

# Release Bulletin Sybase® ETL 4.9

Document ID: DC00789-01-0490-01

Last revised: September 2009

<b>Topic</b>	<b>Page</b>
1. Accessing current release bulletin information	2
2. Product summary	2
2.1 Sybase ETL Development	2
2.2 Sybase ETL Server	3
2.3 Interfaces	5
2.4 Databases	6
2.5 Compatible products	7
3. Changed functionality in this version	7
4. Special installation instructions	8
5. Known problems	8
5.1 Problems with ETL Server	8
5.2 Problems with ETL Development and components	16
5.3 Problems with internationalization	28
5.4 Problems using ETL with Sybase IQ 12.7	30
5.5 Problems with Non-ETL products	31
6. Documentation updates and clarifications	35
6.1 Displaying ETL Development, product manuals, and demos	36
6.2 Installation Guide	36
6.3 Users Guide	36
7. Technical support	38
8. Other sources of information	38
8.1 Sybase certifications on the Web	39
8.2 Sybase EBFs and software maintenance	40

# 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 Sybase ETL and language and click Go.
- 3 Select Sybase ETL 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 ETL 4.9, which includes Sybase ETL Development and Sybase ETL Server.

### 2.1 Sybase ETL Development

The platform and operating systems supported by ETL Development are:

- Windows XP Professional with Service Pack 3 – 32-bit
- Windows XP Professional with Service Pack 2 – 64-bit
- Windows XP Professional N – 32-bit
- Windows Vista – 32-bit and 64-bit Windows Vista Business, Business N, and Enterprise editions
- Windows Vista Ultimate Service Pack 1 – 32-bit and 64-bit
- Windows 2003 Service Pack 2 – 32-bit and 64-bit Windows 2003 Standard Enterprise, and Data Center editions
- Windows 2008 – 32-bit and 64-bit Standard, Enterprise, and Data Center editions

## 2.2 Sybase ETL Server

This section describes platforms, operating systems, databases, and database interfaces that ETL Server supports.

### 2.2.1 Operating systems

Table 1 on page 4 lists the platform and operating systems supported by ETL Server.

**Table 1: Sybase ETL Server platforms and operating systems**

<b>Platform</b>	<b>Version</b>
HP Itanium	<ul style="list-style-type: none"> <li>• 11.23 – 64-bit</li> <li>• 11.31 – 64-bit</li> </ul>
IBM AIX	<ul style="list-style-type: none"> <li>• 5.3 – pSeries 64-bit</li> <li>• 6.1 – pSeries 64-bit</li> </ul>
Microsoft Windows	<ul style="list-style-type: none"> <li>• Windows 2003 Service Pack 2 – 32-bit and 64-bit Windows 2003 Standard, Enterprise, and Data Center editions</li> <li>• Windows XP Professional Service Pack 3 – 32-bit</li> <li>• Windows XP Professional Service Pack 2 – 64-bit</li> <li>• Windows XP Professional N – 32-bit</li> <li>• Windows Vista – 32-bit and 64-bit Windows Vista Business, Business N, and Enterprise editions</li> <li>• Windows Vista Ultimate Service Pack 1 – 32-bit and 64-bit</li> <li>• Windows 2008 – 32-bit and 64-bit Standard, Enterprise, and Data Center editions</li> </ul>
Red Hat Enterprise Linux	<ul style="list-style-type: none"> <li>• 4.0 x86 – 32-bit Advanced Server and Workstation Editions</li> <li>• 4.0 – 64-bit</li> <li>• 5.0 – 32-bit, 64-bit</li> <li>• 4.0 on POWER – 64-bit</li> <li>• 5.0 on POWER – 64-bit</li> </ul>
Sun Solaris	<ul style="list-style-type: none"> <li>• 9 (SPARC) – 64-bit</li> <li>• 10 (SPARC) – 64-bit</li> <li>• 10 x86 – 64-bit</li> </ul>
SuSE Linux Enterprise Server	<ul style="list-style-type: none"> <li>• 9 Service Pack 4 – 32-bit, 64-bit</li> <li>• 10 – 32-bit, 64-bit</li> <li>• 9 Service Pack 4 on POWER – 64-bit</li> <li>• 10 on POWER – 64-bit</li> </ul>

For a complete list of supported operating systems on each platform, see the Sybase platform certifications Web site at <http://certification.sybase.com>.

## 2.3 Interfaces

The interfaces supported by ETL Development for connecting to destination or source databases from a component, are:

- Sybase
- DB2 – only supported for connecting to source databases.
- ODBC – you must install the ODBC driver on the same computer as Sybase ETL Development, and you must define a system data source name (DSN) for the target.
- Oracle – only supported for connecting to source databases.
- OLE DB – only supported for connecting to source databases.
- SQLite Persistent – is only supported for connecting to source databases.

See “Database connection settings,” in Chapter 5, “Components” in the *Sybase ETL 4.9 Users Guide*.

Table 2 on page 6 lists the interface drivers supported by ETL Server.

**Table 2: Interface driver versions for Sybase ETL Server**

Driver	Version
Sybase native (via Client-Library™)	15.0 ESD #6 (Windows)
	15.0 ESD #15 (UNIX and Linux)
	<b>Note</b> These driver version are installed with Sybase ETL installation.
Adaptive Server® Enterprise ODBC	15.00.00.325 (Windows only)
SQL Anywhere® ODBC	11.00.00 (IQ 12.7 only)
	<b>Note</b> As of version 10.0, Adaptive Server Anywhere has been renamed SQL Anywhere.
Sybase IQ 12.7 ODBC	11.00.00.1264
Sybase IQ 15.0 ODBC	11.00.00.1264
	11.00.00.5120
	11.00.01.5027
Sybase IQ 15.1 ODBC	11.00.01.5027
IBM DB2 native	8.1.8.762, 9.01.00.369
IBM DB2 ODBC	8.01.08.762, 9.01.00.369 (Windows only)
Microsoft SQL Server ODBC	2000.86.3959.00 (Windows only)
MySQL	5.1.4
Oracle native (via Oracle Call Interface (OCI))	11.01.00.06
Oracle ODBC	11.01.00.06 (Windows only)

## 2.4 Databases

Table 3 shows the repository, source, and destination databases supported by ETL Server.

**Table 3: Sybase ETL Server database support**

Database	Version	Repository	Source	Destination	Stage
Sybase Adaptive Server Enterprise	15.0.2 ESD #6	No	Yes	No	Yes
	15.0.1 ESD #4				
	12.5.4 ESD #8				
	15.0.3 ESD #1				

Database	Version	Repository	Source	Destination	Stage
Sybase SQL Anywhere Server	10.0	Yes	Yes	No	Yes
	11.0				
	<hr/> <b>Note</b> Sybase recommends to use SQL Anywhere 11.0 as an ETL repository. <hr/>				
Sybase IQ	12.7 ESD #5	No	Yes	Yes	Yes
	15.0				
	15.1				
IBM DB2 UDB	9.1	No	Yes	No	No
	8.2				
Microsoft SQL Server	2000	No	Yes	No	No
	2005 SP2				
MySQL	5	No	Yes	No	No
Oracle	10g	No	Yes	No	No
	11g				

## 2.5 Compatible products

Sybase ETL is compatible with the Sybase products listed in Table 4.

