SYBASE®

Sybase® Adaptive Server™ Enterprise Troubleshooting and Error Messages Guide Volume 2

Sybase Adaptive Server Enterprise

11.0.x–12.0

DOCUMENT ID: 39995-01-1200-01

LAST REVISED: June 2000

Sybase, Inc., 6475 Christie Avenue, Emeryville, CA 94608.

Contents

About This Book

The *Sybase Adaptive Server Enterprise Troubleshooting and Error Messages Guide* is a three-volume set. The first volume contains guidelines for system database recovery and other troubleshooting procedures. This second volume contains detailed error message descriptions. The third volume includes the text of all Component Integration Services, Backup Server, and Adaptive Server error messages.

For information about the intended audience of this guide, related documents, other sources of information, conventions used in this manual, and how to obtain help, please refer to “About This Book” in Volume 1.

CHAPTER 1

Error Message Writeups

This chapter contains explanations of and instructions for resolving some of the most frequently occurring Adaptive Server error messages for version 11.0.x through 12.0. Most of the messages presented here contain error numbers (from the master..sysmessages table), but some error messages do not have error numbers, and which occur only in the Adaptive Server error log.

Error Message Severity Levels

A severity level is associated with each numbered error message. It indicates the type of problem Adaptive Server has encountered. For more information about severity levels, refer to "Severity Levels" in the *System Administration Guide*.

Inform the Sybase System Administrator whenever errors with severity levels over 16 occur. The System Administrator is responsible for resolving them and tracking their frequency. The System Administrator should monitor all errors that generate severity levels of 17 through 26. To assist the System Administrator in resolving problems, print out a hard copy of the error log, which contains the back trace from the fatal error.

Note Messages that ordinarily have severities greater than 16 will show severity 16 when they are raised by dbcc checktable or dbcc checkalloc so that checks will continue with the next object.

Variables in Error Message Text

The following table lists the symbolic constants that appear in the error message text provided with each error message explanation:

Table 1-1: Error message text variables and meanings

Symbol	Stands For
%d,%D	Decimal number
%x, %X, %.*x, %lx, %04x, %08lx	Hexadecimal number
%s	Null-terminated string
%.*s, %*s, %*.s	String, usually the name of a particular database object
%S_type	Sybase-defined structure
%c	Single character
%f	Floating-point number
%ld	Long decimal number
%lf	Double floating-point number

Finding Object Names from Error Message Text

Many Adaptive Server error messages print only the logical page number, not the name of the table or index affected. Chapter 2, “How to Find an Object Name from a Page Number” supplies a method to find the table or index name using the logical page number supplied by the error message. You may want to review that section now so that you are familiar with the necessary procedures before an error occurs.

How to Determine Your Adaptive Server Version

The dataserver, diagserver, backupserver, isql, bcp, buildmaster, defncopy, and langinstall binaries return version information when executed with a -v or /version flag.

Note diagserver is a diagnostic version of dataserver. Sybase Technical Support may ask you to run diagserver to help in problem solving.

The following examples are from UNIX:

```
% /usr/sybase/bin/dataserver -v
Adaptive Server Enterprise/11.5/P/Sun_svr4
/OS 5.5.1/1/OPT/Wed Sep 10 00:01:37 PDT 1997
% /usr/sybase11/bin/dataserver -v
SQL Server/11.0/P/Sun_svr4/OS 5.2/1/OPT
```

```
/Thu Sep 23 12:28:52 PDT 1996
% /usr/sybase11/bin/backupserver -v
Backup Server/11.0.2.1/P/Sun4/4.1.x
/EBF 4009/OPT/Tue Nov 15 14:27:37 PST 1996
% /usr/sybase11/bin/isql -v
isql/11.0/P/sun_svr4/Solaris 2.2/1
/Mon Jun 30 12:11:43 PDT 1997
```

The following example is from OpenVMS:

```
$ server :== $sybase_system:[sybase.bin]dataserver.exe
$ server/version
SQL Server/11.0.1/P/VMS/5.4-1A+/1/OPT
/28-MAY-1997 18:37:38.53
```

The above examples are pre-12.0. In version 12.0, dataserver resides in *$SYBASE/ASE-12_0/bin*.

Field Order in Version String

The order in which the fields appear in the version string depends on the version of the server you are running.

In Version 11.5.x and earlier, the fields, in order, are:

1 Product.

2 Version number.

3 Release type: production (P), beta, or SWR version.

4 Platform identifier.

5 Operating system release under which the binary was compiled. This number is hard coded; it is *not* determined from your operating system.

6 “Build number” – this is a Sybase internal reference.

7 Mode: optimized (dataserver) or debug (diagserver).

8 Compilation date and time.

In Version 11.9.2, the version string appears as follows (this example is from Sun Solaris):

```
Adaptive Server Enterprise/11.9.2/1031/P/Sun_svr4/OS 5.5.1/FBO
/Fri Aug 14 06:26:45 1998
```

The fields, in order, are:

1 Product.

2 Version number.

3 “Build number” - this is a Sybase internal reference.

4 Release type: production (P), beta (B) or SWR version.

5 Platform identifier.

6 Operating system release under which the binary was compiled. This number is hard coded; it is *not* determined from your operating system.

7 Type of post-build optimization server. In the example, FBO stands for Feedback Optimized.

8 Compilation date and time.

An example of the 12.0 version string is as follows:

```
Adaptive Server Enterprise/12.0/P/SWR 8774 ESD 1/Sun_svr4/OS
5.6/1580/32bit/FBO/Tue Dec  7 03:10:20 1999
```

Creating Error Messages

You can create your own error numbers and messages based on Sybase error numbers. For more information, refer to “sp_addmessage” in the *Adaptive Server Reference Manual*.

Reporting Errors

When you report an error to Sybase Technical Support, include the following information:

- The environment in which the error occurred, including:
  - Adaptive Server version (refer to “How to Determine Your Adaptive Server Version”)
  - EBF (Emergency Bug Fix) number, if available
  - Client version
  - Operating system version
  - Hardware platform

  - Production or development environment
- The message number, state, and complete error message text for any Adaptive Server errors. Include any numbers, database object types, or database object names that are included in the error message.
- The context in which the message was generated—what command was running at the time. You can help by providing a hard copy of the error log, including all messages displayed from just prior to start-up to immediately before or immediately after the message. Include the version number and date information printed at the beginning of the error log.
- The number of users who were logged into Adaptive Server when the error occurred.
- The frequency with which the error occurs.
- The client program that encountered the error (isql, Open Client™ DB-Library™, Data Workbench®, and so on).

A checklist for reporting errors follows. The purpose of this checklist is to help you collect the information you need to give to Sybase Technical Support. Your primary site contact should have a copy of this checklist.

Technical Support Checklist

To expedite the resolution of your question and help you track it, please fill out this checklist before you call your Technical Support Center.

Site ID (customer number):

Case Number:

Priority:

Support Engineer:

Date Case Opened:

Date Case Closed:

Sybase Server version (include SWR [EBF] version if applicable):

Operating system type and version:

Front-end (client) product and version (include SWR [EBF] version if applicable):

Client machine operating system type and version:

Description of the problem (include error numbers, error messages, circumstances in which problem occurred):

Other useful information (include copy of error logs, reproducible case, number of users on the system, average percent of CPU usage, any other pertinent information):

Faxing Error Log Fragments

Whenever Technical Support asks you to fax the error log to Sybase, fax the log from start-up to the error message in question. Be sure to include the following:

- The lines just prior to start-up; for example:

```
kernel  Adaptive Server Enterprise/11.5/P/Sun_svr4/OS 5.5.1/1
/OPT/Wed Sep 10 00:01:37 PDT 1997
```

- All the lines between start-up and the error; for example:

```
Error 2559, Severity Level 16: Data page number 0781 is empty.
```

- Any stack traces or unusual messages that occur near the error message in question.

When trace flag 3601 is turned on, Adaptive Server writes a stack trace to the error log every time an Adaptive Server error is raised. For troubleshooting purposes, Sybase Technical Support may ask you to turn on this trace flag. If they do, refer to Chapter 2, "How to Start Adaptive Server with Trace Flags" for instructions for turning on a trace flag.

Reproducing Problems

In general, the time that Sybase Technical Support needs to resolve a problem is greatly reduced if you generate a reproducible test case demonstrating the error. These test cases can range from a single query that is not behaving as expected to a relatively complex code fragment that encounters a problem.

Whenever possible, test cases should be created using the pubs or pubs2 database. In many cases, this is a task that saves you from having to send large amounts of data to Technical Support.

Error Message Writeups

The rest of this chapter contains error message writeups, in order by message number.

Parser Errors

This section includes error messages for the Adaptive Server Parser.

Error 102

Severity 15

Message text `Incorrect syntax near '%.*s'.`

Explanation

This error occurs when Adaptive Server detects a syntax error in a Transact-SQL command or query.

This error can occur when:

- A keyword is spelled incorrectly.

- A keyword or parameter is missing.
- The order of keywords is incorrect.
- You use an editor invoked from isql to write a SQL command or batch, end your SQL with “go” or another terminator, and then run the file from isql. This only causes this error in some cases.
- A script that used to work no longer works after you change the sort order or the character set of an Adaptive Server (for example, a script that was created on a server that originally was not case-sensitive and now that Adaptive Server's sort order has been changed to case-sensitive).
- A reserved word has been used in a query (for example select user = suser_name() raises Error 102 because “user” is a reserved word).
- You have used a variable instead of a specific period of time following a waitfor delay statement. Variables are not supported in this case. For example, the following would raise Error 102:

  ```
  1> declare @t char(9)
  2> select @t = "00:00:01"
  3> waitfor delay @t
  ```

- You have used a variable for a database name and your query contains a parser error; Error 102 is raised because the batch is never executed. Steps for executing a query are:

  a Read the batch.

  b Parse the batch.

  c Optimize the batch.

  d Execute the batch.

  Any step that fails ends the processing of the query.

Action

Check the spelling and syntax of the command specified in the error message. If it is wrong, correct it in your Transact-SQL statement and run it again.

You can invoke an editor such as vi from isql to edit a SQL statement or statements. However, for the statements to execute, you have to enter the command termination string (usually “go”) after you return to isql. If you include the command termination string in the editor, then save the file and return to isql, the statement will not execute. If you enter the command termination string again after returning to isql, you get the following error (“vi” is the editor in this example; you may be using a different editor):

```
1> vi
```

```
1> select * from sysmessages where error = 102
2> go
3> go
Msg 102, Level 15, State 1:
Server 'REL1002_NAME', Line 2:
Incorrect syntax near 'go'.
```

If you are not sure that your query contains a reserved word, use the following query to see a complete list of reserved words:

```
1> select name from spt_values where type = "W"
2> go
```

Additional information

Refer to the documentation for the utility you are using for correct syntax information. Some examples are:

- *ASE Reference Manual*
- *Transact-SQL User's Guide*
- ASE utility programs manual

Version in which this error is raised

All versions

Error 107

Severity

15

Message text

```
The column prefix '%.*s' does not match with a table
name or alias name used in the query. Either the table
is not specified in the FROM clause or it has a
correlation name which must be used instead.
```

Explanation

Tables are specified in the FROM clause of a query. When Adaptive Server parses a query prior to execution, the name by which it knows the table is in the from clause: the table name if given alone, or a correlation name (alias) if one was specified. For example, in the query:

```
1> select c.cityname from cities c
2> go
```

the table is known as c for purposes of the query. Names specified in other clauses refer back to this name to determine what table is intended.

Error 107 can be raised:

- If a column prefix intended as a table name does not match any table name specified in the from clause.
- If a query does not use correlation names consistently. In Adaptive Server, queries that include correlation names must conform to ANSI requirements. Statements that specify correlation names but do not use them consistently return Error 107.

These restrictions apply to views as well as real database tables.

Action

Check your query for these errors and take corrective action:

1 Check for a typing error like the following:

```
1> select * from titles
2> where title.code = 205
3> go
Msg 107, Level 15, State 1:
Server 'mfg', Line 2:
The column prefix 'title' does not match with a table
name or alias name used
```

The column name in the where clause should be titles.code. Likewise this statement is incorrect:

```
1> select t2.title_id from titles t1
```

The correct statement is:

```
1> select t1.title_id from titles t1
```

2 Ensure that correlation names are used correctly. For example, this statement is incorrect:

```
1> select title_id
2> from titles t
3> where titles.type =  "french_cook"
```

The where clause can not use titles, because the from clause defines a correlation name for the table. The correct query is:

```
1> select title_id
2> from titles t
3> where t.type = "french_cook"
```

Additional information

Note the special case where a query that returns Error 107 may report no error when the same type of correlation is used in a subquery. For example:

```
1> select * from mytable
2> where columnA =
3> (select min(columnB) from mytable m
4> where mytable.columnC = 10)
```

This query is a correlated subquery, and mytable.columnC refers to the outer table mytable. This query works because the same table is referred to in the inner and outer queries. In general, however, correlated subqueries can also generate Error 107 when correlation names are used incorrectly.

Version in which this error is raised

All versions

Sequencer Errors

This section includes error messages for the Adaptive Server Sequencer. The Sequencer is the subsystem that controls command execution. The Sequencer handles the processing between steps in commands and also handles the procedural steps between commands such as if, goto, and while.

Error 207

Severity

16

Message text

```
Invalid column name '%.*s'.
```

Explanation

This error occurs when Adaptive Server tries to use a column name that does not exist.

Adaptive Server supports delimited identifiers for table, view, and column names. Delimited identifiers are object names enclosed within double quotation marks. Using them allows you to avoid certain restrictions on object names. Delimited identifiers can begin with nonalphabetic characters, including characters that would not otherwise be allowed. Delimited identifiers may even be Transact-SQL reserved words.

Delimited identifiers may cause some SQL statements to fail. For example, Error 207 would be raised for the column *"none"* in the following example:

```
1> select au_name, "none", au_fname
2> from pubs2..authors
3> go
```

Action

Check the spelling of the column name specified in the error message. If it is wrong, correct it in your SQL statement and run it again.

If this error message is being displayed because of delimited identifiers, either use single quotes instead of double quotes or turn delimited identifiers off. To turn delimited identifiers off, do the following:

```
1> use database_name
2> go

1> set quoted_identifier off
2> go
```

When delimited identifiers is turned on, delimiting strings with double quotes causes Adaptive Server to treat them as identifiers.

Additional information

Refer to the *Transact-SQL User's Guide* for information about delimited identifiers.

Version in which this error is raised

All versions

Error 208

Severity

16

Message text

```
%.*s not found. Specify owner.objectname  or use sp_help
to check whether the object exists (sp_help may produce
lots  of output).
```

Explanation

This error occurs when Adaptive Server tries to use an object name that does not exist. If the object does exist, you may need to include the owner's name in the object name.

If the object is not owned by the user who is attempting to access it, and it is not owned by the database owner ("dbo"), all references to it must include the owner name. For example, if "user1" creates a table called test, all other users must prefix the table name with the owner name "user1" every time they refer to the object. The complete name in this example is user1.test. This requirement is based on Adaptive Server's naming convention for database objects:

```
database.owner.object_name.column_name
```

The default value for *database* is the current database and the default value for *owner* is the current user. Remember that the owner is part of the object name and it is therefore possible for two different users to have two different tables with the same *object_name* in the same database, for example, user1.test and user2.test. Refer to the *Transact-SQL User's Guide* for more details on naming conventions.

Temporary tables reside in tempdb and are automatically dropped when the user process or Adaptive Server connection that created them is exited. Furthermore, users cannot share temporary tables, even if they are created by the "dbo."

Including the fully qualified object name in an application may complicate the maintenance of the application. For example, if all references to a table include the database name, changing the database name could become quite difficult.

Action

To resolve this error, refer to one of the following sections, depending on how much you know about the object in question.

If You Do Not Know Who Owns the Table

You can either use the sp_help procedure to display the owner (if the procedure is executed with no parameters), or query the system catalog to determine the object's owner and type. For example, to determine the name and type of the object table1, use the following query:

```
1> select owner = user_name(uid), name, type
2> from sysobjects where name = "table1"
3> go
```

If no rows are returned by this query, the object may reside in a different database or may not exist at all.

If You Do Not Own the Object In Question

Avoid this error using either procedure:

- Include the owner's name in the object name. For example:

  ```
  1> select * from user1.table1
  2> go
  ```

- Have the "dbo" create the object. This allows any user to find the object without specifying an owner name.

If the Object Does Not Reside in the Database

The error can be avoided by:

- Moving to the correct database with the use *database_name* command.
- Fully qualifying the object name with the database name. For example:

  ```
  1> select * from database1.user1.table1
  2> go
  ```

The owner name is not needed if you own the object or if it is owned by the "dbo." For example:

```
1> select * from database1..table1
2> go
```

The appropriate permissions must also be set to allow access to this object. If these permissions are not provided, a 229 or 230 error results.

Version in which this error is raised

All versions

Error 213

Severity

16

Message text

```
Insert error: column name or number of supplied  values
does not match table definition.
```

Explanation

This error occurs when there is a mismatch between the information specified in an insert statement and the definition of the table.

Following are some examples of insert statements for the pubs2.titles table that cause this error (assuming that the titles table has nine columns):

- No column names are specified after the table name in an insert statement and there are two or more columns in the table.

```
1> insert titles values ('this is a test')
2> go
Msg 213, Level 16, State 4:
Line 1:
Insert error: column name or number of supplied
values
does not match table definition.
```

- No values are specified in an insert statement.

```
1> insert titles (title_id, title) values ()
2> go
Msg 213, Level 16, State 3:
Line 1:
Insert error: column name or number of supplied
values
does not match table definition.
```

Action

Check your insert statement to make sure you have specified column names and values for the columns.

If this error message is being displayed while an application is being run, use the following information to determine which insert statement is causing the problem.

Data Workbench

To see the SQL text being sent by Data Workbench to Adaptive Server, use the environment variable RECTFOS. When RECTFOS is defined, the application records the SQL text in a file that you can review for debugging a new application or diagnosing other problems.

Before starting Data Workbench, set the value of RECFTOS to the file name where you want the text to be saved:

- In UNIX, type:

```
% setenv RECFTOS filename
```

- In OpenVMS type:

```
$ define recftos filename
```

A new file is created for each connection made to Adaptive Server. The files are named sequentially as each connection is opened: *filename*.0;1 (OpenVMS), *filename*.1 (UNIX), and so on.

This feature is present in Data Workbench 2.2 and later releases.

Use unsetenv (UNIX) or deassign (OpenVMS) to turn off logging when you no longer need to save the SQL text being sent.

Open Client Applications

To see the SQL text being sent by Open Client applications, use the dbrecftos call. Refer to the *Open Client DB-Library Reference Manual* for details.

Version in which this error is raised

All versions

Error 216

Severity

```
20
```

Message text

```
Attempt to automatically drop temporary table failed.
```

Explanation

Adaptive Server goes through four phases while processing a query:

1 Parsing – checks the query to make sure the syntax is correct

2 Normalization – resolves the column and table names and generates a query tree

3 Compilation – performs optimization and generates a query plan

4 Execution – carries out the query plan

During the first three phases of query processing, Adaptive Server creates two types of temporary tables:

- Temporary tables to aid in table resolution. This is necessary if the table will not exist until the command has been executed. For example:

```
1> create table x [information about table x]
2> select * from x
3> go
```

  In this case, Adaptive Server creates a temporary copy of table *x* so it can resolve the select command.

- Other temporary tables prefaced with "#", that exist for the life of the session (or the life of the procedure, if they are created in the procedure).

Error 216 occurs when Adaptive Server is unable to drop one of the temporary tables it created. It is probably caused by an Adaptive Server problem.

Error 216 occurs with the following states:

State	Meaning
1	During the renormalization phase, if Adaptive Server cannot drop a temporary table it created, Error 216 occurs with State 1. These are definition-time temporary tables.
2	During the normalization phase, if Adaptive Server cannot drop a temporary table it created during the parsing phase, Error 216 occurs with State 2.
3	During an abort of the normalization phase, if Adaptive Server cannot drop a temporary table, Error 216 occurs with State 3.
4	When a process is killed, if Adaptive Server fails to clean up temporary tables, Error 216 occurs with State 4. These can be definition-time or normal temporary tables.

