IQ and other Sybase products.

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, or to pronounce certain characters or punctuation such as underscore. Consult the documentation for your tool. See the following table for pronunciation guidelines for important terms.

Table 13: ABBR and ACRONYM values

Term	Element value
\$ASDIR/bin	dollar A S D I R slash bin
asiqdemo	A S I Q demo
ASIQPlugin.jar	A S I Q plug in dot jar
DBISQL	D B I sequel
dbname.db	D B name dot D B
dbname.log	D B name dot log
dbname.iqtmp	D B name dot I Q T M P
Note You can apply this technique to other file names with extensions.	
dbspace	D B space
DSEdit	D S edit
IDs	I deez
I/O	I O
iq_dummy	I Q underscore dummy
IQ_SYSTEM_MAIN	I Q underscore system underscore main
libtcl.cfg	L I B T C L dot C F G
start_asiq	start underscore A S I Q
start_asiq.exe	start underscore A S I Q dot e x e
SQL	sequel (not S Q L, except when explaining how to pronounce SQL, or when it is what you type, like a password. JAWS reads SQL as S Q L by default.)
®	registered trademark
™	trademark
→	right arrow

13. Checklist: information for Technical Support

You can run the `getiqinfo` script to collect much of this information.

Information requested	Value
type of hardware	
amount of memory	
number of CPUs	
operating system name and version (for example, HP Itanium HP-UX 11.23)	
operating system patch level	
front end tool used (for example, Brio Query)	
connectivity protocol used (for example, ODBC, JDBC, TDS)	
Open Client version	
configuration type (single or multiuser)	
message log file (dbname.iqmsg)	
log file (.svrlog and .stderr)	
stack trace file (stktrc-YYYYMMDD-HHNNSS_#.iq)	
command or query that produced the error	
startup option settings	
connect option settings	
database option settings	
schema and indexes for the database	
sp_iqstatus output	
query plan: set options (Query_Plan, Query_Detail, Query_Plan_After_Run), rerun command or query	