**Table 4: Compatible products**

Product	Version
Replication Server®	15.0 or later
Open Client/Server™	15.0 or later
Replication Agent™	15.0 or later
Sybase Software Asset Management (SySAM)	2.0

## 3. Changed functionality in this version

See the *Sybase ETL 4.9 New Features Guide* for descriptions of the new and changed functionality in Sybase ETL 4.9.

## 4. Special installation instructions

SQL Anywhere is the only supported repository in Sybase ETL, SQL Anywhere 11 is the recommended version. If you are using any other repository from an earlier version of ETL, you must migrate your existing ETL repository to SQL Anywhere 11.

Included with the Sybase ETL Development installation is a SQL Anywhere 11 repository for migration of an existing repository. See “Migrating an existing repository to SQL Anywhere,” in Chapter 4, “Upgrading” of the *Sybase ETL 4.9 Installation Guide*.

## 5. Known problems

This section describes known problems in Sybase ETL 4.9 and workarounds where available. These problems are marked with the corresponding Change Request (CR) numbers. Provide the CR number when you contact Sybase Technical Support regarding ETL issues.

### 5.1 Problems with ETL Server

This section documents known problems with Sybase ETL Server.

#### 5.1.1 Migrating ETL 4.5 and 4.5.1 repositories in Sybase IQ to SQL Anywhere 11

**[CR #589667]** If you are migrating an ETL 4.5 or 4.5.1 repository in Sybase IQ, the migrate repository job does not write to the repositoryNew data source.

*Workaround:*

- 1 In the SQL Anywhere 11 database, repositoryNew (repositiory49.db), drop the DS\_CHUNK, DS\_DATA, DS\_OBJECT, and TRON\_PERFORMANCE tables.

---

**Note** You may use the tools provided during the installation of SQL Anywhere 11, which is included with ETL Server.

---

To drop these tables using Sybase ETL Development:

- a Start ETL Development and log in to the demo database.
- b Select File | New | Project.

- c From the Destination tab, add the DB Data Sink Insert component and enter these values:

Property name	Value
Interface	ODBC
Host Name	repositoryNew
User	dba
Password	sql
Destination Table	TRON_PERFORMANCE

- d Click Save.
- e Right-click the new DB Data Sink Insert component and select Execute SQL Commands.

f Enter:

- Drop table DS\_CHUNK
- Drop table DS\_DATA
- Drop table DS\_OBJECT
- Drop table TRON\_PERFORMANCE

g Click Execute.

The tables are dropped. Close the project without saving it.

- 2 In the Repository Logon window, set up the repository connection:
- a Click Add and enter `myRepository_ETL49` as the name.
  - b Select ODBC as the interface.
  - c Select *repositoryNew* as the host.
  - d Enter `dba` as the user and `sql` as the password.
  - e Click Test Logon to test the connection. Click Yes when prompted to create repository tables.
- 3 Migrate your 4.5 or 4.5.1 repository using the steps described in “Migrating from an existing repository to a SQL Anywhere repository,” in Chapter 4, “Upgrading” in the *ETL 4.9 Installation Guide*.

### 5.1.2 Incremental loading fails if bulk copy and dynamic SQL options are enabled

**[CR #586690]** Incremental loading fails if bulk copy and dynamic SQL options are set to on in Replication Server.

*Workaround:* Set the bulk copy and dynamic SQL options to off in Replication Server.

### 5.1.3 Port Structure displays incorrect column length

**[CR #585915]** The column length that appears for the port structure in the Size column of the Port Structure Viewer does not match the value defined in the SQL statement used to generate the table.

*Workaround:* None.

### 5.1.4 Application stops responding using named pipes on UNIX

**[CR #585799]** If ETL Server is running on UNIX and you use named pipes to load data to IQ, you may encounter an error that causes the application to stop responding.

*Workaround:* Terminate the process and restart the application.

### 5.1.5 SQL Anywhere installation fails on Solaris

**[CR #585784]** SQL Anywhere 11 does not successfully install on Solaris 10.

*Workaround:* In the PATH variable, include `/bin` before `/usr/ucb/bin`.

### 5.1.6 No verification to check if a table name exists

**[CR #585059]** If you do not provide a new table name in the “schema.name” format and a table with that name already exists, errors occur during table creation.

*Workaround:* Specify the new table name in the “schema.name” format. For example: `dbo.table1`.

### 5.1.7 Scheduled task is not repeated after ETL Development or ETL Server is restarted

**[CR #584982]** If you restart ETL Development or ETL Server, the Repeat Task option specified for a schedule is ignored; the task is not repeated at the specified time.

*Workaround:* Create individual scheduled tasks for each repeat schedule.

### 5.1.8 Replication is not dropped successfully in Replication Server

**[CR #584555]** When you manually drop replication, the replication state changes to “Dropped.” However, sometimes replication may not actually be dropped in Replication Server.

*Workaround:* Manually drop replication again:

- 1 Right-click the CDC (Capture Data Changes) Provider Sybase Replication Server component.
- 2 Select Drop Replication.

### 5.1.9 Alert e-mail notification does not display the total number of rows processed

**[CR #584083]** If you have included the [ProcessRows] Event Property in your alert e-mail definition, the total number of rows that were processed is evaluated, instead of the total number of rows that were loaded into the destination tables. See “Configuring alerts for runtime events,” in Chapter 4, “Advanced Concepts and Tools” in the *Sybase ETL 4.9 Users Guide*.

*Workaround:* None.

### 5.1.10 Connection to Sybase IQ in Sybase Central fails if ETL is uninstalled

**[CR #582762]** If Sybase IQ 15.1 and ETL 4.9 are installed on the same machine, and you uninstall ETL, you cannot then use Sybase Central™ to connect the IQ demo database.

*Workaround:* Either:

- Copy the *dbodbc11.dll* under *c:\IQ\IQ-15\_1\bin32* to *c:\IQ\ETLDevelop49* after uninstalling ETL, or,
- Create an ODBC data source using the Sybase IQ driver, and select “ODBC data source name” from the Identification tab. Enter the data source name created when connecting Sybase Central to the IQ demo database.

### 5.1.11 Cannot start Sybase IQ demo database after installing ETL 4.9

**[CR #582749]** If you install the 64-bit edition of Sybase IQ 15.1 in the same directory where you have installed Sybase ETL Development and ETL Server, the IQ demo database may not start.

*Workaround:* Manually install the 64-bit Visual Studio redistribution package.