Action

Temporary tables left behind are not in themselves a problem. However, they do take up space in tempdb, and tempdb may eventually fill up. If you think this is a potential problem, shut down and restart Adaptive Server. This will re-create tempdb and thus drop all temporary tables.

If Error 216 occurs again, call Sybase Technical Support.

Version in which this error is raised

All versions

Error 225

Severity

11

Message text

```
Cannot run query--referenced object (name %.*s) dropped
during query optimization.
```

Explanation

This error occurs when Adaptive Server fails to build an execution plan for a stored procedure or an ad-hoc query because a referenced object was not found.

Some reasons for this error include:

- The value of the number of open objects configuration parameter is too low.
- A problem exists in Adaptive Server.

This is not a serious error as it does *not* indicate corruption. However, it may prevent you from running some stored procedures and ad-hoc queries.

This error usually means a problem has occurred with Adaptive Server.

Action

Using sp_configure, review the value of the configuration parameter number of open objects and increase this value if feasible. For information about memory-related configuration parameters, refer to:

- "Setting Configuration Parameters" in the *System Administration Guide*
- *Performance and Tuning Guide*

If you cannot increase the value of number of open objects or the 225 error occurs again, call Sybase Technical Support for assistance.

Version in which this error is raised

All versions

Error 226

Severity

16

Message text

```
%s command not allowed within  multi-statement
transaction.
```

Explanation

SQL commands are grouped into the following categories:

- SQL commands that are not allowed in transactions at all.
- SQL commands, such as Data Definition Language (DDL) commands, that are allowed in transactions only if the required database option (ddl in tran) is set to TRUE.

- SQL commands that are allowed only if the transaction affects some other database. These commands include create table, drop table, and other commands that are run across databases to create or drop objects in another database when the database in which the objects are being created or dropped has the database option ddl in tran set to TRUE.

Error 226 occurs when Adaptive Server detects a command that is not allowed in a multi-statement transaction. A multi-statement transaction is a set of commands prefaced with the begin transaction command.

The following commands are never allowed in multi-statement transactions:

- alter database
- create database
- dbcc reindex, dbcc fix_text
- disk init
- drop database
- dump database, dump transaction
- load database, load transaction
- select into
- set transaction isolation level
- truncate table
- update statistics

The following commands are not normally allowed in multi-statement transactions but you can use them if you use sp_dboption to set ddl in tran to TRUE first:

- create default, create index, create procedure, create rule, create schema, create table, create trigger, create view
- drop default, drop index, drop procedure, drop rule, drop table, drop trigger, drop view
- grant
- revoke

Action

If the command is allowed in a multistatement transaction when ddl in tran is set to TRUE, set ddl in tran to TRUE before running the transaction. Setting ddl in tran to TRUE causes locks on system tables and this can affect performance. You can check the current setting of ddl in tran with sp_helpdb.

If the command is never allowed in a multi-statement transaction, execute it outside the multi-statement transaction.

Additional information

Some applications take SQL statements as input and run them for you. If the application uses begin and commit or rollback transaction to surround those statements, Error 226 may occur. Refer to the user guide for your application to determine if this is the case.

Version in which this error is raised

All versions

Error 229

Severity

14

Message text

```
%s permission denied on object %S_OBJID, database
%S_DBID, owner %.*s
```

Explanation

This error occurs when Adaptive Server attempts to access an object for which you do not have the appropriate permission.

Error 229 can occur during any transaction which involves objects, such as bcp, select, or update.

This error causes a rollback of the current transaction. Remaining commands in the batch are not executed and processing continues at the next batch.

Action

Check permission on the object named in the error message:

```
1> use database_name
2> go
```

```
1> sp_helprotect object_name
2> go
```

where *database_name* is the name of the database in the error message and *object_name* is the name of the object in the error message.

Grant the needed permission to the object (you either need to own the object or be the “sa” to do this). For example, to give Mary permission to insert into and delete from the titles table, type the following:

```
1> grant insert, delete
2> on titles
3> to mary
4> go
```

Additional information	Refer to “Managing User Permissions” in the *Security Administration Guide* for detailed information about permissions.
Version in which this error is raised	All versions

Error 232

Severity	16
Message text	`Arithmetic overflow error for type %s, value = %f.`
Explanation	This error occurs when Adaptive Server detects an arithmetic overflow error during execution of a SQL statement. Arithmetic overflows occur when a variable has too few places to store the converted data. In particular, Error 232 is raised when an overflow occurs for a variable of datatype float. Error 232 can occur during execution of the convert function. It can also occur during implicit conversion within Adaptive Server, either during a stored procedure run or during execution of a SQL statement which converts a float value to some other datatype.
Action	Determine which variable is causing the problem based on the datatype reported in the error message output and the SQL statement on which the error occurred. If the error occurs within a stored procedure, an additional error message will be displayed showing the name of the stored procedure and the line number where the error occurred. Increase the size of the variable that is causing the error. If it is in a procedure, re-create the procedure. If this error occurs on a Sybase-supplied system stored procedure, contact Sybase Technical Support.
Version in which this error is raised	All versions

Error 233

Severity	16
Message text	`The column %.*s in table %.*s does not allow null values.%S_EED`

Explanation

During query processing, Adaptive Server performs an action called normalization, during which it resolves column and table names and generates a query tree. During normalization, when Adaptive Server prepares the target list for processing an insert or update, it checks to make sure nulls are allowed. Error 233 will be raised at compile time if, from within a stored procedure, you try to insert null values in a column which does not allow nulls.

Error 233 is raised with the following states:

State	Meaning
1	Nulls are not allowed, the column is in the target list, and the value being inserted is null.
2	No default exists, nulls are not allowed, and the column is not in the target list.
3	For inserts to temporary tables, Adaptive Server checks to be sure all "not null" columns are listed explicitly. If a "not null" column is not listed in the insert statement, the error is raised with State 3.

Action

When you create a table, you can explicitly define whether each column should allow null values. If you do not specify NULL or NOT NULL for a column when you create the table, the default value will be NOT NULL. If you use sp_dboption to set allow nulls by default to TRUE for the database, any new table that is created will have columns with the default value NULL.

To determine whether a column has NULL or NOT NULL defined, enter commands like the following, replacing "pubs2" by the database in which your table resides and "titleauthor" by the table in which your column resides:

```
1> use pubs2
2> go
1> sp_help titleauthor
2> go
Name          Owner   Type
-------------------------------------------------------------
titleauthor   dbo     user table
Data_located_on_segment        When_created
------------------------------ ----------------------------
default                        Oct 27 1994 10:09AM
Column_name Type Length Prec Scale Nulls Default_name Rule_name Identity
-----------------------------------------------------------------------
au_id       id       11 NULL NULL  0     NULL         NULL      0
title_id    tid       6 NULL NULL  0     NULL         NULL      0
au_ord      tinyint   1 NULL NULL  1     NULL         NULL      0
royaltyper  int       4 NULL NULL  1     NULL         NULL      0
```

The Nulls column indicates whether null values are allowed. A value of 0 for the column means nulls are not allowed; a value of 1 means null values are allowed.

To change the default for a column from NOT NULL to NULL:

1 Use bcp to copy the data out of the existing table.

2 If you want the modified table to have the same name as the existing table, drop the old table.

3 Re-create the original table, specifying NULL for the column you want to change.

4 Use bcp to put back the data for the table.

If you want *new* tables in the database to have the default NULL, use the following commands for the database:

```
1> use master
2> go
1> sp_dboption database_name,
2> "allow nulls by default", true
3> go

1> use database_name
2> go

1> checkpoint
2> go
```

where *database_name* is the name of the database whose behavior you want to change.

Version in which this error is raised

All versions

Error 247

Severity

16

Message text

```
Arithmetic overflow during %S_MSG conversion of %s value
'%s' to a %s field.
```

Explanation

Error 247 occurs when Adaptive Server attempts to convert a given value from one data type to another, but is unable to do so because of incompatibilities between the datatypes, or when the receiving field (the "to" field in the message) does not have a range large enough to accomodate the converted value.

The error may be raised during either explicit or implicit conversions. Explicit conversions occur when a query explicitly requests a conversion with the convert, inttohex, or hextoint functions.

Implicit conversions occur when Adaptive Server is required to perform certain types of comparisons between heterogeneous datatypes, and when the server stores and retrieves data. Adaptive Server automatically handles many such conversions from one datatype to another, but may be unable to handle the conversion due to the reasons mentioned above.

For example:

```
1> create table deliver_dates
2> (itemno int,
3> dlydate smalldatetime)
4> go

1> insert into deliver_dates
2> values (42298, '12/12/2080')
3> go
Line 1: Arithmetic overflow during implicit conversion
of
VARCHAR value '12/12/2080' to a SMALLDATETIME field.
```

The date value being inserted is outside the range of SMALLDATETIME.

Action

Check the following problem areas when you see Error 247:

1 Is the range of the receiving datatype large enough to accomodate the converted value?

2 Are the sending and receiving datatypes (the "of" and "to" datatypes in the message, respectively) compatible? For details refer to "Datatype Conversion Functions" in the *Reference Manual*.

3 Is the query attempting to display a FLOAT value by converting it to numeric datatype? FLOATs can only be displayed this way if no decimal digits are lost. Try increasing the precision of the numeric field so that this condition is met.

4 If the query is updating a table, is there a trigger on the table that attempts to place the new value into another table? The receiving field in the second table may not have the correct range for the operation.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level
- Text of all error messages
- Text of the query which produces the error.

Version in which this error is raised

All versions

Error 257

Severity

16

Message text

```
Implicit conversion from datatype '%s' to '%s' is not
allowed.  Use the CONVERT function to run this query.
```

Explanation

When performing certain types of comparisons, and while storing and retrieving data, Adaptive Server automatically handles many conversions from one datatype to another. These are called implicit conversions. You can explicitly request other conversions with the convert, inttohex, and hextoint functions.

Error 257 is raised when Adaptive Server is unable to do an implicit conversion because of incompatibilities between the datatypes.

Action

Corrective action depends on the exact situation in which the error occurred. Here are some typical situations where Error 257 is raised, and how you can correct the problem:

- If you attempt to perform comparisons on integer data with the like keyword; you must use the convert function on integer data if you want Adaptive Server to treat it as character data for comparisons.
- If you are using embedded SQL, and Adaptive Server is unable to perform an implicit conversion between a database column and the corresponding host language variable. Consider using a different host variable to accomplish the conversion. For example, in a COBOL application a column of money datatype requires a COMP-2 host variable.

- If you attempt to insert quoted values for integer data. For example the following will raise Error 257:

```
1> create table citycodes
2> (cityname char(12), citycode smallint)
3> go
1> insert into citycodes
2> values ('Detroit', '123')
3> go
```

  Quotes are not allowed around the city code in this insert statement.

- When a statement implicitly exceeds the maximum length of a character datatype. For example, suppose *columnX* is defined as varchar(255). If you submit a query like

```
1> select * from tableX
2> where columnX like '.....%'
```

and there are 255 characters preceding `%', Error 257 is raised. Adaptive Server treats the comparison string as containing more than 255 characters and assumes that it is a TEXT datatype, which it cannot convert implicitly.

Additional information

For detailed information about datatype conversions, refer to the section "Datatype Conversion Functions" of "Using the Built-In Functions in Queries" in the *Reference Manual*.

Before calling Technical Support, have the following information available:

- The SQL statement generating the error
- sp_help for any tables involved
- Server version and SWR Rollup level
- Server error log
- Text of all error messages

Version in which this error is raised

All versions

Error 259

Severity

16

Message text

```
Ad-hoc updates to system catalogs not enabled. A user
with System Security Officer (SSO) role must reconfigure
```

```
system to  allow this.
```

Explanation

System tables are all the Sybase-supplied tables in the master database and in each user database that begin with “sys” and have an ID value in the sysobjects table of less than or equal to 100. The system catalog is the collection of system tables.

Error 259 occurs when you try to directly modify an Adaptive Server system table without first using sp_configure to set the allow updates configuration parameter. Setting this variable allows the System Administrator (“sa”) to make changes to system tables.

Warning! Serious problems can result from an incorrectly modified system table! Make changes only if Sybase Technical Support or this manual instructs you to do so.

Action

Do not use any of the procedures in this section before reading all warnings and cautions! Incorrect alterations of a system table can result in database corruption and data loss.

To successfully modify the system catalog as directed by Technical Support or this guide, use the following procedure:

1 As “sa,” use sp_configure to allow changes to the system catalog:

```
1> sp_configure "allow updates", 1
2> go
```

2 Make changes to the system catalog as directed by Technical Support or elsewhere in this guide. Always use begin transaction when changing a system table, to protect against errors which could corrupt your database.

Warning! Be sure to include the next step to protect your system.

3 Immediately following the changes, use sp_configure to disallow changes to the system catalog (the normal state for Adaptive Server):

```
1> sp_configure "allow updates", 0
2> go
```

You must return Adaptive Server to its normal state, or anyone who can log on as “sa” can make changes to the system tables, or create stored procedures which do so.

4 If you update sysusages, sysobjects, sydatabases, or sysdevices, restart Adaptive Server. You need to do this because structures related to those tables are stored in cache and will not be updated until you restart Adaptive Server.

Additional information

Refer to "allow updates to system tables" configuration parameter in the *System Administration Guide* and "sp_configure" in the *ASE Reference Manual* for more information about (sp_configure) allow updates.

Version in which this error is raised

All versions

Error 265

Severity

16

Message text

```
Insufficient result space for %S_MSG conversion  of %s
value '%s' to a %s field.
```

Explanation

Error 265 can be raised in the following situations:

- Conversions to character data fails with Error 265 if you attempt a conversion that would cause truncation of data.
- Error 265 can also be raised in some conditions when certain system stored procedures are executed.

Examples and corrective action are provided in the next section.

Action

The following sections describe some scenarios under which Error 265 occurs and ways of dealing with those situations.

Errors in Integer Conversion

Conversions of integer to character data will fail if the target format is not large enough to accommodate the data:

```
1> select convert (char(1), 500)
2> go
Msg 265, Level 16, State 1:
Server 'mfg1', Line 1:
Insufficient result space for explicit conversion of
INT value '500' to a CHAR field.
```

To correct this problem, choose a larger target format (char(3) in the example).

Errors in Floating Point Conversion

Conversions of floating point to character data will fail if the target format is not large enough to accommodate the data:

```
1> select convert (char(10), 3.1415e)
2> go
Msg 265, Level 16, State 1:
Server 'mfg1', Line 1:
Insufficient result space for explicit conversion of
FLOAT value '3.14150000000000002' to a CHAR field.
```

As in this example, the trailing part of the source value reported in the error message can be different from the entered value. This is because FLOAT is an approximate numeric datatype whose internal representation (and rounding upon display) are platform-dependent.

To correct this problem, choose a larger target format for display. Use the str function to determine the necessary format. The required format varies depending upon the number being converted and the accuracy of floating point numbers supported by your platform. To guarantee success, use a target of 25 characters.

If loss of precision (rather than display format) is a concern in the application, consider using an exact numeric datatype such as integer, numeric or decimal.

Additional information

Refer to the *Transact-SQL User's Guide* for information about the str function.

Version in which this error is raised

All versions

Error 266

Severity

10

Message text

```
Transaction count after EXECUTE indicates that  a COMMIT
or ROLLBACK TRAN is missing. Previous count = %ld,
Current count = %ld.
```

Explanation

When a stored procedure is run, Adaptive Server maintains a count of open transactions, adding 1 to the count when a transaction begins, and subtracting 1 when a transaction commits. When you execute a stored procedure, Adaptive Server expects the transaction count to be the same before and after the stored procedure execution. Error 266 occurs when the transaction count is different after execution of a stored procedure than it was when the stored procedure began.

The 266 error occurs most often when stored procedures are executed in chained mode. In chained mode, if there is no explicit begin transaction statement, an implicit begin transaction is issued. A begin transaction statement, whether implicit or explicit, increments the session transaction count by 1.

In chained mode, the following statements cause an implicit begin transaction to be issued: select, update, delete, insert, and fetch. (The fetch command begins a transaction only when close on endtran is off in chained mode and cursors can remain open across transactions. Refer to "set" in the *Reference Manual* for a discussion of the close on endtran option of the set command.)

Each begin transaction statement in a stored procedure, whether implicit or explicit, must be balanced by a commit transaction statement, which decrements the session transaction count by 1. Otherwise, the transaction count remains higher when the stored procedure exits than it was at the beginning of execution.

In addition, Error 266 occurs when you are using nested procedures, and procedures at each level of nesting include begin, commit, and rollback transaction statements. If a procedure at a lower nest level opens a transaction and one of the called procedures issues a rollback transaction, Error 266 occurs when you exit the nested procedure. The following example illustrates this situation:

```
1> create procedure proc1
2> as
3> begin transaction
     .
     .
     execute proc2
4> commit transaction
5> go
1> create procedure proc2
2> as
3> begin transaction
     .
     .
4> rollback/commit transaction
5> go
```

If *proc2* executes a rollback transaction, it returns with Error 266.

Action

In the following example, execution of the stored procedure "test" results in a 266 error:

```
1> use pubs2
2> go
1> create procedure test as
```

```
2> select * from titles
3> go
1> sp_procxmode test, chained
2> go
1> set chained on
2> go
1> execute test
2> go
```

The transaction count is 0 before executing stored procedure “test”. When the stored procedure is executed, its select statement causes a transaction to begin implicitly. Since there is no balancing commit transaction for the implicit begin transaction the transaction count is not decremented in the stored procedure. The stored procedure returns to the caller with a transaction count of 1. A non-fatal 266 error is raised.

To balance the implicit begin transaction within a stored procedure, issue a commit statement before exiting the stored procedure, as in the following example:

```
1> use pubs2
2> go
1> create procedure test as
2> select * from titles
3> commit transaction
4> go
1> sp_procxmode test, chained
2> go
1> set chained on
2> go
1> execute test
2> go
```

The transaction count is 0 before stored procedure execution. When the stored procedure is executed, the select results in an implicit begin transaction and the transaction count is incremented to 1. The commit transaction in the stored procedure decrements the transaction count to 0. The transaction count is the same at the end of stored procedure execution as it was in the beginning.

You can use the *@@trancount* global variable to check the transaction count at the beginning and end of the procedure. If the count is not equal, commit or roll back transactions as appropriate.

Note It is safer to commit within the stored procedure than to open the transaction explicitly before executing the stored procedure. Opening the transaction before executing the stored procedure carries the risk of leaving a transaction open if the execute procedure statement fails and causing other problems as your procedures get more complicated.

Additional information

For more information on transactions within stored procedures and the *@@trancount* global variable, refer to the section "Checking the Transaction Nesting Level with @@trancount" of "Using Batches and Control-of-Flow Language" in the *Transact-SQL User's Guide*.

Version in which this error is raised

All versions

Error 268

Severity

16

Message text

```
You can't run SELECT INTO in this database. Please check
with the Database Owner.
```

Explanation

This error occurs when you try to add rows to a table via the select into clause in a database that has the sp_dboption option select into/bulkcopy parameter disabled.

If this error occurs while trying to select into a temporary table, or while running a Sybase-supplied stored procedure, it is likely that the select into/bulkcopy option is not enabled in tempdb. Some stored procedures, for example sp_help and sp_helpsort, trigger this error because they select into temporary tables (which all belong to tempdb) to get reports from the system tables, and to update them. If tempdb does not have the sp_dboption option select into/bulkcopy enabled, Error 268 will occur.

Note When you install Adaptive Server, the select into/bulkcopy option is enabled in tempdb and turned off in all other databases.