- 1 Restart your machine after installing Sybase ETL Server and ETL Development.
- 2 Download and install one of these:
  - For Visual Studio 2005 redistribution, go to the Microsoft Download Center at <http://www.microsoft.com/DOWNLOADS/details.aspx?familyid=EB4EBE2D-33C0-4A47-9DD4-B9A6D7BD44DA>.
  - For Visual Studio 2008 redistribution, go to the Microsoft Download Center at <http://www.microsoft.com/DOWNLOADS/details.aspx?familyid=BD2A6171-E2D6-4230-B809-9A8D7548C1B6>.
- 3 Restart the Sybase IQ 15 demo database.

### 5.1.12 Executing SQL statements fails when using ODBC connectivity

**[CR #581741]** SQL statement with IQ temporary options are not executed when ODBC is used to connect to IQ.

*Workaround:* Use the Sybase interface to connect to IQ.

### 5.1.13 Project execution fails on Solaris with Oracle as the database server

**[CR #580854]** Sybase ETL Server cannot connect to Oracle on Sun Solaris 10.0 with Oracle as the source database.

*Workaround:* Download and install the latest 11.1.0.7 patch for the Oracle database server.

### 5.1.14 ETL Server fails to restart if *svc.conf* file is not saved in UTF-8 format

**[CR #576388]** ETL Server fails to start if *svc.conf* file in the *etc* subdirectory is updated and saved in a non UTF-8 encoding format.

*Workaround:* Use a text editor to open the *svc.conf* file and change the file format to UTF-8 encoding.

### 5.1.15 Issues using ASE as source and ODBC as the interface

**[CR #575450]** If you are using Adaptive Server Enterprise as the source database and ODBC as the interface for connectivity, not all the tables created in the database are always returned. This issue is related to the user role. If you create tables without the `sa_role`, but later grant it, tables that were created before the `sa_role` was granted do not appear when you request for the list of tables.

*Workaround:*

- Grant access to the users for all tables. Use the ODBC interface to request the tables.
- Alternatively, specify the interface as Sybase and request the list of tables. After selecting the tables you need, change the interface back to ODBC.

### 5.1.16 Executing a job with more than 50 projects causes ETL Server to crash

**[CR #564661]** Sybase ETL Server crashes while executing a job with more than 50 projects, due to excessive memory consumption.

*Workaround:* To limit excessive memory consumption:

- 1 Select File | Preferences from the Sybase ETL Development window.
- 2 Select Performance Logging, and set the level to 'o'.

### 5.1.17 Creating a repository connection to SQL Anywhere 11 database using OCS may fail

**[CR #563581]** If you are using the latest EBF of SQL Anywhere 11 database and you try to create a repository connection to this database using OCS connectivity, you may encounter errors.

*Workaround:* Install the version of SQL Anywhere 11 that is packaged with ETL 4.9 or access SQL Anywhere 11 with ODBC connectivity.

### 5.1.18 Printing a project displays incorrect file path on Windows 2008

**[CR #560347]** On Windows 2008, the message that appears after a project prints successfully as XML, displays an incorrect file path.

*Workaround:* Go to the *reports* folder in the installation directory and manually open the file to view.

### 5.1.19 ETL Server fails to start on Sun Solaris

**[CR #559206]** ETL Server fails to start on Sun Solaris if you select “Install for all users” option during installation, and you have not set the HOME directory.

*Workaround:* Set the HOME directory variable to your correct home directory or a directory on which you have write permission.

### 5.1.20 SySAM license expiration warning message displays only in English

**[CR #557213]** On all UNIX platforms, except HP-UX, the Sybase Software Asset Management System (SySAM) license expiration warning message appears only in English, regardless of the language selected during installation.

*Workaround:* None.

### 5.1.21 ETL Server fails to execute project on HP-UX

**[CR #556856]** Project execution may fail on certain HP-UX 11v23 machines.

*Workaround:* Download and install the following patches:

Patch	Description
<i>PHCO_34974 s700_800 11.23</i>	cumulative SAM patch
<i>PHNE_33283 s700_800 11.23</i>	nettl(1M), netfmt(1M), and nettladm(1M) patch
<i>PHNE_34150 s700_800 11.23</i>	r-commands cumulative mega patch
<i>PHNE_35770 s700_800 11.23</i>	telnet kernel, telnetd(1M), telnet(1) patch
<i>PHNE_34698 s700_800 11.23</i>	ftpd(1M) and ftp(1) patch

### 5.1.22 *getaddrinfo()* function limitation on Solaris

**[CR #556653]** ETL Server displays error messages while using *getaddrinfo()* function on Solaris machine.

*Workaround:* Download and install the *#solaris\_10u5(s10u5\_07)* patch on your machine.

### 5.1.23 System error in JavaScript editor while debugging large scripts

**[CR #555171]** The default maximum size of the JavaScript engine runtime is 20000000 bytes, or 19.07MB. Debugging large scripts or long running scripts such as a script with many loop iterations, you may encounter a memory issue.

*Workaround:* Configure the maximum memory size for the JavaScript engine in the *Default.ini* file. To configure:

- Navigate to the *etc* directory of the installation folder and use a text editor to open the *Default.ini* file.
- In the [Scripting] section, add:

```
Runtime Memory = <a new number for memory in bytes>
```

### 5.1.24 Issue in monitoring loader components

**[CR #555170]** If you use a Web browser to monitor projects that include the IQ Loader File via Load Table and IQ Loader DB via Insert Location components, you may not be able to view the execution progress status.

*Workaround:* None.

### 5.1.25 Unable to monitor projects started from command line

**[CR #553669]** Web based remote monitoring fails to monitor the execution of a project that is started from command line. In ETL 4.9, you can only monitor jobs started from command line.

*Workaround:* Create a simple job and include the project in it that you want to monitor.

### 5.1.26 ETL cannot be installed in the same directory as Sybase IQ

**[CR #551068]** Installation fails when Sybase ETL and Sybase IQ are installed in the same directory. Since Sybase IQ creates an “uninstall” directory during installation, and Sybase ETL also creates a file with the same name, installation fails with an error message.

*Workaround:* Do not install ETL in the same directory as Sybase IQ.

### 5.1.27 Use server name caching to connect to a specific Sybase IQ server

**[CR #540023]** ETL may not connect to a specific Sybase IQ server using ODBC if there are multiple Sybase IQ servers with the same server name, database name, and port number.

*Workaround:* Use server name caching to connect to the target Sybase IQ server. Add the DoBroadcast=DIRECT communication parameter to the CommLinks connection parameter in the *.odbc.ini* file in your ETL Server:

```
CommLinks=tcip (DoBroadcast=DIRECT;host=iq server;port=2638)
```

See “Server name caching for faster connections,” in Chapter 3, “Sybase IQ Connections” in the *Sybase IQ 12.7 System Administration Guide*.

### 5.1.28 ETL Server stops responding if there is insufficient temp space in Sybase IQ

**[CR #539896]** ETL Server can stop responding if there is insufficient temp space on the source Sybase IQ database server when transferring a large amount of data from Sybase IQ.