The sp_dboption option select into/bulkcopy must be enabled to perform nonlogged operations. For example:

- Performing a select into a permanent (nontemporary) table.
- Performing a fast bulk copy with the bcp utility. Tables without triggers or indexes use the fast version of bcp by default, that is, they are not logged in order to save time.
- Executing the Transact-SQL command writetext (using the with log option of writetext causes it to be logged, so the select into/bulkcopy parameter is not needed).
- Executing the DB-Library routine dbwritetext.

Warning! You cannot dump the transaction log to a device after performing a nonlogged operation. Always perform a dump database after the nonlogged operation is complete, since performing nonlogged operations leaves changes recorded in the transaction log unrecoverable.

Action

Ask the System Administrator ("sa") or database owner ("dbo") to enable the sp_dboption option select into/bulkcopy for all databases affected by the query (including tempdb if applicable). The "sa" or "dbo" can use this procedure:

1 Use the master database and change the option:

```
1> use master
2> go
1> sp_dboption database_name,
2> "select into/bulkcopy", true
3> go

1> use database_name
2> go

1> checkpoint
2> go
```

2 To verify that the change is active, use sp_helpdb:

```
1> sp_helpdb database_name
2> go
```

In the output, the phrase select into/bulkcopy should appear, indicating that the option is enabled.

Warning! Dump your database before disabling the select into/bulkcopy option. If you insert nonlogged data (via a nonlogged operation) into the database, you will not be able to dump the transaction log to a device, and will not be able to recover your data.

By default, the select into/bulkcopy option is disabled when a database is first created. To change this default, use the procedure above to allow the option in the model database.

Additional information

If a table has indexes or triggers, then bcp will not run in the fast mode, so you do not need to set the select into/bulkcopy option to run bcp.

You cannot dump the transaction log to a device after performing a minimally logged operation. Trying to dump the transaction log after a minimally logged operation causes error messages that instruct you to dump the database instead.

Version in which this error is raised

All versions

Error 277

Severity

16

Message text

```
There was a transaction active when exiting the stored
procedure '%.*s'. The temporary table '%.*s' was dropped
in this transaction either explicitly or implicitly.
This transaction  has been aborted to prevent database
corruption.
```

Explanation

When exiting a stored procedure, Adaptive Server checks to see whether there is an active transaction (a transaction that has not been committed) and then checks to see whether any temporary objects exist for that uncommitted transaction.

If temporary objects exist, the transaction is rolled back, the temporary objects are dropped, and Error 277 is raised.

Action

In stored procedures, make sure all begin transaction statements have corresponding commit transaction or rollback transaction statements.

In chained mode, if there is no explicit begin transaction statement, an implicit begin transaction is issued. If you are using chained mode, an explicit commit transaction or rollback transaction statement is required to end the transaction.

If you confirm that all begin transaction statements have corresponding commit transaction or rollback transaction statements, check to see whether the stored procedure is exiting without completing its processing.

Version in which this error is raised

All versions

Query Processor Errors

This section contains error messages for the Adaptive Server query processor.

Error 311

Severity

16

Message text

```
The optimizer could not find a unique index which it
could use to scan table '%.*s' for cursor '%.*s'.
```

Explanation

For an explicit updatable cursor scan, Adaptive Server requires that a unique index exist on the table. A unique index ensures that the cursor will be positioned at the correct row the next time a fetch is performed on that cursor. Error 311 occurs when a unique index does not exist during a scan required for a cursor marked for update.

Action

Create a unique index using one of the following methods:

- Drop your existing index. Create an index with the unique option or use alter table add constraint with the unique option.
- Create a new index using create index with the unique option or use alter table add constraint with the unique option.

  Note When you create an index with the unique option, if there is a duplicate key value or if more than one row contains a null value, the command is aborted and Error 1505 is raised, showing the duplicate value. Refer to **Error 1505** if that happens.

- Add an IDENTITY column using alter table and the identity option.

If your table has one or more IDENTITY columns, you can, for future indexes you create for this table, use the database option identity in nonunique index. When this option is set, any nonunique index created on a table with an IDENTITY column will have the IDENTITY column automatically included as the last key field of the index. This allows an otherwise nonunique index to be used for a cursor marked for update.

However, all indexes would then be considered unique by the optimizer since every index that is created would be unique and this could result in poor performance for some queries. Since the identity in nonunique index option does not affect existing indexes, only future indexes you create, the existing 311 error will not be solved.

Warning! Do not use the identity in nonunique index option unless you plan to test your queries.

To use the identity in nonunique index option, the table must already have an IDENTITY column, either from a create table statement or by setting the auto identity database option to true before creating the table.

The commands to set this option are:

```
1> use master
2> go

1> sp_dboption database_name,
2> "identity in nonunique index", true
3> go

1> use database_name
2> go

1> checkpoint
2> go
```

Additional information

Refer to the *ASE Reference Manual* for information about create index, alter table, and sp_dboption.

Version in which this error is raised

All versions

Error 313

Severity

16

Message text

```
A dirty read scan requires a unique index (unless the
user forced a non-unique index or table scan). The
optimizer could not find a unique index to use.
```

Explanation

Isolation level 0 allows transactions to read uncommitted data (*dirty reads*). Since level 0 scans do not acquire locks, modifications by other processes can cause rows being scanned to move. When this occurs, Adaptive Server restarts its scan. To restart properly, the scan must use the same key used to find the modified row, and start again from the next key in the index. Thus, the scan must be on a unique index. Error 313 occurs when a unique index does not exist during a dirty read scan.

Action

Create a unique index using one of the following methods:

- Drop your existing index. Create an index with the unique option or use alter table add constraint with the unique option.
- Create a new index using create index with the unique option or use alter table add constraint with the unique option.

  Note When you create an index with the unique option, if there is a duplicate key value or if more than one row contains a null value, the command is aborted and Error 1505 is raised, showing the duplicate value. Refer to **Error 1505** if that happens.

- Add an IDENTITY column using alter table and the identity option.

If your table has one or more IDENTITY columns, you can, for future indexes you create for this table, use the database option identity in nonunique index. When this option is set, any nonunique index created on a table with an IDENTITY column will automatically have the IDENTITY column included as the last key field of the index. This allows an otherwise nonunique index to be used for a cursor declared for update.

However, all indexes would then be considered unique by the optimizer since every index that is created would be unique and this could result in poor performance for some queries. Since the identity in nonunique index option does not affect existing indexes, only future indexes you create, the existing 313 error will not be solved.

Warning! Do not use the identity in nonunique index option unless you plan to test your queries.

To use the identity in nonunique index option, the table must already have an IDENTITY column, either from a create table statement or by setting the auto identity database option to true before creating the table.

The commands to set this option are:

```
1> use master
2> go

1> sp_dboption database_name,
2> "identity in nonunique index", true
3> go

1> use database_name
2> go

1> checkpoint
2> go
```

forceindex

forceindex with a clustered nonunique index uses key values to restart the scan at the first row that has key values. You can use forceindex to force Adaptive Server to use a nonunique index for your table scan. However, the restart will either be approximate or it will fail altogether if a row becomes invalidated. If it fails, your query will abort.

Warning! It is recommended that you do *not* use forceindex with dirty reads.

Additional information

Refer to the *ASE Reference Manual* for information about create index, alter table, and sp_dboption.

Version in which this error is raised

All versions

Error 314

Severity

16

Message text

```
WARNING: A non-unique clustered index has been forced
on an isolation level 0 scan on table '%.*s'.  If the
scan must restart, the scan will be repositioned at the
beginning of the duplicate key group.  Thus, it is
possible that this can infinitely loop.
```

Explanation

Isolation level 0 allows transactions to read uncommitted data (*dirty reads*). Since level 0 scans do not acquire locks, modifications by other processes can cause rows being scanned to move. When this occurs, Adaptive Server restarts its scan. To properly restart, the scan must use the same key used to find the modified row and start again from the next key in the index. Thus, the scan must be on a unique index.

Error 314 occurs when a unique index does not exist during a dirty read scan and you have used forceindex to force Adaptive Server to use a nonunique clustered index for your table scan. In this case, Adaptive Server will use key values to restart the scan at the first row that has key values. However, the restart will either be approximate or will fail altogether if a row becomes invalidated. If it fails, your query will abort. In addition, it is possible that the scan can go into an infinite loop.

Action

This is a warning message. No action is required. However, since the consequences are severe, consider using the following options instead of forceindex.

Create a unique index using one of the following methods:

- Drop your existing index. Create an index with the unique option or use alter table add constraint with the unique option.
- Create a new index using create index with the unique option or use alter table add constraint with the unique option.

  Note When you create an index with the unique option, if there is a duplicate key value or if more than one row contains a null value, the command is aborted and Error 1505 is raised, showing the duplicate value. Refer to **Error 1505** if that happens.

- Add an IDENTITY column using alter table and the identity option.

If your table has one or more IDENTITY columns, you can, for future indexes you create for this table, use the database option identity in nonunique index. When this option is set, any nonunique index created on a table with an IDENTITY column will have the IDENTITY column automatically included as the last key field of the index. This allows an otherwise nonunique index to be used for a cursor declared for update.

However, all indexes would then be considered unique by the optimizer since every index that is created would be unique and this could result in poor performance for some queries. The identity in nonunique index option does not affect existing indexes, only future indexes you create.

Warning! Do not use the identity in nonunique index option unless you plan to test your queries.

To use the identity in nonunique index option, the table must already have an IDENTITY column, either from a create table statement or by setting the auto identity database option to true before creating the table.

The commands to set this option are:

```
1> use master
2> go

1> sp_dboption database_name,
2> "identity in nonunique index", true
3> go

1> use database_name
2> go

1> checkpoint
2> go
```

Additional information

Refer to the *ASE Reference Manual* for information about create index, alter table, and sp_dboption.

Version in which this error is raised

All versions

Error 403

Severity

16

Message text

```
Invalid operator for datatype op: %s type: %s.
```

Explanation

This error occurs during Adaptive Server expression processing, when an operator in the query is applied to a datatype for which that operator is not valid. An operator is a logical or arithmetic expression such as “+” or “-”.

For example:

```
1> select pub_id * pub_name from publishers
2> go
Msg 403, Level 16, State 1:
Line 1:
Invalid operator for datatype op: MULTIPLY type: CHAR.
```

Action — Correct your query and run it again.

Additional information — Refer to the *Reference Manual* and the *Transact-SQL User's Guide* for information about the use of operators in queries.

Version in which this error is raised — All versions

Error 404

Severity — 19

Message text — `Too many ANDs or ORs in expression (limit %d per expression level). Try splitting query or limiting ANDs and ORs.`

Explanation — This error occurs when you use more than 251 and or or expressions in a query.

This error also occurs if you exceed the limit of 251 values in a where in clause. For example:

```
1> select * from old_table
2> where old_column in (1,2,3...,253,254,...)
3> go
```

Action — Write the query so that the limit is not exceeded. Often, this requires that you split the query.

Additional information — Refer to “Search Conditions” in the *Transact-SQL User's Guide* for more information.

Version in which this error is raised — All versions

Error 414

Severity — 16

Message text — `The current query would generate a key size of %d for a work table.  This exceeds the maximum allowable limit`

```
of %d.
```

Explanation

When you execute a query, Adaptive Server may need to create work tables to temporarily store query results. For example, a work table is used when duplicate rows must be removed in processing a query with an aggregate function. A sysindexes row is built for the work table after checking that the specified command does not violate any limitations on keys for user tables.

Error 414 is raised when you execute a query containing an aggregate function, and the total length of columns named in the group by clause of the query exceeds the maximum limit of 256 bytes.

Action

Check the command for possible violations of the 256 byte size limit. You can correct the problem by doing one of the following:

- Reduce the columns named in the group by clause until the combined column length is no more than 256 bytes.
- If it is necessary to group by a large character column, consider using the substring string function on the column in the group by clause. This allows Adaptive Server to build a composite key to group the result set using only a portion of the character string. For example:

  ```
  1> select * from titles
  2> group by title, substring(notes,1,10)
  ```

  This query uses only the first 10 bytes of notes (a 200 byte varchar column) to group the data.

Version in which this error is raised

All versions

Error 428

Severity

```
20
```

Message text

Version 11.5 and later

```
There are more than %d referential constraints on table
%.*s. Please reduce the number of referential
constraints before trying this query.
```

Version 11.0.x

```
Too many table names or referential constraints in the
query, maximum allowed table references is '%d'.
```

Explanation

During query processing, Adaptive Server checks for the existence of foreign keys and dependent foreign keys (a foreign key is a column or combination of columns whose values match the primary key). A range table entry is created for each reference check and foreign key constraint. “%d” in the error message is the maximum number of table references allowed. Error 428 is raised when this limit is reached.

Action

When setting up constraints on your tables, determine the maximum number of tables that might be touched by an update, insert, or delete statement.

To determine which constraints exist for a table, type:

```
1> use database_name
2> go

1> sp_helpconstraint table_name
2> go
```

If necessary, drop some of the constraints on the table:

```
1> alter table table_name
2> drop constraint constraint_name
3> go
```

Additional information

Refer to the *ASE Reference Manual* for information about constraints.

Version in which this error is raised

All versions

Error 511

Severity

16

Message text

```
Updated or inserted row is bigger than maximum  size (%d
bytes) allowed for this table.
```

Explanation

This error occurs when you try to insert or update a row that is longer than the maximum allowable length. Rows consist of 1962 characters; allow 2 characters of row overhead for APL tables, for a maximum usable row size of 1960. For DOL tables, subtract two characters per varchar column in determining usable row size.

Error 511 is caused by database design errors (for example, a table designed with the potential for rows containing more than the maximum allowable characters).

The following warning is given when you create a table that has the potential for rows exceeding the maximum row size (that is, the maximum length of all columns added up is greater than the allowable number of characters):

```
Msg 1708, Level 16, State 1:Warning: Row size could exceed row size limit,
which is %d bytes.
```

Action

If Error 511 is being caused by a table containing rows with more than the maximum row size, divide the table into two or more tables so that no row length is more than the allowable number of characters.

If the 511 error occurring on your database does not appear to be caused by the above situation, call Sybase Technical Support.

Version in which this error is raised

All versions

Error 512

Severity

16

Message text

```
Subquery returned more than 1 value.  This is illegal
when the subquery follows =, !=, <, <= , >, >=, or when
the  subquery is used as an expression.
```

Explanation

When an expression subquery returns more than one result, it violates the relational operator rule for the outer query, and Error 512 occurs.

An example of an expression subquery that returns one result follows:

```
1> select * from table_one where x =
2> (select sum(a) from table_two
3> where b = table_one.y)
4> go
```

An example of a query that returns more than one result and causes Error 512 follows:

```
1> use pubs2
2> go

1> select authors.au_id from authors where
2> authors.au_id = (select titleauthor.au_id
3> from titleauthor)
4> go
Msg 512, Level 16, State 1:
Line 1:
```

```
Subquery returned more than 1 value.  This is illegal
when the subquery follows =, !=, <, <= , >, >=, or when
the subquery is used as an expression.
```

Action

To correct the problem in the example, use “in” in place of “=”, as in the following example:

```
1> select authors.au_id from authors where
2> authors.au_id in (select titleauthor.au_id
3> from titleauthor)
4> go
au_id
-----------
172-32-1176
213-46-8915
  .
  .
899-46-2035
998-72-3567

(19 rows affected)
```

Version in which this error is raised

All versions

Error 515

Severity

```
16
```

Message text

```
Attempt to insert NULL value into column  '%.*s', table
'%.*s'; column does not allow nulls. Update fails.%S_EED
```

Explanation

When you create a table, you can explicitly define whether each column should allow null values. If you do not specify NULL or NOT NULL for a column when you create the table, the default value will be NOT NULL. If you use sp_dboption to set allow nulls by default to TRUE for the database, any new table that is created will have columns with the default value NULL.

Error 515 occurs at run time when a column has a NOT NULL default and you try to insert a NULL value into that column (for example, if a user does not enter a value for that column). The error message includes:

- The name of the affected column.
- The name of the affected table.

- The EED (extended error data array) includes the database name, owner, table name, and column name. This information is not displayed but you will see it in the message definition if you select Error 515 from sysmessages.

The following sample would cause a 515 error:

```
1> create table table1 (column1 varchar (5))
2> go

1> declare @c varchar(5)
2> insert into table1 values (@c)
3> go
Msg 515, Level 16, State 3:
Server 'SERVER_NAME', Line 2:
Attempt to insert NULL value into column 'column1', table 'test.dbo.table1';
column does not allow nulls. Update fails.
```

Action

To determine whether a column has NULL or NOT NULL defined, enter commands like the following, replacing "pubs2" by the database in which your table resides and "titleauthor" by the table in which your column resides:

```
1> use pubs2
2> go

1> sp_help titleauthor
2> go
Name          Owner   Type
-------------------------------------------------------------
titleauthor   dbo     user table
Data_located_on_segment        When_created
------------------------------ --------------------------
default                        Oct 27 1994 10:09AM
Column_name Type Length Prec Scale Nulls Default_name Rule_name Identity
-----------------------------------------------------------------------
au_id       id       11 NULL NULL  0     NULL         NULL      0
title_id    tid       6 NULL NULL  0     NULL         NULL      0
au_ord      tinyint   1 NULL NULL  1     NULL         NULL      0
royaltyper  int       4 NULL NULL  1     NULL         NULL      0
```

The Nulls column indicates whether null values are allowed. A value of 0 for the column means nulls are not allowed and 1 means null values are allowed.

To change the default for a column from NOT NULL to NULL:

1 Use bcp to copy the data out of the existing table.

2 If you want the modified table to have the same name as the existing table, drop the old table.

3 Re-create the original table, specifying NULL for the column you want to change.

4 Use bcp to put back the data for the table.

If you want *new* tables in the database to have the default NULL, use the following commands for the database:

```
1> use master
2> go

1> sp_dboption database_name,
2> "allow nulls by default", true
3> go

1> use database_name
2> go
1> checkpoint
2> go
```

where *database_name* is the name of the database whose behavior you want to change.

Version in which this error is raised

All versions

Error 530

Severity

16

Message text

```
Attempt to insert NULL value into column %d in work
table (table id %ld); column does not allow NULLS.
UPDATE fails.
```

Explanation

During the run-time phase of updates, Adaptive Server sometimes uses worktables to temporarily store query results. Later in the query processing, Adaptive Server selects the values from those tables.

If, as a result of your query, Adaptive Server tries to insert a null value into a column of a worktable, and nulls are not allowed for the column, Error 530 occurs. The command is aborted and the update fails. Error 530 is caused by an Adaptive Server problem.

Action

Depending on the context of the error, you may be able to rewrite the query (for example, supply a column name or specify isnull). If the error reoccurs, or you cannot write the query in another way, call Sybase Technical Support.

Version in which this error is raised

All versions

Error 539

Severity

20

Message text

Version 11.9.2 and Later

```
Unexpected internal access methods error %d, state %d.
Please report to Technical Support the following
information: spid=%d, dbid=%d, objectid=%ld, curcmd=%d
(%s), plasterror=%ld, pstat=0x%x, p2stat=0x%x,
xactid=(%ld, %d).
```

Version 11.5.x and Earlier

```
Access methods error code %d encountered but run does
not know how to handle it: spid = %d, db id = %d,
objectid = %ld, curcmd = %d (%s), plasterror = %ld,
pstat = 0x%x, p2stat = 0x%x, pattention = %ld, xactid =
(%ld, %d).
```

Explanation

Error 539 is raised during query processing when Adaptive Server attempts to execute an access method, but the attempt fails with an unknown failure in the access method. This is why Error 539 reports error code 0.

Error 539 is due to an Adaptive Server problem.

Action

If additional errors accompanied the 539 error, resolve them using the relevant writeups in this document.If the problem persists, call Sybase Technical Support.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level
- Server error log
- Text of all error messages.

Version in which this error is raised

All versions

Error 540

Severity

16

Message text

```
Schema for table '%.*s' has changed since  compilation
of this query. Please re-execute query.
```

Explanation

A table's schema consists of the table structure along with any constraints, rules, defaults and so on. A change to any of these elements (for example, creating or dropping a constraint) constitutes a change in the table's schema.

Query processing involves the following basic steps:

- Adaptive Server examines the SQL statements and, if it passes syntax checks, the query is parsed to produce an internal representation called the query tree.
- The query tree is read into an area of memory known as the procedure cache. Adaptive Server then determines the optimal access strategy for implementing the query, compiles the query to produce the query plan (the internal code necessary to implement the query), and places the plan in cache.
- The query plan is executed to produce the results.

For efficiency Adaptive Server can reuse an existing query plan in cache.

Error 540 is raised:

- If this was an ad hoc query, and the table's schema has changed since the query plan was generated. Adaptive Server does not attempt to recompile an ad hoc query if there is a mismatch between the query plan and the table's schema.
- If update statistics is run for a table while a query which references that table is executing.

This is not a serious error; however, it may prevent you from running some ad hoc queries.

Action

To address this error:

- Retry the query. This ensures that Adaptive Server can build the appropriate query plan for the revised schema.
- Avoid running update statistics on a table while a query which references the table is executing.

Version in which this error is raised

All versions

Error 547

Severity

16

Message text

```
Dependent foreign key constraint violation in a
referential integrity constraint. dbname =  '%.*s',
table name = '%.*s',  constraint name = '%.*s'.%S_EED
```

Explanation

Adaptive Server provides integrity constraints to help you maintain logical data integrity in a database. Referential integrity (or foreign key) constraints are a type of constraint which require that data being inserted into a given table column already has matching data in another column (the target column), which may be in the same table or another table. The column on which the constraint is declared can be considered the child, or dependent, in a parent-child relationship.

Error 547 is raised when updating or deleting rows from a parent table would remove target columns matching dependent data in the child table. For example:

```
1> create table parent
2> (a int primary key,
3> b int,
4> unique (b))
5> go

1> create table child
2> (c int primary key,
3> d int references parent(b))
4> go

1> insert parent values (11,22)
2> insert parent values (13,26)
3> insert child  values (101,22)
4> go

1> delete parent where a=13
2> go
(1 row affected)
1> delete parent where a=11
2> go
Msg 547, Level 16, State 1: Line 1:
Dependent foreign key constraint violation in a referential integrity
constraint. dbname =  'hrdb', table name = 'parent', constraint name =
'detail_d_1088006907'.
Command has been aborted.
(0 rows affected)
```

```
                    1> update parent
                    2> set b=29
                    3> where a=11
                    4> go
Msg 547, Level 16, State 1: Line 1:
Dependent foreign key constraint violation in a referential integrity
constraint. dbname =  'hrdb', table name = 'parent', constraint name =
'child_d_1088006907'.
Command has been aborted.
(0 rows affected)
```

Notice that you may drop or update rows in parent provided you do not affect the referential integrity constraints.

Action

Delete or update the dependent data in the child table before deleting or changing the parent data.

Additional information

Use the system procedure sp_helpconstraint to view the referential constraints in effect for a table.

Version in which this error is raised

All versions

Error 551

Severity

20

Message text

An unknown EVAL was sent to the execution module.

Explanation

The instructions Adaptive Server creates to run a query are contained in an evaluation list that contains instruction-argument pairs. Error 551 occurs when Adaptive Server receives an illegal instruction. It is probably caused by an incorrectly compiled query tree.

Action

Drop and re-create the procedure or trigger being run:

1. If you do not have a script for re-creating your procedure or trigger, get the text of the procedure or trigger:

   ```
   1> use database_name
   2> go

   1> sp_helptext object_name
   2> go
   ```

   where *database_name* is the name of the database in which the procedure or trigger resides and *object_name* is the name of the procedure or trigger.

Alternatively, use the defncopy program to copy the procedure or trigger definition to a file. Refer to the Adaptive Server utility programs manual for details.

2 Drop the procedure or trigger:

```
1> drop procedure object_name
2> go
```

or:

```
1> drop trigger object_name
2> go
```

3 Re-create the procedure or trigger. Refer to the *ASE Reference Manual* for information about the create procedure and create trigger commands.

If the 551 error occurs again, create a scenario to reproduce the problem and call Sybase Technical Support.

Additional information

Have the following information ready before calling Technical Support:

- Adaptive Server version and SWR level
- Copy of the query which raises the error

Version in which this error is raised

All versions

Error 584

Severity

20

Message text

```
Explicit value specified for identity field  in table
'%.*s' when IDENTITY_INSERT is set to OFF.
```

Explanation

Each table can include a single IDENTITY column. IDENTITY columns store sequential numbers that are generated automatically by Adaptive Server. The value of the identity column can uniquely identify each row in a table.

The Adaptive Server query processing option set identity_insert determines whether explicit inserts into a table's IDENTITY column are allowed. Inserting a value into the IDENTITY column allows you to specify a “seed” value for the column or to restore a row that was deleted in error. Setting identity_insert on allows the table owner, database owner, or System Administrator to explicitly insert a value into an IDENTITY column. Unless a unique index exists on the IDENTITY column, any positive value without regard to uniqueness may be inserted in IDENTITY columns when identity_insert is set to on.

Setting identity_insert off prohibits inserts to IDENTITY columns.

Error 584 is raised if you attempt to insert an explicit value into an IDENTITY column when identity_insert is set to off.

Action

Set identity_insert on before attempting to insert an explicit value into an IDENTITY column. identity_insert should be reset to off once the insert operation is complete.

Additional information

The syntax for setting the identity_insert option is as follows:

```
1> set identity_insert `table_name'  {on | off}
2> go
```

where *table_name* is the base table for the column. Only the table owner, database owner, or System Administrator can set this option.

Version in which this error is raised

All versions

Access Method Errors

This section contains error messages for problems in Adaptive Server access methods.

Error 601

Severity

21

Message text

```
Descriptor for system table '%ld' in database  '%d' not
found in the descriptor hash table.
```

Explanation

Adaptive Server expects every database to have a systhresholds table. Error 601 occurs when the Adaptive Server Threshold Manager could not open the systhresholds system table in a database because the descriptor for systhresholds was not found. The failure to open systhresholds is probably due to corruption in the database named in the error message.

Action

1 Determine the name of the database from the database ID in the error message:

```
1> use master
2> go
1> select * from sysdatabases where dbid = ID
2> go
```

where *ID* is the database ID from the error message.

2 Determine whether the system table object ID displayed in the error message exists:

```
1> use database_name
2> go

1> select * from sysobjects where id = object_ID
2> go
```

where *database_name* is the name from step 1 and *object_ID* is the system table ID listed in the error message.

3 If the ID exists for the table displayed in the message, then the 601 error may be due to corruption in the memory structure used to hold the descriptor. Shutting down and restarting Adaptive Server should clear the problem. If this does not work, go on to step 4.

4 If the ID does not exist, there is probably corruption in your database. Do the following:

   a Run dbcc checkalloc, dbcc checkcatalog, and dbcc checkdb for the database.

   b Determine whether hardware problems exist by checking your operating system error log.

   c Call Sybase Technical Support. They may be able to help you recover your database, but you will probably have to recover your database from backups.

Additional information

Before contacting Technical Support, be prepared to provide:

- Adaptive Server version and SWR version level
- Complete text of all error messages

- Adaptive Server error log output
- Operating system error log output
- dbcc checkalloc, dbcc checkcatalog, and dbcc checkdb output for the database listed in the error message

Version in which this error is raised

All versions

Error 603

Severity

19

Message text

```
There are not enough system session descriptors
available to run this query. The maximum number
available to a process is  %d. Split query and rerun.
```

Explanation

This error occurs when Adaptive Server runs out of *system* session descriptors when running a query. Error 603 can occur on commands which create, drop, and update objects, as these commands often have many dependencies on system tables which Adaptive Server must handle to complete the command.

Session descriptors are the data structures used to manage access to Adaptive Server objects. Session descriptors come in three forms: *user* session descriptors (including worktable session descriptors), *system* session descriptors, and *worktable* session descriptors.

Session descriptors are used in the following situations:

- When your query is compiled or executed, a *user* session descriptor is reserved and assigned to each table involved in the query. Currently, a maximum of 16 user session descriptors are possible, limiting the number of tables involved in a query to that number. If a single table is referenced more than once in the from clause (a “self join”), a user session descriptor is reserved for each reference. User session descriptors correspond to the number of open objects configuration parameter.
- *System* session descriptors are used for system tables that are not named in the query. For example, the following query:

  ```
  select x from t where y = 3
  ```

  prompts Adaptive Server to open sysobjects and syscolumns to look up the table t and its columns x and y. A maximum of six system session descriptors are available.

- A *worktable* session descriptor is used to manage references to worktables. A worktable is required for order by, group by, distinct, and some other operators. Also, the query optimizer may choose a plan involving reformatting which copies rows from a table into a worktable, creates an index, then performs joins between the worktable and other tables in the query. There are 12 worktable session descriptors that can be used at any one time.

You cannot change the number of session descriptors available to a query.

Action

When you get Error 603, you do indeed have to simplify your query or command. To recover from this error, split up your query or simplify the command and run it again.

For example, if you are entering a drop procedure command and that procedure contains other procedures or triggers, you can simplify your command by dropping the other procedures or triggers manually before running the drop procedure command.

Version in which this error is raised

All versions

Error 605

Severity

21

Message text

```
An attempt was made to fetch logical page  '%ld' in
database '%.*s' from cache '%.*s'. Page belongs to
object with id  '%ld', not to object '%.*s'.
```

Explanation

This serious error occurs when Adaptive Server discovers page allocation corruption. Adaptive Server tries to access a particular object but discovers a page in the page chain for that object whose object ID is different than that of the object being accessed. There is probably either a damaged page chain or an invalid entry in the system tables for that object.

Usually this error is detected after the corruption has been written to the database on disk (hard 605), but it can also occur entirely in cache without the damage ever being written to disk (transient 605), and is not associated with data corruption. See the next section for more information about the different kinds of 605 errors and their causes.

Note When an error is transient, it means it is reported even though no error condition exists or it exists only in memory. Generally, a transient error is cleared by a reboot of Adaptive Server, whereas a hard error does not go away with an Adaptive Server reboot.

605 Error in Version 11.0.3 and Later

Starting with Version 11.0.3, 605 error handling was enhanced to allow for more accurate diagnosis of object ID mismatch problems. The 605 error is now raised only when the objects in question are user objects. Other related types of errors, previously reported as 605 errors, now result in one of the errors 691 through 697. Refer to the writeups for these errors for details.

How Much Corruption?

Assuming the error is not a transient 605, corruption usually exists in the second object specified in the 605 error text. If the dbcc checkdb and dbcc checkalloc commands do not report additional errors, the first object mentioned is not corrupt. Because the 605 error can mask the existence of other errors, you must run the dbcc checks to determine the extent of the damage.

Instead of listing two object names, this error might contain other information:

- The error displays a number greater than zero. An attempt was made to refer to an object ID that does not exist in the system table sysobjects.
- If the error states that a page belongs to object “ALLOCATION,” some of the allocation structures used by the database may be corrupted.

Hardware and Software Causes for 605 Errors

Table 1-2: Hardware causes of Error 605

Hardware Cause	Additional Information	Transient or Hard
Overlapping partitions on disk (UNIX only).[a]	Often this occurs when the server is installed on a partition that overlaps other partitions.	Hard
Problem with device driver	--	Hard
Problem with controller	--	Hard

Hardware Cause	Additional Information	Transient or Hard
Bad disk	"Data retry" messages in */var/adm/messages* (UNIX).	Hard
a. For more information about overlapping partitions, refer to Chapter 2, "Correct Use of Raw Partitions".		

Table 1-3: Software causes of Error 605

Software Causes	Additional Information	Transient or Hard
Loading a master database on a master device that does not have rows in sysusages for dbid 1,2,3 identical to the device it was dumped from.	--	Hard
Loading a master database that has syscharsets in a location that is different from the default.	Caused by changing Adaptive Server sort order after operating it for some time.	Hard
Operating system notifies Adaptive Server that an I/O operation is complete when it is not. Any time buffer information about a page is different from the disk information about the page (and the disk has the correct information).	Adaptive Server connection often dies. Running dbcc checkdb and dbcc checkalloc should produce no more 605 errors. Corruption is not on disk. First, checkpoint the other databases. Then, perform a shutdown...with nowait and restart Adaptive Server.	Transient

Action

1 Run the dbcc checktable command on the second object specified in the error message.

2 To help determine the full extent of the corruption, run the dbcc checkdb and dbcc checkalloc commands as soon as feasible.

3 Check the Adaptive Server error log for other errors which often accompany a 605 error. If the 605 error is not transient, the problem is severe and you will probably need to restore from known clean backups.

4 Either of the following situations may also apply to your Adaptive Server:

- If this error occurs after rebuilding the master device or reconfiguring the configuration block to the defaults, then the sort order on the new device may not match the sort order in the dump. If the sort orders do not match, call Technical Support for further assistance.

- To check if the error occurred as the result of hardware failure, examine your operating system error log and correct hardware problems.

If you have dial-in facilities, Technical Support may be able to help you recover from this error without having to restore from backups. However, this recovery procedure often results in a loss of data, and therefore is often unacceptable, especially for production systems.

Additional information

For information on locating device fragments, refer to Chapter 2, "How to Determine Which Physical Devices a Database is On".

Have the following information ready before contacting Technical Support:

- Adaptive Server version and SWR version level
- Complete text of all error messages
- Adaptive Server error log output
- Operating system error log output
- Output of dbcc checkdb, dbcc checkalloc, and dbcc page for the corrupted page

Related Page Mismatch Errors in Version 11.0.3 and Later

Starting with Version 11.0.3, many page mismatch problems no longer raise 605 errors, but are instead reported with the following error numbers:

Error	Description
691	Page 0 is read incorrectly, or an attempt is made to read a negative page number.
692	Attempt to read an uninitialized page.
693	I/O request was made prematurely, causing an object ID mismatch.
694	The device or operating system may be suspect.
695	The page has been overwritten, for example page 10 on disk may actually hold page 20. Hard error.
696	Objects are crossed in tempdb. Similar to a 605 except that objects are in tempdb.
697	Results from memory corruption.

Refer to the writeups for these errors in this manual for more information.

Version in which this error is raised

All versions

Error 611

Severity

21

Message text

```
Attempt made to end a transaction that is idle or in the
middle of an update.
```

Explanation

This error occurs when a termination request event occurs during the processing of a transaction. Error 611 can be caused by:

- Massive processing using group by or order by.
- A stored procedure that references other stored procedures (and is in the midst of calling those procedures) being dropped.
- An Adaptive Server problem.

Action

To determine what might have caused the above error message, check for the following:

- dbcancel calls from within DB-Library exit processing routines that fail to test all error conditions
- Users typing Ctrl-c during heavy use of temporary tables

Error 611 may produce a stack trace that gives a portion of the SQL that caused the error. Obtain the entire stack trace, along with any errors that might have occurred prior to or following the 611 error, to determine the internal state of the process at the time of the failure.

Version in which this error is raised

All versions

Error 614

Severity

21

Message text

```
A row on page %ld was accessed that has an  illegal
length of %d in database '%.*s'.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server accesses a data or index row whose length is smaller than the minimum row size or greater than the maximum row size.

The minimum length of a row for each object is stored in the minlen column of sysindexes and in each data or index page header. The size of a data row or index row is 1962 characters; allow 2 characters of row overhead for APL tables, for a maximum usable row size of 1960. For DOL tables, subtract two characters per varchar column in determining usable row size.

This error can occur under the following conditions:

- During normal processing, when Adaptive Server tries to access the row specified by the error message.
- During database recovery (database recovery occurs during Adaptive Server start-up or when a load database or load transaction command is processed).

Error 614 can be caused by data corruption during normal processing (for example, an operating system panic occurs, causing interruption in disk writes when using UNIX files for Sybase database devices). This may be due to a problem with Adaptive Server, the operating system, or hardware.

Action

Error 614 is usually the result of a more serious underlying problem, and recovering from this error depends on when the error occurred. Determine whether the error occurred during normal processing or during database recovery, then follow the appropriate set of instructions in this section.

If the Error Occurred During Normal Processing

1 Use the procedure in Chapter 2, “How to Find an Object Name from a Page Number” to identify which table and index correspond to the page number from the error message text.

2 If the object encountering the error is *not* a system table (a system table's object ID is less than 100), continue with step 3.

If the object encountering the error is a system table and the index ID is *not* 0, refer to Chapter 2, “How to Fix a Corrupted Index on System Tables” for instructions on how to repair the system table index.

If the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore the database from clean backups.

3 For user tables, if the index ID is 0, continue with step 4.

If the index ID is *not* 0, translate it into an index name:

```
1> use database_name
2> go
1> select name from sysindexes
```

```
2> where id = object_ID and indid = index_ID
3> go
```

To ensure that the information needed to re-create the index is available, run sp_helpindex on the index prior to dropping it.

Drop the index.

Re-create the index. This clears the corruption in most cases.

Run dbcc checktable on the table to verify that the corruption is gone.

4 If the index ID is 0, do one of the following:

- Restore the database from clean backups.
- Refer to Chapter 2, “How to Rescue Data from a Corrupted Table”.

Warning! Some data might be lost on this page if you recover your table using bcp or select into (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to see which rows are different (corrupted).

Before dumping your database, make sure it works correctly. Run the following commands prior to each dump:

1 dbcc checkdb.

2 dbcc checkalloc or dbcc checkalloc with the fix option. (Refer to Chapter 2, “How to Fix and Prevent Allocation Errors” for information about how to run these commands in multi-user mode and how to prevent spurious allocation errors from dbcc commands.)

If the Error Occurred During Database Recovery

When this error occurs during recovery, the database is marked suspect and is not accessible. Usually, you must load the database from backup. To do this, follow the instructions below:

1 If the database in the 614 error is tempdb, check the page number. If this page is not valid for tempdb, the problem may be on another database (possibly due to a bad device). Stop here and contact Sybase Technical Support.

If the page is valid for tempdb, or the 614 error refers to a database other than tempdb, continue with step 2.

2 Drop the database. If the drop fails, follow the instructions in Chapter 2, “How to Drop a Database When drop database Fails”.

3 Create a database for load. Make sure the database you create has sizes at least as large as those in sysusages for the original database (and that all other sysusages values match the original values). Refer to “create database” in the *ASE Reference Manual* for information on create database for load.

4 Load the database from backup. (Refer to “load database” in the *ASE Reference Manual*.)

5 Use the online database command to make the database available for use.

If loading from backups is not feasible, call Sybase Technical Support. Have the following documentation ready:

- Adaptive Server error log
- Text of all error messages
- select * from master..sysusages output

Preventing This Error on Recovery

To prevent Error 614 occurring on recovery, checkpoint each database that is being used before shutting down Adaptive Server.

Additional information

Refer to “Developing a Backup and Recovery Plan” in the *Adaptive Server Enterprise System Administration Guide* for information about how to safely create, dump, load, and re-create databases.

Version in which this error is raised

All versions

Error 622

Severity

```
20
```

Message text

```
Opentable was passed a varno of %d. Object  '%.*s'
already has that session descriptor in use.
```

Explanation

A session descriptor is an internal data structure in Adaptive Server which contains information about a table. All open objects in Adaptive Server require a descriptor structure. Whenever a connection to Adaptive Server attempts to open a table, it gets a session descriptor. *varno* refers to a variable number that Adaptive Server uses to maintain an array of session descriptors for a query. A unique *varno* is assigned to each session descriptor tracking the usage of a given object for the query.

Error 622 is raised when Adaptive Server attempts to open a table, but the *varno* points to a session descriptor that is already in use.

Error 622 breaks the user's connection to Adaptive Server.

Action

Error 622 is due to an Adaptive Server problem. Contact Sybase Technical Support for assistance when this error is raised.

Additional information

When calling Technical support, have the following information available:

- Adaptive Server version and SWR version level
- Adaptive Server error log
- Text of all error messages
- Text of the query that raises the error

Version in which this error is raised

All versions

Error 623

Severity

```
21
```

Message text

```
Attempt to retrieve row from page via RID  failed
because logical page %ld is not a data page. %S_RID.
%S_PAGE.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server tries to retrieve a row from a data page by specifying the row ID (RID) and the retrieval fails because the requested page was not a data page. This error is usually caused by a server problem.