*Workaround:* Increase the temp space for the Sybase IQ server to a minimum of 1GB.

### 5.1.29 Cannot run earlier versions of ETL Development after installing ETL 4.9

**[CR #535684]** You cannot run ETL 4.9 concurrently with earlier versions of ETL Development due to registry and resource conflicts. If you have an earlier version of ETL Development installed on your machine, then installing ETL 4.9 on the same machine will prevent the earlier version from starting or connecting with the Demo data sources.

*Workaround:* To continue using the earlier version of ETL Development, do not install ETL 4.9 on the same machine that has the earlier versions of ETL installed.

### 5.1.30 Exceeding connection limit locks connections to Sybase IQ

**[CR #496226]** When multiple connections from ETL to Sybase IQ remain open after projects and jobs execute, connections to Sybase IQ can lock when open connections exceed the Sybase IQ connection limit. Users receive error messages that ETL cannot retrieve column descriptions, ETL has exceeded the database connection limit, and ETL cannot connect to Sybase IQ.

*Workaround:* Increase the connection limit in Sybase IQ. See “Managing IQ user accounts and connections,” in Chapter 12, “Managing User IDs and Permissions” in the *Sybase IQ 12.7 System Administration Guide*.

## 5.2 Problems with ETL Development and components

This section documents known problems with ETL Development and components.

### 5.2.1 Structure Viewer does not display data type for the XML via SQL Data Provider component

**[CR #590001]** When you right-click the OUT-port of the XML via SQL Data Provider component to view its structure, the Data Type field in the Structure Viewer appears empty. This prevents subsequent components with an editable structure from getting the correct port structure.

*Workaround:* Manually assign port structure of subsequent components.

### 5.2.2 Oracle data changes are received even if commit time is later than project or job start time

**[CR #589603]** If the primary database is Oracle, the CDC Provider Sybase Replication Server component receives data changes even if the commit time is later than the job or project start time.

*Workaround:* Do not write to the primary database during project execution.

### 5.2.3 Replication Server does not recognize the interfaces name

**[CR #588669]** If the SYBETL\_VIR\_RDBMS entry in the interfaces files for the source database, Replication Server, and ETL Server, includes more than 20 lines, Replication Server does not recognize the interface name.

*Workaround:* Make sure that the SYBETL\_VIR\_RDBMS entry in the interface file contains fewer than 20 lines.

### 5.2.4 Demo database points to the data source of the old repository

**[CR #588134]** If you have previously installed ETL Development, and you uninstall it to install ETL Development 4.9 for all users, when any user other than the one who installed it tries to connect to the repository, the demo database connects to the data source of the ETL Development 4.8 repository.

*Workaround:*

- If the previous version of ETL Development was installed for all users, before upgrading, log in as each user and use the ODBC Data Source Administrator to delete the DEMO\_Repository, ETLDEMO\_DWH, ETLDEMO\_GER, and ETLDEMO\_US data sources. If you do not plan to uninstall the earlier version of ETL Development, you may want to rename the data sources instead of deleting them. For example: *<old name\_version number>*.

See Chapter 4, “Upgrading,” in the *ETL 4.9 Installation Guide*.

- If you have already installed the new version without renaming or deleting the old data sources, manually delete the old data sources. Follow the steps described in “Restoring the initial set of ODBC data sources of the demo repository,” in Chapter 3, “Postinstallation Tasks” in the *ETL 4.9 Installation Guide*.

---

**Note** Make sure that you specify the database files in the user data location, not in the installation location.

---

If you encounter errors when trying to delete the old data sources, you can overwrite the entries for each user, as follows:

- Save a copy of the *userdata.conf* file located in the *etc* folder of the installation directory.
- Edit the *userdata.conf* file to replace all occurrences of “once” with “always” in the ODBC section.
- For each user:
  - 1 Log in to the machine and start ETL Development, then shut down the application.
  - 2 Verify that the ODBC Data Source entries use SQL Anywhere 11 and point to the appropriate database files in your user data directory.

After all users have logged in to ETL Development, you can restore the original *userdata.conf* file.

See “Checking for a valid installation,” in Chapter 3, “Postinstallation Tasks” in the *ETL 4.9 Installation Guide*.

### 5.2.5 Errors encountered if interface used for connecting to IQ is changed

**[CR #587862]** If you change the interface used for connecting to the Sybase IQ database, an error may occur, and the component may fail to connect to IQ.

*Workaround:* Restart the grid engine after changing the interface.

### 5.2.6 Upsert function is processed even when Truncate is selected

**[CR #587639]** In spite of the Truncate option being selected, the target tables do not get cleared before the Pre-processing SQL code executes. Thus, the Upsert or Delete function is processed on a non-empty table, leading to errors.

*Workaround:* To clear and repopulate target table before processing Upsert or Delete functions, make sure Truncate is unselected, and perform clearing of the table in pre-SQL script.

### 5.2.7 Error encountered using DB Data Provider Full Load with schema option

**[CR #587092]** The schema of the DB Data Provide Full Load component does not support the “automatically add missing schema information” option. You may encounter an error message when a project with the DB Data Provide Full Load component is executed.

*Workaround:* Specify the schema in your query statement.

### 5.2.8 Text Data Provider cannot correctly read quote character and delimiters

**[CR #586711]** For the Text Data Provider component, if you specify a quote character for a value, and the first character of the value is same as the specified quote character, data is read incorrectly.

*Workaround:* None.

### 5.2.9 CDC Provider Replication Server component fails if grid engines are in multiple subnets

**[CR #585575]** The CDC Provider Replication Server component may not work if the grid engines are running in more than one subnet, or if the subnet is unstable.

*Workaround:* All grid engines that start the Replication CDC Service must be in the same subnet.

### 5.2.10 Error connecting to the CDC Provider Replication Server component

**[CR #584015]** Replication Server throws a connection error while connecting to the ETL CDC Provider Replication Server component, if you have not properly configured Replication Server, interfaces file, and Open Client/Server (OCS) library.

Although replication is created successfully, an error is encountered when the project is executed. In such a scenario, you are also unable to drop replication.

*Workaround:* Verify that:

- The interfaces file of Replication Server and ETL Server are configured correctly. See Chapter 5, “Components,” in the *Sybase ETL 4.9 Users Guide* for more information.
- All grid engines that start the Replication CDC Service are in the same subnet.
- The primary database has already been added to Replication Server.
- The grid engine name and the Replication CDC Service names are unique. To check, run:

```
GridNode --repcdcinstancename <my rep cdc instance name> --nodename <my node name>
```

### 5.2.11 Selecting the Delete function displays inconsistent result

**[CR #581370]** If you have selected the Delete function for the DB Bulk Load Sybase IQ component, the results that appear during simulation and execution may not be consistent.

*Workaround:* None.