After loading a database on top of an existing database, Error 623 may appear during recovery or while doing a checkpoint.

This error may or may not be transient (transient errors disappear after Adaptive Server is restarted). In the worst case, Error 623 may indicate corruption in your database. Perform the action below as soon as possible to determine how serious the error is and to avoid further corruption if corruption has occurred.

Action

Determine whether this error is transient by performing the following steps:

1 Refer to Chapter 2, “How to Find an Object Name from a Page Number” for instructions on how to use the dbcc page command to determine the table involved. Save this information and use it in the appropriate section below.

2 Using the object ID found in step 1, issue the following query to determine the table name:

```
1> select object_name(object_ID)
2> go
```

3 Using the table name displayed in the output from step 2, issue the following query to determine whether corruption has occurred:

```
1> dbcc checktable(table_name)
2> go
```

If no errors occur when you run dbcc checktable, the 623 error was probably transient. Shut down and restart the server.

If the Error is Not Transient

If errors *do* occur when you run dbcc checktable, refer to the writeups for those errors in this manual for instructions on how to recover from them or restore your database from a clean backup. To restore from backup, follow the instructions below:

1 Drop the database. If the drop fails, follow the instructions in Chapter 2, “How to Drop a Database When drop database Fails”.

2 Create a database for load. Make sure the database you create has sizes as least as large as those in sysusages for the original database and that all other sysusages values match the original values. Refer to “create database” in the *ASE Reference Manual* for information on create database for load.

3 Load the database from backup. (Refer to “load database” in the *ASE Reference Manual*.)

4 Use the online database command to make the database available for use.

If loading from backups is not feasible, call Sybase Technical Support. Have the following information ready:

- Adaptive Server version and SWR version level
- Adaptive Server error log
- Text of all error messages
- select * from master..sysusages output

Additional information

Refer to "Developing a Backup and Recovery Plan" in the *System Administration Guide* for complete information about how to safely create, dump, load, and re-create databases.

Version in which this error is raised

All versions

Error 624

Severity

21

Message text

```
Attempt to retrieve row from page via RID  failed
because the requested RID has a higher number than the
last RID on  the page. %S_RID.%S_PAGE.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server tries to retrieve a row from a data page by specifying the row ID (RID), but the retrieval fails because the requested row ID was a higher number than the last row ID on the page. This can happen:

- During normal processing, if the leaf page of a corrupt non clustered index points to an incorrect or nonexistent row ID on a data page.
- During database recovery at Adaptive Server startup (database recovery also occurs when a load database command is processed). Recovery fails because the data structure which contains the last checkpoint record in the transaction log is pointing to an incorrect or nonexistent record in the log.

Errors 624 and 625 occur under similar conditions.

Action

Recovery from this error depends on when the error occurred. Determine whether the error occurred during normal processing or during database recovery and then follow the appropriate instructions in this section:

- "If the Error Occurred During Normal Processing"
- "Index Unknown, Query Known"
- "Index and Query Both Unknown"
- "If the Error Occurred During Database Recovery"

If the Error Occurred During Normal Processing

Refer to Chapter 2, "How to Find an Object Name from a Page Number" for instructions about how to use the dbcc page command to determine the table involved. Save this information and use it in the appropriate section below.

The specific action you take depends on whether or not you know which index or query caused the problem. In general, dropping and re-creating the index should resolve the index corruption.

Index Unknown, Query Known

If you do not know which index is causing the problem, but you do know which query encounters the problem, use this section to correct it. If you do not know the index or query, go to "Index and Query Both Unknown".

1 Determine which index should be dropped by reading the showplan output for the query that encounters the error and determining which nonclustered index the query is using to access the table in question.

Turn on showplan and use the no execute mode:

```
1> set showplan on
2> go
1> set noexec on
2> go
```

Enter the query that was causing the problem. For example:

```
1> select title from titles where title > "Cooking"
2> go
STEP 1
The type of query is SELECT
FROM TABLE
titles
Nested iteration
Index: titlein
```

Turn showplan and the no execute mode off again:

```
1> set noexec off
2> go
1> set showplan off
```

```
2> go
```

2 If the index identified in step 1 is on a system table (object ID is less than 100), refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on repairing the system table index. Otherwise, drop and re-create the index identified in step 1 (in this example "titleind").

Index and Query Both Unknown

If you do not know either the index or the query, rebuild all nonclustered indexes on the table using the instructions in this section.

1 Examine the dbcc page output you obtained earlier.

If the object is a system table (the object ID is less than 100), refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index. Then go to step 3.

2 For a user table, use sp_helpindex to list all indexes on the table, and then rebuild all the nonclustered indexes using one of two methods:

- Drop and re-create each nonclustered index on the table.

Or:

- If a clustered index also exists on the table, drop and re-create it; this causes all nonclustered indexes to be rebuilt. If your table is large, you may not have the space to do this (a rough rule of thumb is 120 percent to 150 percent of your table size must be available).

3 Run dbcc checktable on the table to confirm that the problem has been resolved.

If the Error Occurred During Database Recovery

When this error occurs during recovery, the database is marked suspect and is not accessible. Usually, you must load the database from backup. Follow the instructions below, depending on the version level of your server (if you have dial-in facilities, Sybase Technical Support may be able to correct the problem so that recovery of this database can proceed):

1 Drop the suspect database using the instructions in Chapter 2, "How to Drop a Database When drop database Fails".

2 Create a database for load. Make sure the database you create has sizes as least as large as those in sysusages for the original database (and that all other sysusages values match the original values). Refer to "create database" in the *ASE Reference Manual* for more information on create database for load.

3 Load the database from backup (refer to “load database” in the *ASE Reference Manual*).

4 Use the online database command to make the database available for use.

Version in which this error is raised

All versions

Error 625

Severity

21

Message text

```
Could not retrieve row from logical page  %ld via RID
because the entry in the offset table for that RID is
less than  or equal to 0.
```

Note This error may be caused by a hardware problem.

Explanation

Every row in a database has a unique row ID (RID). The two parts of a row ID are a logical page number and a row number. The row number part of row IDs are stored at the end of every data page and indicate where a certain row is located on that page.

Commands that use row IDs include dbcc checkdb, dbcc checktable, and dbcc checkcatalog. They verify the consistency of indexes and use the row IDs stored in nonclustered indexes to quickly access the data rows.

Error 625 occurs when Adaptive Server tries to retrieve a row from a data page by specifying the row ID and it fails because the requested row ID has an illegal value (smaller than or equal to zero).

This error can occur under the following conditions:

- The data page of a nonclustered index is corrupt.
- Hardware failure.

Errors 624 and 625 occur under similar conditions.

Action

Examine your operating system error log file and the Adaptive Server error log to determine if hardware errors may have corrupted your database devices.

Refer to Chapter 2, “How to Find an Object Name from a Page Number” for instructions about how to use the dbcc page command to determine the table and the index involved. Save this information and use it in the appropriate section that follows.

The specific action you take depends on whether or not you know which index or query caused the problem. In general, dropping and re-creating the index will resolve the index corruption.

Index Unknown, Query Known

If you do not know which index is causing the problem, but you do know which query encounters the problem, use this section to correct it. If you do not know the index or query, use the next section.

1 Determine which index should be dropped by reading the showplan output for the query that encounters the error and determining which nonclustered index the query is using to access the table in question.

Turn on showplan and use the no execute mode:

```
1> set showplan on
2> go
1> set noexec on
2> go
```

Enter the query that was causing the problem. For example:

```
1> select title from titles where title > "Cooking"
2> go
STEP 1
The type of query is SELECT
FROM TABLE
titles
Nested iteration
Index: titleind
```

Turn showplan and the no execute mode off again:

```
1> set noexec off
2> go
1> set showplan off
2> go
```

2 If the index identified in step 1 is on a system table (object ID is less than 100), refer to Chapter 2, “How to Fix a Corrupted Index on System Tables” for instructions on how to repair the system table index.

3 Otherwise, drop and re-create the index identified in step 1 (in this example “titleind”).

If this did not clear the error, create a clustered index on the table or drop and re-create the existing clustered index. Note that creating a clustered index requires an amount of space equal to 120 percent to 150 percent of the size of the table where you want to create the index.

Both Index and Query Unknown

If you do not know either the index or the query, rebuild all nonclustered indexes on the table using the instructions in this section.

1 Look at the dbcc page output you obtained earlier.

2 If the object is a system table (object ID is less than 100), refer to Chapter 2, “How to Fix a Corrupted Index on System Tables” for instructions on repairing the system table index. Then go to step 4.

3 For a user table (object ID is 100 or greater), use sp_helpindex to list all indexes on the table, and then rebuild all the nonclustered indexes using one of two methods:

   - Drop and re-create each nonclustered index on the table.

   Or:

   - If a clustered index also exists on the table, drop and re-create it, which causes all nonclustered indexes to be automatically rebuilt. If your table is large, you may not have the space to do this (a rough rule of thumb is 120 percent to 150 percent of your table size must be available).

   If the above methods did not clear the error, create a clustered index on the table (if space is available). You may then drop the clustered index.

4 Run dbcc checktable on the table to confirm that the problem has been resolved.

Additional information

For more information, refer to “create database” in the *ASE Reference Manual*.

If you need to call Technical Support for assistance, have the following documentation ready:

- Adaptive Server version and SWR version level
- Adaptive Server error log
- Text of all error messages
- select * from master..sysusages output

Version in which this error is raised

All versions

Error 629

Severity

21

Message text

Version 11.9.2 and Later

```
Clustered index row entry for data page %ld is missing
in index page %ld of table '%S_OBJID' in database
'%S_DBID'; index row contains data page %ld instead. You
can recreate the clustered index to fix this error.
Xactid is (%ld,%d).
```

Version 11.5 and Earlier

```
Fatal attempt to delete clustered index entry for page
%ld - index row contains page %ld - referenced by index
page %ld.
```

Explanation

This error occurs when Adaptive Server fails to delete a clustered index entry because the index entry did not point to the expected page.

In the error message text, the first page number refers to the data page and the last page number refers to the leaf-level index page that points to the data page.

Error 629 can occur when you attempt to delete a row in a table that has a clustered index, but Adaptive Server cannot find a pointer from the leaf page of the clustered index to the data page as expected.

The error is caused by data corruption that occurred during Adaptive Server processing (for example, an operating system panic occurs, causing interruption in disk writes when using UNIX files for Sybase database devices). This may be due to a problem with the server, the operating system, or hardware.

Action

1 Use the procedure in Chapter 2, “How to Find an Object Name from a Page Number” to identify which table and index correspond to the first page number in the error message text.

2 If the object encountering the error is *not* a system table (a system table's object ID is less than 100), continue with step 3.

If the object with the error is a system table, refer to Chapter 2, “How to Fix a Corrupted Index on System Tables” for instructions on how to repair the system table index. Then go to step 6.

3 Determine the index name and the first page in the page chain.

```
1> use database_name
2> go
1> select name, first from sysindexes
```

```
2> where id = object_ID and indid = 1
3> go
```

4 Check that the page linkage is intact.

```
1> dbcc pglinkage(<dbid>,<first page>,0,1,0,1)
2> go
```

If you see a message indicating “End of chain reached”, the page linkage is intact; continue with Step 5. If you do not see such a message, the page linkage is bad; stop here, do *not* drop the clustered index, and contact Sybase Technical Support.

5 To ensure that the information needed to re-create the index is available, run sp_helpindex on the index prior to dropping it.

Drop the index.

Re-create the index. This clears the corruption in most cases.

6 Run dbcc checktable on the table to verify that the corruption is gone. If corruption still exists, call Sybase Technical Support.

Additional information

Refer to “drop index” and “create index” in the *Reference Manual* for information about dropping and re-creating indexes.

Version in which this error is raised

All versions

Error 631

Severity

21

Message text

```
The length of %d passed to delete row routine  for the
row at offset %d is incorrect on the following page:
%S_PAGE.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server attempts to delete a row (via a direct delete or inherently through updating) from an index or data page by specifying the row offset and the row length, and the action fails because the specified values of the offset or row length did not match the actual values.

Error 631 can happen under the following conditions:

- During normal processing, when Adaptive Server tries to delete the row specified by the error message.
- During database recovery. Database recovery occurs:
  - During Adaptive Server start-up
  - When a load database or load transaction command is processed

Some potential causes of Error 631 are:

- Data corruption during normal processing (for example, an operating system panic occurs, causing interruption in disk writes when using UNIX files for Sybase database devices). This may be due to a problem with Adaptive Server, the operating system, or hardware.
- Hardware failure during loading or dumping.

Action

Error 631 is probably the result of a more serious underlying problem, and recovering from this error depends on when the error occurred. Follow the instructions in this section, selecting the correct set depending on whether the error occurred during normal processing or during database recovery.

If the Error Occurred During Normal Processing

1 Use the procedure in Chapter 2, “How to Find an Object Name from a Page Number” to identify which table and index correspond to the page number from the error message text.

2 If the object encountering the error is *not* a system table (a system table's object ID is less than 100), continue with step 3.

If the object with the error is a system table and the index ID is *not* 0, refer to Chapter 2, “How to Fix a Corrupted Index on System Tables” for instructions on how to repair the system table index.

If the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption but you may have to restore from clean backups.

3 For user tables, if the index ID is 0 or 255, continue with step 4.

If the index ID is *not* 0 or 255, translate it into an index name:

```
1> use database_name
2> go
```

```
1> select name from sysindexes
2> where id = object_ID and indid = index_ID
3> go
```

To ensure that the information needed to re-create the index is available, run the sp_helpindex procedure on the index prior to dropping it.

Drop the index.

Re-create the index. This clears the corruption in most cases.

Run dbcc checktable on the table to verify that the corruption is gone.

4 If the index ID is 255, delete the bad data row.

If the index ID is 0, do one of the following:

- Restore the database from clean backups.
- Refer to Chapter 2, “How to Rescue Data from a Corrupted Table”.

Warning! Some data might be lost on this page if you recover your table using bcp or select into (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

Before dumping your database, make sure it works correctly. Run the following commands prior to each dump:

1 dbcc checkdb.

2 dbcc checkalloc or dbcc checkalloc with the fix option. (Refer to Chapter 2, “How to Fix and Prevent Allocation Errors” for information about how to run these commands in multi-user mode and how to prevent spurious allocation errors from dbcc commands.)

If the Error Occurred During Database Recovery

When this error occurs during recovery, the database is marked suspect and is not accessible. Usually, you must load the database from backup. To do this, follow the instructions below:

1 Drop the database. If the drop fails, follow the instructions in Chapter 2, “How to Drop a Database When drop database Fails”.

2 Create a database for load. Make sure the database you create has sizes as least as large as those in sysusages for the original database (and that all other sysusages values match the original values). Refer to “create database” in the *Reference Manual* for more information on create database for load.

3 Load the database from backup. (Refer to "load database" in the *Reference Manual*.)

4 Use the online database command to make the database available for use.

If loading from backups is not feasible, call Sybase Technical Support. Have the following documentation ready:

- Adaptive Server error log
- Text of all error messages
- select * from master..sysusages output

Preventing This Error on Recovery

One of the following items may help to prevent Error 631 from occurring on recovery:

- Use shutdown to shut down Adaptive Server instead of shutdown with nowait after a period of heavy update activity in your databases.

Or:

- checkpoint each database that is being used before shutting down Adaptive Server.

Examine both your operating system error log and the Adaptive Server error log to determine if hardware errors may have affected your database devices. Look for any kernel messages reporting I/O errors and check the hardware error log or diagnostics utilities for I/O errors.

Additional information

Refer to "Developing a Backup and Recovery Plan" in the *System Administration Guide* for complete information about how to safely create, dump, load, and re-create databases.

Version in which this error is raised

All versions

Error 644

Severity

21

Message text

Version 11.5

```
Index row entry for data row id (%ld, %d) is missing
from  index page %ld of index id %d of table '%S_OBJID'
in database '%S_DBID'. Xactid  is (% ld,%d). Drop and
re-create the index.
```

Version 11.0.3 and Earlier

```
Index row entry for index id %d of table '%.*s' in
database  '%.*s' is missing. Drop and re-create the
index. (index page %ld, row %d,  data page %ld)
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when the nonclustered index indicated by “index id” is corrupt. This corruption is detected when a process tries to delete a nonexistent row.

Action

To recover from this error, use the following steps to drop and re-create the index.

1. Record the value of “index page” and “ index id” specified in the 644 error text. Then follow the instructions in Chapter 2, “How to Find an Object Name from a Page Number” to identify which table and index correspond to the index page number. Also note the object ID.
2. If the object with the error is a system table (its object ID is less than 100), refer to Chapter 2, “How to Fix a Corrupted Index on System Tables” for instructions on how to repair the system table index. Then go to step 4.
3. If the object ID is greater than 100, drop and re-create the index, using the table name and the index name obtained in step 1. This clears the corruption in most cases.
4. To verify that all problems have been resolved on this table, run dbcc checktable and dbcc tablealloc on the affected table.

Additional information

If problems persist, this procedure may not be sufficient to clean up the index corruption, and you should contact Sybase Technical Support. Have the following information ready:

- Server version and SWR version level
- Server error log
- Output of dbcc tablealloc and dbcc checktable
- Text of all error messages

Version in which this error is raised

All versions

Error 678

Severity

20

Message text

```
Invalid Buffer Cache Access: Could not open object
'%ld' in database '%d'
```

Explanation

This error is raised when Adaptive Server is unable to locate and read an object into a data cache. The error occurs in the following situations:

- When dbcc checkalloc finds that an allocation structure or extent belongs to a non-existent object.

  The error does not prevent you from continuing other operations. However, each occurrence of this error represents a loss of up to eight pages on disk, or 16K. The pages cannot be used until the error is corrected.

- When the configuration parameter sort page count is set to a very high value, and a stack overflow occurs during subsequent operations.

Action

If the error is raised by dbcc checkalloc, it is due to an Adaptive Server problem. contact Sybase Technical Support for assistance.

If the configuration parameter sort page count is set to a very high value, try using a smaller value for this parameter.

Version in which this error is raised

All versions

Error 691

Severity

20

Message text

```
Encountered invalid logical page '%ld' while accessing
object '%ld' in database '%d'. This is an internal
system error.  Please contact Sybase Technical Support.
```

Explanation

Note This message was introduced in Version 11.0.3 to improve the ability to diagnose 605 object ID mismatch errors. Some errors that previously raised the 605 error are now reported as 691 errors. These errors may be due to hardware problems.

This error occurs when Adaptive Server tries to access an object but requests an invalid page number. There are two possibilities:

- A negative page number is requested.
- Page number 0 is requested. To be more specific, a request is issued for page 0, which is an allocation page, as if it was a data page. Data pages do not use page ID 0. Corrupt page linkages or incorrect entries for the object in sysindexes can cause the server to request page 0.

Error 691 results in a stack trace and breaks your connection to Adaptive Server.

Action

Take the following steps:

1 Check the server error log for other errors that may have been raised prior to the 691 error. Save the entire error log.

2 Identify the object named in the error message using the following isql commands:

```
1> use <database name>
2> go
1> select name from sysobjects where id=<object id>
2> go
```

3 To determine the full extent of the corruption, run the dbcc checkdb and dbcc checkalloc commands (or dbcc checktable and dbcc tablealloc) as soon as possible. On version 11.5 and higher, you can also run dbcc checkstorage.

4 To check if this error is a result of hardware failure, examine your operating system error log and correct hardware problems.

If the 691 error is not a soft error (one that exists only in memory), the problem is severe. Sybase Technical Support may be able to help you recover from the error. However this recovery procedure often results in loss of data; if this is unacceptable, you will probably need to restore from backups.

Additional information

Refer to the writeup for **Error 605** for a discussion of potential causes of hardware error.

Have the following information ready before calling Sybase Technical Support:

- Adaptive Server version and SWR version level
- Server error log
- Operating System error log
- Output of the dbcc checks from Step 3 above.

Version in which this error is raised: 11.0.3 and later

Error 692

Severity: 20

Message text:

```
Uninitialized logical page '%ld' was read while
accessing object '%ld' in database '%d'. Please contact
Sybase Technical Support.
```

Explanation

Note This message was introduced in Version 11.0.3 to improve the ability to diagnose 605 object ID mismatch errors. Some problems that previously raised the 605 error are now reported as 692 errors. These errors may be due to hardware problems.

Error 692 occurs when Adaptive Server tries to read a page that is not formatted (initialized) or is improperly formatted. The server retries the read if the first read attempt failed, and raises the 692 error if the second read attempt also fails.

Error 692 results in a stack trace and breaks your connection to Adaptive Server.

Action

Take the following steps:

1. Check the server error log for other errors that may have been raised prior to the 692 error. Save the entire error log.
2. Obtain page information by running dbcc page as soon as possible, using the database ID and page number shown in the message:

```
1> dbcc page (<db_id>, <page_number>, 0, 1, 1, -1)
2> go
```

Warning! Use the dbcc page command only as directed above.

3. Identify the object named in the error message using the following isql commands:

```
1> use <database name>
2> go
1> select name from sysobjects where id=<object_id>
2> go
```

4. Find information about this object in sysindexes:

```
1> select first, root, doampg, ioampg from sysindexes
2> where name = object_name(<object_id>)
3> go
```

5 To determine the full extent of the corruption, run the dbcc checkdb and dbcc checkalloc commands (or dbcc checktable and dbcc tablealloc) as soon as possible. On version 11.5 and higher, you can also run dbcc checkstorage.

6 To check if this error is a result of hardware failure, examine your operating system error log and correct hardware problems. See Chapter 2, "Checking the Operating System Error Log" for assistance.

If the 692 error is not a soft error (one that exists only in memory), the problem is severe. Sybase Technical Support may be able to help you recover from the error. However, this recovery procedure often results in loss of data; if this is unacceptable, you will probably need to restore from backups.

Additional information

Refer to the **Error 605** writeup for a discussion of potential causes of hardware error. Refer to Chapter 2, "Useful dbcc Commands" for more dbcc information.

Have the following information ready before calling Sybase Technical Support:

- Adaptive Server version and SWR version level
- Server error log
- Operating System error log
- Output of the sysindexes query and dbcc checks from Steps 2, 4, and 5 above.

A database dump and transaction log dumps may also be required.

Version in which this error is raised

11.0.3 and later

Error 693

Severity

20

Message text

```
Multiple copies of logical page '%ld' from database '%d'
reside in more than one cache. This is an internal
```

```
system error. Please contact Sybase Technical Support.
```

Note This message was introduced in Version 11.0.3 to improve the ability to diagnose 605 object ID mismatch errors. Some errors that previously raised the 605 error are now reported as 693 errors.

Explanation

This error occurs when Adaptive Server tries to access an object but is unable to establish the identity of a logical page for the object. Copies of the page are found in more than one cache, and the current session is accessing the wrong version.

Error 693 is caused by an Adaptive Server problem. It results in a stack trace and breaks your connection to Adaptive Server.

Action

1 Check the server error log for other errors that may have been raised prior to the 693 error. Save the entire error log.

2 Obtain cache information by running dbcc page as soon as possible, using the database ID and page number shown in the message:

```
1> dbcc page (<db_id>, <page_number>, 0, 1, 1, -1)
2> go
```

Warning! Use the dbcc page command only as directed above.

3 Contact Sybase Technical Support with the information you collected.

Additional information

Refer to **Error 605** for a complete list of related object mismatch errors.

Have the following information ready before calling Sybase Technical Support:

- Adaptive Server version and SWR version level
- Server error log
- Operating System error log
- Output of dbcc page.

Version in which this error is raised

11.0.3 and later

Error 694

Severity

24

Message text

```
An attempt was made to read logical page '%ld', virtpage
'%ld' from virtual device '%d' for object '%ld' in
database '%d'. The page was not read successfully. You
may have a device problem or an operating system
problem.
```

Note This message was introduced in Version 11.0.3 to improve the ability to diagnose 605 object ID mismatch errors. Some errors that previously raised the 605 error are now reported as 694 errors. These errors may be due to hardware problems.

Explanation

This error occurs when Adaptive Server tries to access an object but is unable to read the page mentioned in the message. Adaptive Server automatically issues a second read request to verify the consistency of the first read attempt. One of the following cases will be true:

- The second I/O request did not take place either, and the page header is unchanged by the I/O operation. Adaptive Server writes the following message into the error log:

```
I/O did not occur, buffer contents are unchanged. Previous page
in buffer = %ld Previous objid = %ld
```

- The second read request returned new information, but that information was incorrect. This indicates an unreliable disk or controller.
- The second read request was successful. The first read request may have failed due to a timing or caching problem on the device.

In all cases, Adaptive Server writes the following message into the error log:

```
Suspect a Device or OS problem, %s
```

and raises Error 694, resulting in a stack trace and breaking your connection to Adaptive Server.

This error indicates hardware or operating system problems.

Action

- Save the entire Adaptive Server error log from the last boot to the time of the error.
- If the second read was successful, check whether dbcc checkalloc or dbcc checkdb is running concurrently with other activities on the server. Error 694 may occur when too many I/O's are requested from the server.
- If the second I/O attempt failed or if it returned a bad read, check if the device in question provides a caching mechanism. If so, turn off caching to see if this resolves the 694 error.

- If the error persists, determine what device(s) are responsible for the errors. Examine your operating system error log and correct hardware problems. See Chapter 2, “Checking the Operating System Error Log” for assistance.

Additional information

Refer to the writeup for **Error 605** for a discussion of potential causes of hardware error.

Version in which this error is raised

11.0.3 and later

Error 695

Severity

```
20
```

Message text

```
An attempt was made to read logical page '%ld' for
object '%ld' in database '%d' from cache '%.*s'. Wrong
logical page '%ld' was brought into cache.
```

Note This message was introduced in Version 11.0.3 to improve the ability to diagnose 605 object ID mismatch errors. Some problems that previously raised the 605 error are now reported as 695 errors. These errors may be due to hardware problems.

Explanation

Error 695 occurs when Adaptive Server tries to read a page from disk (the first logical page in the message), but the page number in the header of the returned page does not match the requested page. This means that the page is corrupt on disk, and has been overwritten by another page. The error may be due to an Adaptive Server problem but could also be caused by problems such as overlapping partitions or hardware errors.

Error 695 results in a stack trace and breaks your connection to Adaptive Server. Additional errors may be reported on the console and in the server error log prior to this error.

Action

1 Check the server error log for other errors that may have been raised prior to the 695 error. Save the entire error log.

2 Obtain page information by running dbcc page as soon as possible, using the database ID and page number shown in the message:

```
1> dbcc page (<db_id>, <page_number>, 0, 1, 1, -1)
```

```
2> go
```

Warning! Use the dbcc page command only as directed above.

3 To determine the scope of the problem, locate the device on which the database resides, and check if any other databases use the same device. Refer to Chapter 2, “How to Determine Which Physical Devices a Database is On”.

4 Check that the device is partitioned correctly. See Chapter 2, “Correct Use of Raw Partitions”.

5 Run dbcc log to find the history of transactions against this page.

```
1> dbcc log (<dbid>, 0, <page_no>, 0, 0, -1)
2> go
```

Warning! Use the dbcc log command only as directed above.

6 To determine the full extent of the corruption, run the dbcc checkdb and dbcc checkalloc commands (or dbcc checktable and dbcc tablealloc) as soon as possible. On version 11.5 and higher, you can also run dbcc checkstorage.

7 To check if this error is a result of hardware failure, examine your operating system error log and correct hardware problems. See Chapter 2, “Checking the Operating System Error Log” for assistance. Also see **Error 605** for a discussion of potential causes of hardware error.

Sybase Technical Support may be able to help you recover from the error. However this recovery procedure often results in loss of data; if this is unacceptable, you will probably need to restore from backups.

Additional information

See Chapter 2, “Useful dbcc Commands” for more dbcc information.

Have the following information ready before calling Sybase Technical Support:

- Adaptive Server version and SWR version level
- Server error log
- Operating System error log
- Output of the dbcc checks from Steps 2, 5, and 6 above.

Version in which this error is raised

11.0.3 and later

Error 696

Severity

21

Message text

```
An attempt was made to fetch logical page '%ld' in
tempdb from cache '%.*s'. Page belongs to object '%ld'
and not to object '%ld'. Restart of SQL Server will
clear the error. Please contact your System
Administrator for help.
```

Note This message was introduced in Version 11.0.3 to improve the ability to diagnose 605 object ID mismatch errors. Some errors that previously raised the 605 error are now reported as 696 errors.

Explanation

This error occurs when Adaptive Server discovers page allocation corruption. Adaptive Server tries to access a particular object but discovers a page in the object's page chain whose object ID is different than that of the object being accessed.

This error is similar to the 605 error. However, Error 696 always occurs in the context of the tempdb database. Since the object ID mismatch is detected in cache, this is a transient (soft) error.

Action

Shut down and restart Adaptive Server to clear the cache problem. If Error 696 occurs again, call Sybase Technical Support.

Additional information

Have the following information ready before calling Sybase Technical Support:

- Adaptive Server version and SWR version level
- Text of all error messages
- Text of the query which raises the error.

Version in which this error is raised

11.0.3 and later

Error 697

Severity

20

Message text

```
An attempt was made to fetch logical page '%ld' for
object '%ld' in database '%d' from cache '%.*s'. Wrong
```

```
logical page '%ld' was found in cache.
```

Note This message was introduced in Version 11.0.3 to improve the ability to diagnose 605 object ID mismatch errors. Some errors that previously raised the 605 error are now reported as 697 errors. These errors can be due to hardware problems.

Explanation

This error occurs when Adaptive Server tries to access an object in cache but arrives at an invalid page. The page number is correct on disk, but it does not match the page number of the page in memory.

Error 697 results in a stack trace and breaks your connection to Adaptive Server.

Action

Since the problem occurs in cache, it is a soft error. Shut down and restart Adaptive Server to clear the memory corruption.

Version in which this error is raised

11.0.3 and later

Memory Manager Errors

This section contains error messages for the Adaptive Server Memory Manager.

Error 701

Severity

```
19
```

Message text

```
There is not enough procedure cache to run  this
procedure, trigger, or SQL batch. Retry later, or ask
your SA to reconfigure  SQL Server with more procedure
cache.
```

Explanation

This error occurs when there is not enough procedure cache available to execute a stored procedure, trigger, or batch of one or more SQL statements. This can happen when loading a query plan into procedure cache or when resolving or compiling a procedure, trigger, or batch.

The amount of available procedure cache is based on Adaptive Server use, so this error may occur intermittently as procedure cache is used and released. Also, one or more procedures can repeatedly fail with this error while others complete normally. This is because query plans vary greatly in size. Therefore, while there may not be sufficient space to load a large query plan, several smaller plans may fit.

Action

Correct this error by increasing the size of procedure cache. You can increase the amount of total procedure cache in three ways:

- Use the total memory configuration parameter to increase the total amount of memory requested by Adaptive Server.
- Use the procedure cache percent configuration parameter to increase the percentage of total cache space dedicated to procedure cache.
- Decrease the amount of memory required for other resources; for example, unneeded user connections.

Increasing the total memory configuration parameter for Adaptive Server is the most straightforward method. It increases both the procedure and data caches, although it can waste memory.

Note On some operating systems, especially VMS and most UNIX systems, be sure to make any necessary operating system memory resource adjustments. For example, verify the current kernel value for the maximum size of a shared memory segment, usually *SHMMAX*, and make sure you have adequate memory and swap space on the system for the additional memory configured for the Server. For more information, refer to a) your operating system documentation, and b) the Adaptive Server installation and configuration guide.

If your memory resources are limited, you can increase the procedure cache percent configuration parameter for the Adaptive Server without changing the value of the total memory configuration parameter. This shifts space from the data cache to the procedure cache and may result in performance degradation or other problems if not enough data cache remains.

You can also increase the size of the procedure cache without substantially changing the amount of data cache. This requires some combination of increasing both the total memory and procedure cache percent configuration parameters in such a way that the majority of the new memory goes to procedure cache. Although this method requires more planning, it allows you to control where the additional memory goes. Refer to the *Performance and Tuning Guide* and "Configuring Memory" in the *System Administration Guide* for more information about configuring Adaptive Server memory.

Version in which this error is raised

All versions

Error 702

Severity

20

Message text

```
Memory request for %d bytes exceeds the size of single
page of %d bytes.
```

Explanation

This error occurs when you exceed any of the following limits:

- 128 search conditions or join operations in a SQL statement (version 11.0.x).

  A search condition sets the conditions in a where or having clause. For more information about, and examples, of search conditions, refer to "Search Conditions" in the *Reference Manual*.

  A join operator compares two or more tables or views. For more information about and examples of join operators, refer to the *Reference Manual*.

- 128 columns in a Data Workbench update or delete table.

  When modifying database tables, Data Workbench sends every updated table column as a search condition to Adaptive Server. Therefore, any Data Workbench table that has more than 128 columns and is updated though the "Modify Data" option generates Error 702.

- The total width of the columns in a temporary table during a union exceeds the allowed row width (rows consist of 1962 characters; allow 2 characters of row overhead for APL tables, for a maximum usable row size of 1960. For DOL tables, subtract two characters per varchar column in determining usable row size).

After displaying Error 702, Adaptive Server terminates the current process.

Action

If your query exceeds the limit of 128 search conditions or join operations, rewrite the query so that the limit of 128 statements is not exceeded (version 11.0.x).

If the total width of the columns in a temporary table during a union exceeds the allowed row width, rewrite your query so that the limit of 1962 bytes is not exceeded.

Data Workbench

If a table has more than 128 columns, avoid updating or deleting rows from that table via the "Modify Data" option of Data Workbench. Instead, run queries using SQL statements and take into consideration the primary, unique keys on that table when you define the rows that are to be updated or deleted.

For example, if a table has a unique key on column *column1*, run the following query in order to delete the row in the table that contains the unique key *unique_key1*:

```
1> begin transaction
2> delete table_name
3> where column1 = unique_key1
4> go
```

If the key is unique in the table, only one row will be deleted by the above query. If only one row is being returned by the above query, commit the transaction with the following query:

```
1> commit transaction
2> go
```

Otherwise, roll back the transaction:

```
1> rollback transaction
2> go
```

If you are not sure if the key *unique_key1* is unique in *column1*, you can check by running the following query:

```
1> select * from table_name
2> where column1 = unique_key1
3> go
```

Similarly, if the table has a unique index on columns *column1* and *column2*, you can delete a row in that table by running the following query:

```
1> delete table_name
2> where column1 = unique_key1
3> and column2 = unique_key2
4> go
```

Version in which this error is raised	All versions

Error 703

Severity	17
Message text	`You cannot run this procedure, trigger, or SQL batch because it requires more than %ld pages of memory. Break it up into shorter queries, if possible.`

Explanation

This error occurs when a stored procedure or trigger cannot be executed because it requires more memory than is allowed for execution.

Two different stages of execution can trigger this error: "resolution," in which the query tree is built, or " compilation," in which the query plan is generated. If a query has been executed successfully but later fails with this error, it means that the query tree used to be less than the allowed memory limit but has since grown. If a query encounters this error every time it is executed, then it is too complex to be executed in its present form without exceeding the memory limit. Refer to the *Performance and Tuning Guide* for detailed information about query trees and query plans.

Action

Note Errors 701 and 703 are very similar. Error 703 is raised when the procedure requires multiple process headers (control structures) but there is not enough memory to allocate the next header; Error 701 is raised when there is not enough memory to extend the current header which already has some pages allocated. Refer to **Error 701** for additional troubleshooting information.

If the procedure or trigger causes this error every time you try to execute it, divide it into smaller pieces. This division process varies greatly with the type of SQL statement.

If the object has successfully executed before, the 703 error probably occurred because the query tree grew beyond the specified page limit. This growth occurs each time the query tree is re-resolved. Once you have encountered the 703 error in this situation, you can drop and re-create the procedure or trigger in order to shrink the query tree to a legal size, or you can break the object up into smaller modules. Restarting Adaptive Server has no effect on the size of the query tree because query trees are stored on disk.

If this error occurs frequently on an object, you can either periodically drop and re-create the object as part of regular database maintenance, or break it into smaller modules.

Version in which this error is raised: All versions

Error 706

Severity: 20

Message text:

```
Process %d tried to remove PROC_HDR 0x%lx  that it does
not hold in Pss.
```

Explanation: The Memory Manager allocates, deallocates, and manages memory for Adaptive Server. It manages an array of structures in the procedure header, each of which represents a physical page of memory in procedure cache and the current byte allocation of that page.

As part of its memory deallocation process, Adaptive Server stops tracking procedure headers when they are no longer needed and attempts to remove them from procedure cache. Error 706 occurs when Adaptive Server fails to deallocate a procedure header.

Error 706 is caused by memory corruption or an Adaptive Server problem.

Action: Since procedure headers are stored in procedure cache, restarting Adaptive Server should clear Error 706. If it does not clear the 706 error or if the 706 error occurs again, contact Sybase Technical Support.

Version in which this error is raised: All versions

Error 707

Severity: 20

Message text:

```
System error detected during attempt to free memory  at
address 0x%lx. Please consult the SQL Server error log
for more details.
```

Explanation: The Memory Manager allocates, deallocates, and manages memory for Adaptive Server. It manages an array of structures in the procedure header, each of which represents a physical page of memory in procedure cache and the current byte allocation of that page.

As part of its memory deallocation process, Adaptive Server tries to release the pages of memory allocated to a procedure header when they are no longer needed. When Adaptive Server is unable to release that section of memory, Error 707 occurs.

Error 707 is caused by memory corruption or an Adaptive Server problem.

Action

Since procedure headers are stored in procedure cache, restarting Adaptive Server should clear Error 707. If it does not clear the 707 error or if the 707 error occurs again, contact Sybase Technical Support.

Version in which this error is raised

All versions

Error 709

Severity

17

Message text

```
There is insufficient system memory to continue login
process for spid %d.
```

Explanation

The Memory Manager allocates and deallocates memory for Adaptive Server processes and manages memory requirements for the system. For each process connecting to the server, the Memory Manager allocates a memory structure called a procedure header and ensures that there is enough memory to dedicate to the process.

Error 709 is raised when a user process attempts to log in to the server, but there is not enough memory available to complete the login.

Action

Correct this error by increasing the amount of available memory. You can do this in three ways:

- Use the total memory configuration parameter to increase the total amount of memory requested by Adaptive Server.
- Use the procedure cache percent configuration parameter to increase the percentage of total cache space dedicated to procedure cache.
- Decrease the amount of memory required for other resources; for example, unneeded user connections.

Increasing the total memory configuration parameter for Adaptive Server is the most straightforward method. It increases both the procedure and data caches, although it can waste memory.

Since total memory and procedure cache percent are static parameters, you must restart Adaptive Server after changing these parameters.

Additional information

Check the value of the number of user connections configuration parameter. Since there is memory overhead associated with each user connection, setting this parameter too high can reduce the amount of space available for the data and procedure caches, and contribute to 709 errors.

Version in which this error is raised

All versions

Buffer Manager Errors

This section contains error messages for the Adaptive Server Buffer Manager.

Error 803

Severity

20

Message text