### 5.2.12 Wait time ignored by projects before acquiring an exclusive lock

**[CR #580679]** For DB Data Sink, DB Staging, Loader, and DB Bulk Load Sybase IQ components, if you select the “IQ Lock Table in Exclusive Mode” option and specify the maximum blocking time that the project should wait before acquiring a lock, the project attempts to acquire the lock without waiting, and fails if the lock is acquired by another project.

*Workaround:* Serialize project execution in a job.

### 5.2.13 DB Staging component cannot bulk load large amount of data

**[CR #577621]** DB Staging component fails to bulk-load large amounts of data to Sybase IQ 15.0 database with the ODBC interface. Sybase ETL Server generates this error message:

```
std::bad_alloc
```

*Workaround:* Use small write block size for bulk-loading data.

### 5.2.14 INSERT operation does not work when the Upsert function is processed

**[CR #575891]** For the DB Bulk Load Sybase IQ component, when the Upsert function is processed, the DELETE operation succeeds, but the INSERT operation fails because of a datatype conversion error in one of the columns being inserted. If you select “Execute post-processing for successful execution” to commit the transaction, the unsuccessful INSERT operation is not committed. The original row that was supposed to have been updated is deleted from the table.

*Workaround:* Select “Execute post-processing for failed execution” to roll back the Delete operation and to reset the table to its original state.

### 5.2.15 Evaluate setting is enabled for Insert Options and Update Options property of DB Data Sink components

**[CR #572569]** If the Evaluate option is not enabled during project execution, square bracket expressions (SBN) are evaluated incorrectly.

*Workaround:* Select the Evaluate option and save your project to allow expressions within component property values.

### 5.2.16 DB Staging component fails to connect to Sybase IQ database with large data

**[CR #566615]** DB Staging component fails to connect to Sybase IQ database with ODBC interface when the source data is large.

*Workaround:* Set the Force\_No\_Scroll\_Cursors option to off in your Sybase IQ 15.0 database. See the *Sybase IQ Reference Manual*.

### 5.2.17 Insert Location component does not support Preprocessing SQL

**[CR #564099]** The Source Pre Processing SQL property of the IQ Loader DB via Insert Location component is ignored if the “Use remote server definition for source database” option is enabled.

*Workaround:* None.

### 5.2.18 Multiplex execution fails on Linux and UNIX machines that do not have ODBC drivers installed

**[CR #560953]** Project execution fails if you select the “Use IQ Multiplex” option on a UNIX or Linux machine that does not have SQL Anywhere 11 or Sybase IQ 15 ODBC drivers installed.

*Workaround:* Use the SQL Anywhere 11 ODBC drivers installed with Sybase ETL.

### 5.2.19 DB Bulk Load Sybase IQ component cannot load data on SuSE machines

**[CR #560814]** On SuSE, the DB Bulk Load Sybase IQ component fails to load data if the ODBC environment is not properly configured. When the project executes, Sybase ETL Server exits unexpectedly without any error message.

*Workaround:* Follow these steps:

- 1 Install the SQL Anywhere 11 or Sybase IQ 15 ODBC drivers.
- 2 Add ODBC to the LD\_LIBRARY\_PATH environment variable.
- 3 Link *libodbc.so* to *libdbodbc11.so* in the *ASA11/IQ15* installation directory.
- 4 Restart ETL Server.

### 5.2.20 Issues using multiple writers with IQ Loader DB via Insert Location component

**[CR #560036]** To use the multiple writers, you must have execute permission on the *sp\_iqtable* and *sp\_iqcolumn* stored procedure privilege in the target IQ database, as well as the appropriate permissions on create table and execute *sp\_iqstatistics*.

*Workaround:* Set the required permission using Sybase Central:

- 1 In Sybase Central, connect to the Sybase IQ 15.0 server as DBA user or a member of dbo role.
- 2 Expand Users & Groups, then select the user or group for whom you want to set the permission.
- 3 Right-click the user or group and select Properties.
- 4 Select the Permissions tab, then select Procedures & Functions to see a list of all the available permissions.
- 5 Select *sp\_iqtable* and *sp\_iqcolumn*, then click the corresponding Execute column to grant the user permission to execute the stored procedure in the IQ database.
- 6 Click OK to save the settings.

### 5.2.21 Cannot execute projects if database value is specified

**[CR #556364]** If you enable the “Use IQ Multiplex” option in the DB Bulk Load Sybase IQ, IQ Loader File via Load Tables, and IQ Loader DB via Insert Location components, to select writers in an IQ multiplex environment, specifying a database in the “Database” property field may display an error during project execution.

*Workaround:* Do not specify a database in the “Database” property field when “Use IQ Multiplex” option is selected.

### 5.2.22 Error messages display when you migrate tables with large number of columns

**[CR #549882]** If you migrate wide tables using the Migration wizard, you may see this error message at the end of the migration process:

```
An error occurred during execution of the engine.  
Commit Failed: Client Library Message: severity(0)  
layer(1) origin(1) number(50) Text: ct_cmd_drop(): user  
api layer: external error: The connection has been  
marked dead.
```

```
Cannot execute the last command.
```

```
Please refer to the Message section.
```

```
The migration failed due to the previous error. Unable  
to open the generated job.
```

Despite the error message, the table is successfully migrated to the target database, and you can manually open the generated job to transform the data.

*Workaround:* This error displays if you are using the Sybase interface to connect to the source database and if the OCS version being used is 15.0 ESD #13 or later. To prevent this error:

- Use the ODBC interface to connect to the source database.
- If you are using Sybase interface to connect to the source database, use OCS version 15.0 ESD #7.

### 5.2.23 Issues using named pipes on Windows Vista

**[CR #549539]** On Windows Vista, you cannot specify a pipe name for the Load Stage property of the DB Bulk Load IQ component.

*Workaround:* Change the firewall setting on the Vista machine.

- 1 Go to Start | Control Panel | Security Center | Windows Firewall.
- 2 Click the Exceptions tab.
- 3 Under Programs and Services, select the File and Printer Sharing checkbox and then click OK.

### 5.2.24 File names with multibyte characters cannot be loaded using DB Bulk Load Sybase IQ component

**[CR #549397]** The DB Bulk Load Sybase IQ component fails to load data from ETL generated files with multibyte characters in the name. You see this error message:

```
Could not execute statement. Right truncation of string data.
```

*Workaround:* Do not use multibyte characters as the Load Stage file name for the DB Bulk Load Sybase IQ component if your target database is Sybase IQ 15.0.

### 5.2.25 Project reports printed as XML fail to open on Windows Vista

**[CR #546658]** On Windows Vista Business 32-bit, Sybase ETL may generate errors when displaying project reports that are generated using XML output.

*Workaround:* Install Service Pack 1 on Windows Vista.