```
Unable to place buffer '0x%lx' from cache '%.*s' holding
logical page '%ld' in sdes for object '%.*s' - either
there is no room in sdes or buffer already in requested
slot.
```

Note "sdes" is an abbreviation for "session descriptors." Refer to Error 603 for information about session descriptors.

Explanation

A session descriptor is an internal data structure in Adaptive Server which contains information about a table. All open objects in Adaptive Server require a descriptor structure. Whenever a connection to Adaptive Server attempts to open a table, it gets a session descriptor.

To ensure that a buffer read by the session remains in the buffer cache and does not age out, Adaptive Server "keeps" the buffer. The session descriptor contains eight slots where buffers that are "kept" by the session are stored.

Error 803 occurs during recovery or regular Adaptive Server operation when one of the following is true:

- Adaptive Server runs out of session descriptors
- A buffer was already in the session descriptor's requested slot

Note The utility sybload also has an Error 803, which can occur during installation or upgrade. The last part of this writeup details the sybload 803 error.

Action

When an 803 error occurs:

- Look for other errors in the Adaptive Server error log to indicate the specific source of the problem (for example, not enough locks, not enough open objects) and clear those errors first.
- Restart Adaptive Server so that the configuration change will take effect. Restarting will also clear the buffer cache and the 803 error.

If problems persist, contact Technical Support for assistance. Before contacting Technical Support, be prepared to fax or provide the complete text of all error messages.

Examine your operating system error log file as well as the Adaptive Server error log to determine if hardware errors may have caused the problem.

sybload Error 803

Error Message Text

```
Please inform Sybase Customer Service that there may be
a problem with this tape and cite error condition 803.
```

sybload Error 803 can occur when unloading software because of one of the following:

- When loading Sybase software from tape, the operator made one of the following errors:
  - The nonrewinding device was not specified when loading Sybase software from media with sybload.
  - When loading the sybload utility for a remote installation, the operator typed the commands for a local installation instead of those provided for a remote installation.

    A local installation is one in which the Sybase installation directory and tape drive are on the same machine. Remote installations are the ones where the Sybase installation directory and tape drive are on different machines.
- The version of the Adaptive Server to be installed requires a different version of the operating system.

Action for sybload *Error 803*

For more information, refer to the Adaptive Server installation and configuration guide and the *Release Bulletin*.

Use appropriate procedures from the following:

- Check that the Adaptive Server platform and version you are trying to install or upgrade is compatible with your operating system version.
- Check that you specify a nonrewinding device when instructed to do so.

- If you are loading Sybase software for a remote installation, check that you typed in the correct commands for loading the sybload utility from tape. Remote load commands are different from the commands for a local installation.

Version in which this error is raised

All versions

Error 804

Severity

20

Message text

```
Unable to find buffer 0x%lx holding logical  page %ld
in sdes 0x%lx kept buffer pool for object '%.*s'.
```

Note "sdes" is an abbreviation for "session descriptors." Refer to Error 603 for information about session descriptors.

Explanation

A session descriptor is an internal data structure in Adaptive Server which contains information about a table. All open objects in Adaptive Server require a descriptor structure. Whenever a connection to Adaptive Server attempts to open a table, it gets a session descriptor.

To ensure that a buffer read by the session remains in the buffer cache and does not age out, Adaptive Server "keeps" the buffer. The session descriptor contains eight slots where buffers that are "kept" by the session are stored.

Error 804 occurs when Adaptive Server cannot find the pointer to a buffer header in a session descriptor. This error can be transient.

Action

When an 804 error occurs, restart Adaptive Server as soon as possible. If additional 804 errors occur, call Sybase Technical Support.

Additional information

When calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log
- Text of all error messages

Version in which this error is raised

All versions

Error 806

Severity

21

Message text

```
Could not find virtual page for logical page  %ld in
database '%S_DBID'.
```

Note This error may be caused by a hardware problem.

Explanation

A virtual page is a page within a Sybase device. A logical page is a page in an Adaptive Server database. There is a one-to-one correspondence between these two types of pages.

Error 806 occurs when Adaptive Server fails to convert a logical page number to a virtual page number. Depending on what caused the error, it can be serious or transient.

If Error 806 is accompanied by the message "Page %d is not in the range of pages for database id %d," it means you used dbcc page with an invalid parameter. This is *not* a serious problem.

If Error 806 occurs on recovery, it may be transient or serious (see information under "Action" for specifics).

Error 806 can occur during normal processing, such as creating an index or running a stored procedure. In this case, the error is probably caused by corruption or a problem with Adaptive Server and it is a serious error.

Action

If Error 806 specifies tempdb in the message output, restart Adaptive Server. Since tempdb is rebuilt each time Adaptive Server is restarted, this may clear the error. If the error occurs again (on tempdb), call Sybase Technical Support.

During Recovery

If Error 806 occurs on recovery, the database will be marked suspect. If the error is transient, resetting the suspect status will solve the problem.

To resolve this problem:

1 Bypass recovery by starting Adaptive Server with status -32768.

2 Run dbcc checkdb and dbcc checkalloc on the database listed in the error message output.

3 Set status back to 0.

4 Shut down Adaptive Server:

```
1> shutdown with nowait
```

```
2> go
```

5 Try recovery again.

If Error 806 occurs again, call Sybase Technical Support. Technical Support may be able to help you recover from this error if the corruption is not too widespread. However, recovery from clean backups might be necessary.

During Normal Processing

Run dbcc checkdb and dbcc checkalloc on the database listed in the error message output.

Call Sybase Technical Support. Technical Support may be able to help you recover from this error if the corruption is not too widespread. However, recovery from clean backups might be necessary.

Additional information

Refer to "Checking Database Consistency" in the *System Administration Guide* for information about dbcc commands.

Before contacting Technical Support, be prepared to provide:

- Server version and SWR version level
- Complete text of all error messages
- Server error log output
- Operating system error log output
- dbcc checkdb, dbcc checkalloc, and dbcc page output

Version in which this error is raised

All versions

Error 813

Severity

```
20
```

Message text

```
Logical page '%ld' in database '%S_DBID',  cache '%.*s'
is already hashed.
```

Note This error may be caused by a hardware problem.

Explanation

This error may be serious, especially if it occurs on a table's data page. It means that a page is currently in the data cache and is in use by a table or index but is not marked as allocated. An attempt to allocate it causes Error 813.

The same page could be allocated again after it is removed from the data cache, resulting in a loss of whatever data resides on the page.

Warning! Pages encountering Error 813 will not be included in a database dump. This is because database dumps are performed by reading allocation pages and not by traversing page chains. Therefore, this error should be corrected before dumping the database.

After the page is removed from the data cache, further attempts to access this page may raise 2500 series errors (allocation errors) when running dbcc checkalloc, dbcc tablealloc, or dbcc indexalloc.

Action

To clear the error, do the following:

1 Execute the checkpoint command on the database specified in the error message. This will cause all modified pages in the data cache to be flushed to disk.

2 Run dbcc checkalloc or dbcc checkalloc with the fix option on the database. As a result, you might get allocation errors because the page displayed in the error message will be linked but not allocated. Refer to “dbcc” in the *Reference Manual* for information about running these commands.

3 If you do get allocation errors:

- Check your hardware error log or diagnostics utilities for I/O errors.
- Check your operating system error log file and the Adaptive Server error log to determine if hardware errors may have caused the corruption. (Look for messages reporting I/O errors.)
- Refer to the writeups in this manual.

4 If you still have 813 errors after fixing any other errors reported during step 2, shut down and restart Adaptive Server. This will probably clear Error 813.

If errors still occur, call Sybase Technical Support. They might be able to help you recover from this error if the corruption is not too widespread. However, recovery from backups might be necessary.

Using dump transaction with no_log

Using the dump transaction with no_log command can result in an 813 error. Therefore, do not use dump transaction with no_log unless it is absolutely necessary (when dump transaction with truncate_only will not truncate the log). Try using dump transaction with truncate_only first and use the no_log option only as a last resort. Refer to Error 1105 for details.

Check for any occurrences of dump transaction with no_log in any of your scripts or procedures and replace them with dump transaction with truncate_only.

Additional information

For more information on the 2500 series errors, refer to the information in this section and in Chapter 2, "How to Fix and Prevent Allocation Errors".

Version in which this error is raised

All versions

Error 820

Severity

21

Message text

```
Attempt to dirty non-log buffer %S_BUF which is in I/O.
```

Explanation

Adaptive Server uses buffers to manage physical pages while they are in memory. When a page is read into memory from the disk, a buffer header is assigned to it. The header tracks the usage and contents of the page image.

The buffers themselves are grouped together to form Memory Address Space Segments (MASSes). A buffer cache consists of MASSes linked in MRU/LRU (most recently used/least recently used) chains.

A buffer that was changed while in cache and has not yet been written to disk is known as a dirty buffer.

Error 820 is raised when the server detects that an invalid operation has been performed on a buffer or MASS. The error is raised in the following states:

State	Meaning
1	Attempted to mark as dirty a MASS that is currently being written.
2	Object ID mismatch detected when linking a buffer to the chain.
3	MASS found to belong to a different database. This state is only raised by diagserver.

State	Meaning
4	Buffer is being changed but the page residing in the buffer does not belong to the transaction log as expected.
5	Buffer pinning to a cache is invalid with respect to transaction state.
6	Attempted to pin a buffer to a cache while the buffer is pinned to a different cache.

Error 820 is due to an Adaptive Server problem. Some scenarios in which this error may be raised include:

- Parallel select into with union all.
- drop role when 15 or more databases exist on the server.

Action

Call Sybase Technical Support.

Additional information

Before calling Technical Support, have the following information available:

- Adaptive Server version and SWR rollup level
- Server error log
- Operating system error log

Version in which this error is raised

All versions

Error 821

Severity

20

Message text

```
Attempt to unhash buffer at 0x%lx with a  buffer pageno
of %ld and database id %ld with HASHED status set failed
- buffer  was not found. %S_PAGE.
```

Note This error may be caused by a hardware problem.

Explanation

Every data or index page in the data cache has a corresponding buffer pointing to it, controlling its stay in the cache. As these pages are read and modified, they are scheduled to be written to disk and removed from the data cache. Error 821 occurs when Adaptive Server attempts to remove a nonexistent or corrupted page from the data cache.

Some potential causes of this error are:

- Overlapping partitions (refer to Chapter 2, "Correct Use of Raw Partitions").
- Hardware failure.
- An Adaptive Server problem (such as can occur when using the dump transaction with no_log command after a period of heavy user activity in the database).

This error can occur during database recovery as well as during normal operation of Adaptive Server.

Action

1 Check the Adaptive Server error log to determine whether there are other indications of hardware problems, such as kernel messages reporting I/O errors. If problems exist, solve them.

2 Check the operating system error log or run diagnostic utilities to check for I/O errors. If problems exist, solve them.

Refer to Chapter 2, "Checking the Operating System Error Log" for more information.

3 Once the actual cause of corruption has been removed, shut down and restart Adaptive Server to clear the 821 error. If restarting is not feasible and if you have dial-in facilities, Sybase Technical Support may be able to help resolve the immediate problem.

Using dump transaction with no_log

Using the dump transaction with no_log command can result in an 821 error. Therefore, do not use dump transaction with no_log unless it is absolutely necessary (when dump transaction with truncate_only will not truncate the log). Try using dump transaction with truncate_only first and use the no_log option only as a last resort. Refer to Error 1105 for details.

Check for any occurrences of dump transaction with no_log in any of your scripts or procedures and replace them with dump transaction with truncate_only.

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log
- Operating system error log
- Output of dbcc checkalloc and dbcc checkdb

- Text of all error messages

Version in which this error is raised

All versions

Error 822

Severity

20

Message text

```
Could not start I/O for request BLKIO flags  = 0x%lx,
size = %ld, errcode = 0x%lx, %S_BUF.
```

Note This error may be caused by a hardware problem, usually a missing/offline device.

Explanation

This error occurs when a read or write operation was requested and Adaptive Server encountered an error while attempting to initiate the request. Error 822 can occur for any of the following reasons:

- A database device is offline
- A database device or file has been removed or renamed
- A database device or file is inaccessible to Adaptive Server for any other reason, such as insufficient permissions or a device number higher than the number of devices configured for Adaptive Server (such a device will not be activated at startup).

The information in the error message refers to an internal Sybase structure and does not help determine which database device or file is involved. Other error messages appearing in the Adaptive Server error log just before the 822 error usually indicate the nature of the problem and the device involved. Refer to "Additional Information" for examples of this type of error message.

Action

1 Examine the availability and condition of the device involved using your standard operating system procedures, and make sure the device is accessible.

2 Adaptive Server will not be able to recover the database upon restarting because an 822 error will mark the database suspect. If you know that the database was marked suspect because the device was unavailable, reset the status of the database to allow recovery to continue.

Warning! Do not use these procedures if other errors in the error log near the 822 error message suggest that the database may have been marked suspect for some other reason.

Reset the suspect status using one of the methods supplied in Chapter 2, "How to Reset a Database's "suspect" Status". After you reset the status, and start Adaptive Server, if the database is again marked suspect, examine the Adaptive Server error log for indications of other errors.

Additional information

Here are two examples of the 822 error and associated messages from the Adaptive Server error log. Examining the error messages just before the error usually helps determine the type of problem.

Example 1: Permission Problems

In this example, the permissions on *disk1* (virtual device 6) were incorrect, so Adaptive Server was unable to access this device:

```
kernel: initializing virtual device 6, "/work/disk1"
kernel: dopen: open "/work/disk1", Permission denied
kernel: dinit: failed to open primary device /work/disk1 for vdn6
kernel: udstartio: vdn 6 has not been set up
server: Error: 822, Severity: 20, State: 3
server: Could not start I/O for request BLKIO ....
```

To correct an error like this:

1 Change the ownership to user "sybase". Make sure that this user has read and write permissions on the device. In our example, check the user's permission on */work/disk1*.

2 Become the "sybase" user.

3 Start Adaptive Server.

Example 2: File Missing

In this example, the UNIX file corresponding to *disk1* (virtual device 6) did not exist when Adaptive Server started, causing access to the device to fail:

```
kernel: initializing virtual device 6, "/work/disk1"
kernel: dopen: open "/work/disk1", No such file or directory
kernel: dinit: failed to open primary device /work/disk1 for vdn6
```

```
kernel: udstartio: vdn 6 has not been set up
server: Error: 822, Severity: 20, State: 3
server: Could not start I/O for request BLKIO ....
```

To correct an error like this, make sure that the virtual device in the error message exists, and correct whatever error caused Adaptive Server not to find it. (For example, it may have been renamed or moved while Adaptive Server was shut down.)

If the virtual device no longer exists, you will have to restore from known clean backups. Merely creating an empty device with the right name will not solve this problem.

Version in which this error is raised

All versions

Error 823

Severity

24

Message text

```
I/O error detected during %S_MSG for %S_BUF.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server encounters an I/O error on a read or write request made to a Sybase device. It usually means you have disk problems. The parameters in the error message refer to internal Sybase structures and do not often help determine which device is involved. However, additional kernel messages in Adaptive Server's error log recorded before the 823 error should indicate which device is involved.

Causes of the 823 error can include the following:

- If the sum of vstart and size exceeds the device size. This can happen on AIX Systems but is also possible on other UNIX platforms if sysusages or sysdevices have been manually altered.
- When using UNIX files as devices, if the filesystem is full but the device appears to have space as tracked by sysusages and sysdevices. This raises the 823 error during create database or alter database.

Action

Check the accessibility and condition of the device in question. Once you have identified and corrected the problem, execute the dbcc checkdb and dbcc checkalloc commands to ensure that no other damage was caused by the bad device.

SunOS 5.x (Sun Solaris 2.x)

The 823 error can also appear if the Sun operating system is not configured correctly for asynchronous I/O.

Adaptive Servers running on Sun machines can use asynchronous I/O on raw partitions without any special kernel enhancements such as the MtXinu DBM or Sun DBE. An Adaptive Server running on a SunOS 4.1 or later operating system attempts to do asynchronous I/O on raw partitions. Successful use of asynchronous I/O requires that the UNIX kernel be configured to support it. Two options must be included in the Sun operating system kernel configuration file when the kernel is built. These are:

```
options  LWP    # kernel threads
options  ASYNCHIO        # asynch I/O (requires LWP)
```

You can find examples of how to set these options in the original GENERIC configuration file usually found in */sys/{sun3, sun4}/conf*. The exact location of these files may vary. Do not comment out these options.

An improperly configured kernel causes Adaptive Server to fail the first time it attempts to do I/O to a raw partition. For example, if your master device is on a raw partition, and you attempt to start Adaptive Server with that master device, Adaptive Server displays the following message:

```
kernel: read error on virtual disk 0 block 28
kernel: invalid argument
server: Failure to open master database for the first time
server: Error 823, Severity 24, State 1
```

If the master device is located on a UNIX file for which standard disk I/O rather than asynchronous I/O is used, Adaptive Server starts even if the Sun kernel is not configured for asynchronous I/O. However, if you attempt a disk init command on a raw partition, it fails with a 5115 error. Refer to Error 5115 for more information.

AIX Systems Only

Error 823 may occur if the disk device is run under the Logical Volume Manager (LVM) , vstart is equal to 2 and the size is not shortened by 2. vstart needs to be equal to 2 because the first AIX block (1 AIX block equals 2 Sybase pages) has to be available for the LVM control block (LVCB).

Version in which this error is raised: All versions

Error 834

Severity: 20

Message text: `Illegal attempt to clean buffer: %S_BUF.`

Explanation

The Adaptive Server Cache Manager works in conjunction with the Buffer Manager to manage physical disk pages while they are in memory. When a page is read into memory from the disk, a buffer header is assigned to it. The header tracks the usage and the contents of the page image. The page and the buffer header together define a buffer.

The buffers themselves are grouped together to form Memory Address Space Segments (MASSes). A buffer cache consists of MASSes linked in MRU/LRU (most recently used/least recently used) chains. The Cache Manager manages MASSes and buffer caches.

When a buffer or MASS is no longer needed (for example during a drop database, drop table, drop index, or deallocation of pages), Adaptive Server performs some cleanup tasks and releases the resources taken up by the buffer or MASS. Before these actions are carried out, Adaptive Server does a number of checks to make sure the buffer or MASS is no longer in use. Error 834 occurs when one of these checks fails. It is caused by an Adaptive Server problem.

Error 834 occurs with the following states:

State	Meaning
1	If the MASS is actively being used or if Adaptive Server is in the process of writing out the MASS when Adaptive Server tries to clean it up, Error 834 occurs with State 1. This state is only raised by diagserver.
2	If the buffer is still hashed or has not been cleaned up when Adaptive Server is ready to release it, Error 834 occurs with State 2. This state is only raised by diagserver.
3	If the buffer is actively being used or if Adaptive Server is in the process of writing out the buffer when Adaptive Server is ready to release it, Error 834 occurs with State 3. This state is only raised by diagserver.
4	If the buffer belonging to a database has been set aside for later use and Adaptive Server tries to clean it out of a cache, Error 834 occurs with State 4.

5	If a buffer belonging to an object has been set aside for later use and Adaptive Server tries to clean it out of a cache, Error 834 occurs with State 5.

Action

State 4, LTMs Running

834 errors with State 4 can occur if the Log Transfer Manager (LTM), a Replication Server component, is running during server recovery. The LTM process attempts to log in before the database is fully recovered. To avoid this error, shut down the LTM before restarting the server. Restart LTM only after all the databases have been recovered.

All Other States

Shut down and restart Adaptive Server to clear the buffer structure.

If Error 834 occurs again, call Sybase Technical Support.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level
- Text of all error messages.

Version in which this error is raised

All versions

Error 835

Severity

```
20
```

Message text

```
Illegal attempt to change contents of buffer:  %S_BUF.
```

Explanation

The Adaptive Server Cache Manager works in conjunction with the Buffer Manager to manage physical disk pages while they are in memory. When a page is read into memory from the disk, a buffer header is assigned to it. The header tracks the usage and the contents of the page image. The page and the buffer header together define a buffer.

When a client process completes, Adaptive Server performs some cleanup tasks such as closing the buffers and releasing the resources taken up by the buffers. If the client process terminates abnormally, however (for example if the process is killed during execution), Adaptive Server may be unable to carry out the appropriate cleanup, buffers are left open, and Error 835 is raised. The error can also occur due to an Adaptive Server problem, when a page being updated is deallocated before the modification can complete.