### 5.2.26 Using SQL Anywhere as the DB Staging database displays errors

**[CR #546257]** For the DB Staging component, if you are using a SQL Anywhere database with Sybase connectivity, and if you are using an OCS version later than 15.0 ESD #7, you may encounter this error:

```
"The connection has been marked dead."
```

---

**Note** If you encounter a similar error while using Sybase IQ as a staging database, follow the same workaround.

---

*Workaround:*

- Use SQL Anywhere with ODBC connectivity, or use Adaptive Server Enterprise.
- If you are using the Sybase interface, use OCS versions that is packaged with ETL 4.9:

- For Windows – 15.0 ESD #6.
- For UNIX and Linux – 15.0 ESD #15.

### 5.2.27 LOB in File mode does not work correctly

**[CR #543229]** Transferring LOB (large object) data into Sybase IQ database in File mode, converts the file's original content while extracting.

*Workaround:* Use “IQ Loader File via Load Table” or “IQ Loader DB via Insert Location” component to transfer LOB data into Sybase IQ.

See “LOAD TABLE statement” and “INSERT statement” sections, in the “SQL Statements” chapter of the *Sybase IQ Reference Manual* for more details.

### 5.2.28 ETL may stop responding if tables contain a large number of columns

**[CR #541647]** When you execute a project with tables containing thousands of columns, ETL Server may stop responding intermittently. Attempts to save the project may also cause ETL Development to stop responding intermittently.

*Workaround:* None.

### 5.2.29 Text Data Provider fails if “Skip First Rows=0” for FTP source

**[CR #540626]** The Text Data Provider component does not transfer data if you keep the default value of 0 (zero) in the Skip First Rows field when you enter a FTP URL as the text source for the component.

*Workaround:* Enter 1 in the Skip First Rows field, save the configuration, change the value for Skip First Rows to 0, and save the configuration again.

### 5.2.30 Insert Location project fails if network packet size setting is incorrect

**[CR #536684]** Projects that use Adaptive Server as the source and that contain the IQ Loader DB via Insert Location component fail if the network packet size for the component is larger than or equal to the maximum network packet size of the Adaptive Server.

*Workaround:* Before you execute the project:

- 1 Connect to the Adaptive Server and execute `sp_configure 'max network packet size'` to display the maximum network packet size Run Value for your Adaptive Server.

- 2 In ETL, enter a packet size for IQ Loader DB via Insert Location that is less than or equal to the Adaptive Server maximum network packet size Run Value obtained in step 1.

For example, if your Adaptive Server maximum network packet size Run Value is 2048 bytes, your ETL packet size for the IQ Loader DB via Insert Location component must be less than to equal to 4 as 4 translates to a Adaptive Server Run Value of  $4 * 512 = 2048$  bytes.

### 5.2.31 Insert Location component does not support encrypted passwords

**[CR #536482]** Sybase ETL 4.9 does not support the Encrypted Password option in the IQ Loader DB via Insert Location component connection information if the source database is Sybase IQ 12.7 or Sybase IQ 15.0.

*Workaround:* None.

### 5.2.32 The `uSetLocale` JavaScript function works only on Windows platforms

**[CR #531483]** The `uSetLocale` javascript function only works on Windows platforms. On UNIX platforms, setting `uSetLocale` to any language will still display output in English for `uMonthName`, `uMonthNameShort`, `uWeekdayName`, and `uWeekdayNameShort`.

*Workaround:* None.

### 5.2.33 Table and column names cannot be database keywords in Migration wizard

**[CR #496346]** In the Migration wizard, you cannot use the "variable" database keyword as the name of a table or column.

*Workaround:* None.

### 5.2.34 Extra characters truncated from the column length of the source table

**[CR #493550]** If the column length of a source table is more than the column length of a target table, ETL truncates the extra characters without any warning. This occurs only if you use Adaptive Server as the stage database in a Staging component, using Sybase interface.

*Workaround:* When transferring data, make sure the column length of the target table is equal to or more than the column length of the source table.

### 5.2.35 Splitter components cannot inherit port structure

**[CR #492443]** The Data Splitter JavaScript and Copy Splitter components do not inherit the structure of the existing ports when they are reassigned.

*Workaround:* Select the port and assign the structure.

### 5.2.36 Improving performance in accessing Sybase IQ

**[CR #447948]** Using components such as DB Staging, DB Data Sink Insert, DB Data Sink Update, or DB Data Sink Delete to access Sybase IQ can degrade performance.

*Workaround:* To improve performance, you can do one of the following:

- Use the IQ Loader File via Load Table and IQ Loader DB via Insert Location loader components to speed up loading into Sybase IQ.
- Use Adaptive Server Enterprise, Adaptive Server Anywhere, or Microsoft Access instead of Sybase IQ for the staging portion of your project.
- If you still want to use Sybase IQ for staging, split the project into new staging and loading projects. Use DB Bulk Load Sybase IQ instead of DB Staging in the new staging project and use IQ Loader File via Load Table and IQ Loader DB via Insert Location to speed up loading into Sybase IQ in the new loading project.
- Use IQ Loader File via Load Table component for loading binary data that are extracted from Sybase IQ. You must create a custom script to do the binary extract from a Sybase IQ table, using isql utility:

- Extract Data on Source IQ system

```
set TEMPORARY OPTION
Temp_Extract_Name1='C:\myfolder\mybinfile.bin';
set TEMPORARY OPTION Temp_Extract_Binary='ON';
SELECT * FROM mytable
```

- Create a custom load script in IQ Loader File via Load Table component by setting a dummy file name in the Text Source property window.

```
LOAD TABLE mytable
{
[myCol1] BINARY WITH NULL BYTE,
[myCol2] BINARY WITH NULL BYTE,
[myCol3] BINARY WITH NULL BYTE
)
FROM '<path&filename on destination system>'
```

```
QUOTES OFF
ESCAPES OFF
FORMAT binary
```

See Chapter 7, “Moving Data In and Out of Databases” of the *Sybase IQ 12.7 System Administration Guide*.

## 5.3 Problems with internationalization

This section documents known problems with multibyte and non-Western characters.

### 5.3.1 Grid engine cannot detect IPv6

**[CR #588535]** If you have installed ETL Development 4.9 on Windows Vista and enabled IPv6, the grid engine may not detect IPv6. As a result, ETL Development may not be able to connect to the grid engine.

*Workaround:*

- 1 Edit *ETLDev.lap* file in the ETL Development installation directory.
- 2 Change the `java.net.preferIPv6Addresses` property from true to false.
- 3 Restart ETL Development.

### 5.3.2 Project execution fails if the *locale.dat* file is not configured

**[CR #563610]** On Unix and Linux, if you execute a project without providing a value representing the language, for your platform in the *\$ETLinstall/ocs/locale/locales.dat* file, project execution may fail.

*Workaround:* Add the correct language setting in the *locales.dat* file. For example, if you are using a Solaris machine, and your locale is `ja_JP.utf8`, open the *locales.dat* file and search for the `[sun_svr4]` section, then set:

```
locale = ja_JP.utf8, japanese, utf8
```

### 5.3.3 IPv6-only environment does not work with Microsoft Vista

**[CR #552407]** By default, Microsoft Vista supports IPv6. However, if you are using Vista in an IPv6-only network environment, you see this error:

```
ERROR : Cannot create socket. An address incompatible
with the requested protocol was used.
```

This is mainly due to an incompatibility between the JDK and Windows Vista. However, the same error does not occur if you are working in a mixed (dual-stack) environment of IPv4 and IPv6.

*Workaround:* None.

### 5.3.4 Issues accessing Chinese file or folder names on some UNIX platforms

**[CR #549891]** On most UNIX platforms, the LANG environment variable is set to "zh\_CN.GBK" or "zh\_CN.gbk" to support Simplified Chinese characters in the file or folder names. However, on some UNIX platforms, you may encounter errors unless you set the environment variable LANG to the locale supported on your machine.

*Workaround:* To view the list of available locales and set the appropriate one as the LANG environment variable, run the locale -a command. For example, running the locale -a | grep zh command may display the following supported locales:

- zh.GBK
- zh.UTF-8
- zh\_CN.EUC
- zh\_CN.GBK

You can set the environment variable LANG to zh.GBK or zh\_CN.GBK to support Simplified Chinese.

If you cannot see the list of locales, install the required language package on your machine.

### 5.3.5 Source files with byte-order mark are parsed incorrectly

**[CR #543716]** If you are using the Fixed by Bytes property to parse your file, make sure the source file does not include the byte-order mark. If it does, the file will be incorrectly parsed.

*Workaround:* Use a text editor to remove the byte-order mark from the source file before parsing it.

### 5.3.6 Japanese character does not convert correctly from Oracle 10g

**[CR #539726]** The Japanese character “ ~ ”, also known as UTF-16LE: 0x5EFF, does not convert correctly from the source Oracle 10g database if you use the Oracle 10g or 11g native interfaces, or the Oracle 10g or 11g ODBC interfaces in ETL Development on Windows XP.

*Workaround:* In ETL Development, use the Microsoft ODBC driver for Oracle that ships with Windows XP and set the NLS\_LANG environment variable for locale to “AMERICAN\_AMERICA.JA16SJIS” in the registry of your ETL Development machine running on Windows XP.

### 5.3.7 Insert Location component cannot load Unicode data from SQL Anywhere 10

**[CR #531902]** The IQ Loader DB via Insert Location component cannot load non-Western and multibyte character data from SQL Anywhere 10 using the nvarchar datatype.

*Workaround:* None.

### 5.3.8 Cannot enter multibyte characters in password fields

**[CR #530806]** You cannot directly type in non-Western or multibyte characters in some ETL password fields; for example, in the Database Configuration window for components. You can type only ASCII characters directly into password fields.

*Workaround:* You can paste non-Western or multibyte character passwords directly into the password field.

### 5.3.9 Insert Location component cannot load Unicode data from Microsoft SQL Server and IBM DB2

**[CR #530253]** The IQ Loader DB via Insert Location component cannot load non-Western and multibyte character data from Microsoft SQL Server and IBM DB2 to Sybase IQ using Sybase Enterprise Connect™ Data Access (ECDA).

*Workaround:* None.

## 5.4 Problems using ETL with Sybase IQ 12.7

This section documents known problems when using ETL with Sybase IQ 12.7 only.

### 5.4.1 Special characters in Sybase ETL installation paths

**[CR #454526]** Special characters in the installation path of Sybase 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 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
```

*Workaround:* 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.

### 5.4.2 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.

## 5.5 Problems with Non-ETL products

This section documents known problems with third-party products and components.

### 5.5.1 Installing SQL Anywhere 11 in Japanese or Chinese on Windows displays an error

**[CR #561733]** While installing SQL Anywhere 11 on a Japanese or Chinese Windows operating systems, you may encounter an error.

*Workaround:* Ignore the error message.

### 5.5.2 Retrieving or inserting multibyte data using ODBC drivers displays invalid characters

**[CR #550309, #550305]** On Windows and Solaris, if you are executing a project using SQL Anywhere 10 ODBC drivers to retrieve multibyte data from, or insert data to, a UTF8-encoded Sybase IQ 15.0 database, the multibyte data is retrieved from the destination database as invalid characters.

*Workaround:* On Windows:

- 1 Select Start | Settings | Control Panel | Administrative Tools | Data Sources (ODBC).
- 2 Select SQL Anywhere 10 DSN name from the list of User Data Sources or System Data Sources.
- 3 Click Configure.
- 4 Click the Advanced tab.
- 5 Enter UTF8 in the Character Set field.
- 6 Click OK until you exit the ODBC Data Source Administrator window.

On Solaris:

Add "Charset=utf8" in the ODBC configuration file.

### 5.5.3 Creating a repository connection to SQL Anywhere 10 or 11 database using ODBC may fail

**[CR #548368]** You may encounter errors while creating a repository connection to a SQL Anywhere 10 or SQL Anywhere 11 database using ODBC connectivity.

*Workaround:* To successfully create the repository connection, manually enter the Schema Name in the Schema field in the Add Repository Connection window.

### 5.5.4 Transferring large amount of data consumes excessive memory

**[CR #545145]** When you transfer large amount of data rows from Adaptive Server Enterprise to text, memory usage increases, and memory is not released until the project is executed successfully.

*Workaround:* Select "Use Cursors" when you configure the ODBC data source for Adaptive Server.

### 5.5.5 ETL may stop responding if the query has to retrieve large result sets

**[CR #540683]** Reinitializing the DB Data Provider Full Load component may cause ETL Development to intermittently stop responding, if you are running a query to retrieve records from a table containing large number or rows. This may also cause Sybase ETL Server to run out of virtual memory due to excessive memory consumption.

*Workaround:* Select “Use Cursors” when you configure the ODBC data source for Adaptive Server.

### 5.5.6 Projects on Windows platforms fail with pipe error messages

**[CR #539346]** Projects may fail with a “pipe does not exist” or “pipe permission” error when ETL Server and Sybase IQ server are on different Windows platforms.

*Workaround:*

- 1 Verify that the Sybase IQ and ETL Server hosts can both access the shared directories on each other:
  - a If Sybase IQ is running on Windows 2003, which is a multiuser operating system, ensure that you are logged in to the session where the target Sybase IQ server was started
  - b Ensure you can see the target Sybase IQ server when you try to access the ETL Server host from within the Sybase IQ host.
- 2 If pipe errors still exist after following workaround 1, specify a data file instead of a pipe name for the Load Stage option in the IQ Bulk Load Sybase IQ component to execute the project containing the component. See “DB Bulk Load Sybase IQ,” in Chapter 5 “Components” in the *Sybase ETL 4.9 Users Guide*.

### 5.5.7 Wrong date format written to DB2 on IBM AIX

**[CR #538539]** Using DB Data Provider Index Load and Text Data Sink components, the output date format to DB2 on IBM AIX is incorrect.

*Workaround:* There are two workarounds:

- 1 Set the Read Block Size to 1. Be aware that reducing the read block size decreases performance.

- 2 Use the CHAR function to change the query to convert the date or time column to the char datatype. For example, if col\_1 is a date column, col\_2 is a time column, and col\_3 is a timestamp (which does not exhibit the problem), this SQL statement corrects the problem. This statement does not force the user to change the Read Block Size and avoids performance degradation associated with larger block processing:

```
select CHAR(col_1), CHAR(col_2), col_3 from DATE_TIME_TBL
```

Text Data Sink component output is:

```
COL_1, COL_2, COL_3

1963-12-08, 12.00.00, 1991-03-02 08:30:00.000

1967-04-10, 12.01.01, 1991-04-02 08:30:00.000
```

If you do not use the CHAR function, the Text Data Sink output is:

```
COL_1, COL_2, COL_3

1963-12-08, 12:00:00, 1991-03-02 08:30:00.000

0004-10-00, 01:01:00, 1991-04-02 08:30:00.000
```

To view the data written to the DB2 table, for example the DATE4\_TBL table, enter:

```
select * from DATE4_TBL
```

The content in DATE4\_TBL in DB2 is:

```
COL_1      COL_2      COL_3
-----
12/08/1963 12:00:00 1991-03-02-08.30.00.000000

04/10/1967 12:01:01 1991-04-02-08.30.00.000000
```

### 5.5.8 No databases display if you have ODBC driver 15.0.105

**[CR #531861]** If you use Adaptive Server ODBC driver version 15.0.105 or earlier, ETL does not display any database in the Host Name field of the Properties window when you select ODBC for your database interface.

*Workaround:* Upgrade to Sybase Adaptive Server ODBC driver version 15.0.305 or later.

### 5.5.9 Specify “Use Cursors” for Adaptive Server ODBC data source

**[CR #500832]** ETL cannot display the binary, varbinary, or time value from Adaptive Server Enterprise using the ODBC interface.

*Workaround:* Select “Use Cursors” when you configure the ODBC data source for Adaptive Server.

### 5.5.10 Unable to connect to a SQL Anywhere repository created in ETL 4.2

**[CR #480747]** If you create a SQL Anywhere repository in ETL 4.2 using ODBC connectivity, when you migrate to ETL 4.9, you may encounter errors while opening it using the Sybase interface.

*Workaround:* Use ODBC connectivity to connect to the ETL 4.2 SQL Anywhere repository.

### 5.5.11 Recommended settings for IQ database options

**[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 using ETL.

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

## 6. Documentation updates and clarifications

This section contains updates and clarifications to the product documentation, online help, and demos.

## 6.1 Displaying ETL Development, product manuals, and demos

To ensure you can properly view and use the ETL Development GUI and the ETL product manuals, and run the ETL Flash demos that are installed with the product:

- Set the screen resolution to 1024x768 pixels as the Content Explorer dialog is too large to display when the screen resolution is set to 800x600 pixels.
- To run ETL Flash demos in Windows Vista, install the latest version of Adobe Flash Player from the Adobe Web site at <http://www.adobe.com>.
- To access the New Features Guide from Help | What's New in ETL Development, install the latest version of Adobe Reader from the Adobe Web site at <http://www.adobe.com>.

## 6.2 Installation Guide

This section describes updates and clarifications to the *Sybase ETL 4.9 Installation Guide*.

### 6.2.1 Supported Oracle ODBC interface driver version is incorrect

The supported Oracle ODBC interface driver version is Oracle ODBC 11.01.00.06 (Windows only) and not Oracle ODBC 11.01.01.06 (Windows only), as specified in “Interfaces” in Chapter 1, “Introduction” of the *ETL 4.9 Installation Guide*.

## 6.3 Users Guide

This section describes updates and clarifications to the *Sybase ETL 4.9 Users Guide*.

### 6.3.1 Chapter 4, “Advanced Concepts and Tools”

The following information is not documented in “Managing jobs and scheduled tasks” in Chapter 4, “Advanced Concepts and Tools.”

### Last result column of the Runtime Manager

**[CR #584973]** The Last Result column in the Runtime Manager displays a value of either Failed or Succeeded for any scheduled task that has already run. The same information appears in the Execution Result column of the Schedule Task List when viewed in the Web Monitor.

- Succeeded – means either the task completed as expected or was prematurely terminated in one of these ways:
  - By clicking the Terminate a Running Schedule icon on the toolbar of the Runtime Manager, or by selecting Actions | Terminate.
  - By clicking Cancel Execution on the Execution Monitor.
  - By clicking Cancel in the Web Monitor.

The value of Succeeded in the Last Result column does not always mean that the task ran to completion without errors.

- Failed – means that the task was unexpectedly terminated, for example, by killing the grid engine process using the Windows Task Manager or through operating system command line.

---

**Note** Sybase recommends that you do not use the Windows Task Manager or operating system command line to terminate scheduled tasks.

---

### Configuring *.odbc.ini* file

**[CR #589842]** If you are using ETL Scheduler to set up a repeated task using a server that is running on UNIX, you must configure the *.odbc.ini* file on the UNIX machine to include a pointer to the repository data server.

#### ❖ Configuring *.odbc.ini* file

- 1 Use a text editor to open the *etc/.odbc.ini* file.
- 2 Add an entry similar to:

```
[repository_data_server]
uid=dba
pwd=sql
EngineName=demo
CommLinks=tcpip(host=<hostname>;port=<portnumber>)
AutoStop=no
DatabaseName=demo
```

where *repository\_data\_server* is the name of the SQL Anywhere data server.

---

**Note** Make sure that the entry does not include the “AutoPreCommit” connection parameter.

---

- 3 Save the file.
- 4 In the ETL Scheduler, set up the repeated task, then run the job.

### 6.3.2 Chapter 5, “Components”

**[CR #578575]** The Continue on Error property, which is documented in Chapter 5, “Components,” is no longer available for any of the job components.

### 6.3.3 Appendix D, “Best Practices”

#### Simulation does not roll back data when error happens in post-SQL

**[CR #584682]** The following information is not documented in “Best practices for working with ETL Server” in Appendix D, “Best Practices.”

Do not select the “Execute post-processing as for successful execution” option, in case a post-SQL error occurs during execution. If you do so, the project commits all transactional components, regardless of the error.

## 7. 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.

## 8. 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.

- 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/>.

## 8.1 Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

### ❖ Finding the latest information on product certifications

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click Certification Report.
- 3 In the Certification Report filter select a product, platform, and timeframe and then click Go.
- 4 Click a Certification Report title to display the report.

### ❖ Finding the latest information on component certifications

- 1 Point your Web browser to Availability and Certification Reports at <http://certification.sybase.com/>.

- 2 Either select the product family and product under Search by Base Product; or select the platform and product under Search by Platform.
- 3 Select Search to display the availability and certification report for the selection.

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

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

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

## 8.2 Sybase EBFs and software maintenance

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

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

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

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