This error is only raised by diagserver.

Action

If the problem recurs, shut down and restart Adaptive Server to clear the buffer structure.

If Error 835 continues to occur, call Sybase Technical Support.

Version in which this error is raised

All versions

Error 840

Severity

17

Message text

```
Device '%.*s' (with physical name '%.*s',  and virtual
device number %d) has not been correctly activated at
startup  time. Please contact a user with System
Administrator (SA) role.
```

Note This error may be caused by a hardware problem.

Explanation

This error can occur when Adaptive Server is unable to access a device during startup.

This error can occur for any of the following reasons:

- A database device is offline.
- A database device or file has been removed or renamed.
- A database device or file is inaccessible to Adaptive Server for any other reason, such as insufficient permissions.

Action

1 Examine the availability and condition of the device involved using your standard operating system procedures, and make sure the device is accessible.

2 Adaptive Server will not be able to recover any databases that have space on the device upon restarting because an 840 error will mark the databases suspect. If you know that the databases were marked suspect because the device was unavailable, resolve the problem with the device and reset the status of the databases to allow recovery to continue.

Warning! Do not use these procedures if other errors in the error log near the 840 error message suggest that the databases may have been marked suspect for some other reason.

Reset the suspect status using one of the methods supplied in Chapter 2, "How to Reset a Database's "suspect" Status".

After you reset the status, execute the dbcc checkdb and dbcc checkalloc commands for the affected databases to ensure that no other damage was caused by the device being unavailable. If other errors occur, refer to the writeups for those errors in this manual. If errors still persist, call Sybase Technical Support.

Version in which this error is raised

All versions

Error 842

Severity

20

Message text

```
Attempt to set bufclaims in PSS structure for process
%d to negative value while unclaiming cache.
```

Explanation

Adaptive Server manages physical disk pages while they are in memory by maintaining buffer pools of logical pages in cache. It allocates buffers to server tasks as needed, keeps track of the buffers assigned to caches and tasks, and releases (unclaims) buffers that are no longer needed by a task. It also accesses the Process Status Structure (PSS), which is a memory structure associated with each process describing the state of the process.

Error 842 occurs when a buffer is released, and the task is found to have a negative buffer count as a result. The error is sometimes seen during sort operations, for example during index creation.

Error 842 breaks your connection to Adaptive Server.

Action

Reconnect to the server. If the error persists:

1 Configure the server to obtain a shared memory dump for future occurrences of Error 842; and

2 Contact Sybase Technical Support after collecting the memory dump and related error data.

To configure the server to dump shared memory upon future occurrence of Error 842:

```
1> sp_configure "dump on conditions",1
2> go
1> sp_shmdumpconfig "add", error, 842, 1, <directory>, <file_name>
2> go
```

Additional information

Have the following information ready when you call Sybase Technical Support:

- Adaptive Server version and SWR rollup level
- Server error log
- Operating System error log
- Output of sp_shmdumpconfig

Version in which this error is raised

All versions

Error 847

Severity

10

Message text

```
Checkpoint process detected hardware error writing
logical page '%ld', virtual page '%ld' for dbid '%ld',
cache '%.*s'.  It will sleep until write completes
successfully.
```

Note This error may be caused by a hardware problem.

Explanation

Adaptive Server's automatic checkpoint mechanism guarantees that data pages changed by completed transactions are regularly written from memory to the database device. To accomplish this, the checkpoint process periodically checks the number of log records in the transaction log for each database. The number of log records, and a recovery interval defined with sp_configure, are used to decide when to checkpoint each database. During the checkpoint all pages that have been modified in memory, but not on disk, since the last checkpoint are written out to the database device. checkpoint is performed on a per database basis for all the databases on an Adaptive Server.

Error 847 occurs when the checkpoint process is unable to write out pages to the database device due to a hardware problem. The error is preceded by the following message in the error log:

```
bufwritedes: write error detected - spid=%ld, ppage=%ld,
bvirtpg=%ld, dbid=%ld
```

Action

This is an informational message and the checkpoint task will periodically retry writing pages to the device. To identify the affected devices so you can investigate the hardware problem, refer to “Creating and Managing User Databases” in the *System Administration Guide* and read “Getting Information About Database Storage”.

Note If hardware problems were detected, execute the dbcc checkdb and dbcc checkalloc commands for the affected database to ensure that no database corruption occurred due to the problem.

Version in which this error is raised

All versions

Error 849

Severity

```
24
```

Message text

```
I/O error detected for device '%.*s' (with  physical
name '%.*s', and virtual device number %d).
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server encounters an I/O error on a read or write request made to a Sybase device. The error is accompanied by Error 823 and is seen during dbcc operations.

The parameters in the error message refer to the Sybase logical device name, the physical operating system location of the device, and the virtual device number (a unique identifier for the device).

Error 849 is related to one of the following:

- Hardware problems on a disk where a Sybase device resides.
- A disk mirroring problem in Adaptive Server. Disk mirrors allow an Adaptive Server database device to be duplicated, that is, all writes to the device are also copied to a separate physical device. If an I/O to a mirrored device fails, Adaptive Server "unmirrors" the bad device and continues to run unmirrored. Due to the disk mirroring problem, while a disk failure in the primary device unmirrors the device, it does not stop the server from initiating further I/O's to the unmirrored device, raising Error 849.

Action

Check your operating system logs for any information about hardware problems on a disk containing a Sybase device. Since Error 849 accompanies and provides supporting data for 823 errors, refer to the writeup on Error 823 for more troubleshooting information. In the event of hardware failure, you will need to restore the affected database(s) from clean backups.

If your site uses disk mirroring, Error 849 may indicate that Adaptive Server continued to issue I/O's to a failed primary device. Contact Sybase Technical Support for assistance.

Additional information

When calling Technical support, have the following information available:

- Server version and SWR version level
- Server error log
- Text of all error messages
- Text of operating system logs

Version in which this error is raised

All versions

Error 852

Severity

20

Message text

```
Command Failed: Invalid request to move '%ld'  buffers
from the '%ld'K pool of the '%.*s' cache to the '%ld'K
pool. Source  pool only contains '%ld' buffers. Retry
the command specifying fewer buffers.
```

Explanation

The Adaptive Server Cache Manager works in conjunction with the Buffer Manager to manage physical disk pages while they are in memory. When a page is read into memory from the disk, a buffer header is assigned to it. The header tracks the usage and the contents of the page image. The page and the buffer header together define a buffer.

The buffers themselves are grouped together to form Memory Address Space Segments (MASSes). A buffer cache consists of MASSes linked in MRU/LRU (most recently used/least recently used) chains. The Cache Manager manages MASSes and buffer caches.

The Adaptive Server command sp_poolconfig allows you to create, drop, resize, and get information about memory pools within named data caches. Error 852 occurs when you use sp_poolconfig or a configuration file to create or resize memory pool size and Adaptive Server is unable to find contiguous memory to fulfill the request.

Error 852 occurs with the following states:

State	Meaning
1	Error 852 occurs with State 1 when the size of the destination pool is smaller than the requested memory size.
2	Error 852 occurs with State 2 when the requested pool size is larger than the source pool size.

Action

1 Use sp_cacheconfig to display information about the current configuration of the affected data cache and memory pools in the cache. For example:

```
1> sp_cacheconfig test_cache
2> go
Cache Name    Status    Type      IO Sz  Wash Sz   Config Value Run Value
------------- --------- --------- ------ --------- ------------ --------
test_cache    Active    Mixed                          2.00 Mb 2.00 Mb
test_cache    Active    Mixed       2 Kb    512 Kb      0.00 Mb 1.50 Mb
test_cache    Active    Mixed       4 Kb    100 Kb      0.50 Mb 0.50 Mb
                                                     ------------ --------
                                              Total      2.00 Mb 2.00 Mb
```

- The first line identifies the cache. Lines after the first line show defined memory pools.
- The *IO sz* column shows the I/O size for a memory pool.
- The *Config Value* column shows the size the cache or pool will have after the next time Adaptive Server is restarted. If the value is 0, the size has not been explicitly configured and a default value will be used.

- The *Run Value* column shows the size of the cache or pool now in use on Adaptive Server.

2 Try the sp_poolconfig command again with appropriately-sized values or use sp_poolconfig to reconfigure the source pools to have an appropriate number of buffers available.

Additional information

Refer to the *Adaptive Server Enterprise Reference Manual* for information about sp_cacheconfig and sp_poolconfig.

Version in which this error is raised

All versions

Error 855

Severity

16

Message text

```
Bad prefetch size of %d encountered.
```

Explanation

On select, update, and delete statements, you can specify a prefetch size that specifies the I/O size in kilobytes for tables bound to caches for which large I/Os are configured. Valid values for size are 2K, 4K, 8K, and 16K. You must have configured the I/O size for the cache before you can specify it for a prefetch size. The procedure sp_helpcache shows the valid sizes for the cache to which an object is bound, or for the default cache.

When you specify a prefetch size, the Adaptive Server Parser passes that value to the Buffer Manager. If the value you specified is not available, Error 7380 (a warning) is raised, and processing continues with the prefetch size that the Adaptive Server optimizer determines to be the “best” value.

Error 855 occurs when:

- A problem occurs with the passing of the prefetch value and
- The prefetch size that is passed to the Buffer Manager is smaller than the minimum allowed size or larger than the maximum allowed size.

This error is only raised by diagserver.

Action

Call Sybase Technical Support.

Version in which this error is raised

All versions

Error 861

Severity

16

Message text

```
Command Failed: Cannot bind object '%ld',  dbid '%d' to
cache because you have an open cursor on the target
object.   Retry the command after closing the cursor.
```

Explanation

You can use the stored procedure sp_bindcache to bind a database, table, index, or text or image object to a named data cache. If there is an open cursor on the target object, the cursor structures have an open descriptor for the object and there may be cache information stored there. Error 861 occurs when you use sp_bindcache or the configuration file to change object binding and Adaptive Server finds that the object has been opened by a cursor.

Action

Close the cursor and try the command again:

1 Get information about cursors you have open:

```
1> sp_cursorinfo null
2> go
```

2 If the output from step 1 shows an open cursor for the object you are trying to bind, close the cursor:

```
1> close cursor_name
2> go
```

3 If the output from step 1 did not show an open cursor for the object you are trying to bind, this probably means someone else opened it. Use sp_lock to determine who has locks on the object. If you can determine who has the open cursor, ask them to close it.

4 If you cannot determine who has the open cursor on the object you are trying to bind, shut down and restart Adaptive Server or wait and try again later.

5 Try your command again.

Version in which this error is raised

All versions

Error 863

Severity

20

Message text

```
Buffer resources in cache %s, id %d are unavailable.
Please re-run this query or ask the system administrator
to re-configure buffer cache memory.
```

Explanation

When a page is read into memory from disk, a buffer header is assigned to it. The header tracks the usage and the contents of the page image. The page and the buffer header together define a buffer. The buffers themselves are grouped together to form Memory Address Space Segments (MASSes). A buffer cache consists of MASSes linked in MRU/LRU (most recently used/least recently used) chains. Error 863 occurs when you run a query and Adaptive Server is unable to find an empty buffer in any buffer pool in the named cache.

Error 863 occurs with the following states:

State	Meaning
1	Adaptive Server could not find a buffer pool that had any available buffers in it.
2	When Adaptive Server makes a prefetch request for a MASS of a specific size, it checks the page range for the request to determine whether any of the pages within the range are already in cache. If any of the pages are already in cache, Error 863 occurs with State 2. This state is caused by an Adaptive Server internal error.

Action

Rerun the query later (when buffer resources in that cache might be free) or use sp_cacheconfig to increase the size of the cache.

Additional information

Refer to the *Adaptive Server Enterprise Reference Manual* for information about sp_cacheconfig.

Version in which this error is raised

All versions

Open Database Manager Errors

This section contains error messages for the Open Database Manager.

Most 9xx errors are encountered during recovery. Many of these errors may be raised as a result of the master database being corrupt. If this is the case, refer to Chapter 1, System Database Recovery for help.ollolol

Error 903

Severity

23

Message text

```
Unable to find row in sysindexes for clustered  index
on system catalog %d in database %d. This index should
exist in all  databases. Run DBCC CHECKTABLE on
sysindexes in the database.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server is unable to access the sysindexes table for a database.

Error 903 can occur under the following circumstances when Adaptive Server accesses sysindexes:

- Actions that involve opening a database since those actions access the sysindexes table for that database (for example, database recovery)
- Actions that need to access columns in sysindexes. Such actions include:
  - order_by clause compilation.
  - Running the update statistics command.
  - Threshold Manager space calculations. Refer to “Managing Free Space with Thresholds” in the *System Administration Guide* for information about thresholds.

Possible causes of Error 903 are:

- A problem has occurred with the device on which sysindexes resides and sysindexes has become corrupted.
- The partition on which the database resides has been mounted as a file system or overlapping partitions exist. Refer to Chapter 2, “Correct Use of Raw Partitions”.

- An Adaptive Server problem.

Action

Determine whether the device on which the database resides is damaged and solve any problems that exist.

If the database named in the error message is accessible, run dbcc checkdb and dbcc checkcatalog to determine whether the database is corrupted. If the database is corrupted, restore it from backup.

If you cannot find any device problems and the database is not corrupted, call Sybase Technical Support.

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- select * from database_name..sysindexes where id=object_id("sysindexes") output
- dbcc checkdb and dbcc checkcatalog output
- Text of all the error messages

Version in which this error is raised

All versions

Error 905

Severity

17

Message text

```
Unable to allocate a DBTABLE descriptor to  open
database '%S_DBID'. Another database must be closed or
dropped before  opening this one.
```

Explanation

As part of open database processing, the Database Table Manager controls and synchronizes access to database tables. When Adaptive Server tries to open a database and there are not enough database descriptors available, Error 905 occurs. (The total number of database descriptors available is controlled by the **number of open databases** configuration parameter.)

The action in the text of the 905 message only applies to environments where:

- You cannot have any more open databases (for instance, because of memory limitations), or
- You cannot increase the value of the number of open databases configuration parameter until a later date because you cannot shut down and restart Adaptive Server at this time.

Action

When Error 905 occurs, select one of the following strategies to solve the problem.

Short Term

In the short term, you can close or drop another database. This is not a long-term solution. You can use this option until an Adaptive Server restart can be scheduled, as a restart is necessary for the long-term solution.

To close a database, make sure no users are accessing the database or put the database into single-user mode. This does not guarantee that the database will not be accessed by Adaptive Server (for example, for an automatic checkpoint), so the 905 error may still occur.

Long Term

As a long-term solution, use sp_configure to increase the value of the number of open databases configuration parameter:

1 Determine your current value:

```
1> sp_configure "number of open databases"
2> go
```

2 Determine what your new value should be. One way to do this is to determine the number of databases you have now and then add a padding factor for future growth.

For example:

```
1> select count(*) from sysdatabases
2> go
-----------
         12
```

3 Set the new value. For example:

```
1> sp_configure "number of open databases", 15
2> go
```

where “15” is the new value (12 existing databases plus a padding factor of 3 for future growth).

4 Shut down and restart Adaptive Server to put the new value in effect.

Note If you increase the value of the number of open databases configuration parameter, make sure you reconfigure Adaptive Server memory appropriately. 17,408 bytes are required for each additional open database. Refer to "Configuring Memory" in the *System Administration Guide* and the *Performance and Tuning Guide* for information about Adaptive Server memory use.

Version in which this error is raised

All versions

Error 906

Severity

23

Message text

```
Could not locate row in sysobjects for system  catalog
%d in database %d. This system catalog should exist in
all databases.  Run DBCC CHECKTABLE on sysindexes in the
database.
```

Note This error may be caused by a hardware problem.

Explanation

Adaptive Server expects to find certain information in static locations within a database or on the Server:

- Page 1 of each database is the first page for sysobjects
- Page 24 of each database is the first page for sysindexes

As part of open database processing, the Descriptor Manager searches page 1 for entries for sysobjects and page 24 for entries for sysindexes. If the entries are not found on the expected page, Error 906 will occur.

Error 906 occurs during open database actions such as recovery and create database.

Possible causes of Error 906 are:

- The partition on which the database resides has been mounted as a file system or overlapping partitions exist. Refer to "Correct Use of Raw Partitions".

- A database device has been created that exceeds the operating system maximum size limit (for example, 2GB on some 32-bit machines).

Action

Error 906 means that, for the database specified in the message, either page 1 or page 24 is corrupted. Restore that database from backup.

Version in which this error is raised

All versions

Error 908

Severity

22

Message text

```
Unable to find any entries in Sysusages for  dbid
'%d',database '%.*s'. Run DBCC CHECKTABLE on Sysusages
in the master  database.
```

Explanation

The master..sysusages table keeps track of the space that you assign to Adaptive Server databases. create database and alter database allocate new space to the database by adding a row to sysusages for each database device or device fragment.

The Adaptive Server Database Manager maps logical pages (pages in an Adaptive Server database) to virtual pages (pages in a database device). There is a one-to-one correspondence between these two types of pages.

Error 908 occurs during recovery when Adaptive Server attempts to generate the mapping of logical to virtual pages for a specified database, but is unable to find any rows for the database in the sysusages table.

Action

Check the sysusages table:

```
1> use master
2> go
1> dbcc checktable(sysusages)
2> go
```

If errors are reported, try to fix them using the sections in this manual that describe those errors.

If dbcc checktable reported no errors, and you do not need to recover the database in question, remove the database from your Adaptive Server installation. For further information, refer to Chapter 2, "How to Drop a Database When drop database Fails". After removing the database, restart Adaptive Server and check that the 908 error is eliminated.

If dbcc checktable reported no errors, but you need to recover the database in question, start by removing the database from your Adaptive Server installation as described above. Next, rebuild the database using your creation scripts, and reload it from a clean backup.

Version in which this error is raised

All versions

Error 911

Severity

16

Message text

```
Attempt to locate entry in sysdatabases for  database
'%.*s' by name failed - no entry found under that name.
Make sure  that name is entered properly.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs during a use command when Adaptive Server tries to access a database that is not listed in sysdatabases. For example:

- A user tries to access a database not listed in sysdatabases:

  ```
  1> use database_name
  2> go
  ```

- A user tries to connect to Adaptive Server when an invalid database is indicated by the defaultdb column of master..syslogins. The user is connected only to master, or not connected at all, and Error 911 occurs.

Action

1 Obtain a list of all databases in sysdatabases by executing sp_helpdb or the following:

```
1> select name from master..sysdatabases
2> go
```

The list returned contains the only valid databases recognized by Adaptive Server and available for use.

2 Either create the missing database or change the reference to a database from the list created in step 1:

- If a user-issued use command created the error, either create a database following regular procedures (refer to "Creating and Managing User Databases" in the *System Administration Guide)* or change the use command to refer to one of the databases from the list created in step 1.
- If the defaultdb option of sp_addlogin refers to an invalid database, the System Administrator ("sa") or user needs to specify a new, valid, default database. To specify a new default database, choose a database name from the list in step 1 and use the following procedure:

```
1> sp_modifylogin user_name, defdb, database_name
2> go
```

Although you can successfully change your own default database at any time with sp_modifylogin, you can access the database only if the database owner ("dbo") has added you as a user in the database.

Additional information

Refer to the *Adaptive Server Enterprise Reference Manual* for information about sp_addlogin and sp_modifylogin.

Version in which this error is raised

All versions

Error 913

Severity

22

Message text

```
Could not find row in sysdatabases with database  id %d.
Run DBCC CHECKTABLE on sysdatabases.
```

Explanation

This error occurs when Adaptive Server cannot find an entry in sysdatabases when looking for a specific database ID.

Possible causes of Error 913 are:

- Accessing a stored procedure or view that refers to a table in a database that has been dropped.
- Accessing a stored procedure or view that refers to a table in a database that has been dropped and re-created.
- A nonclustered index on the sysdatabases table in the master database has been corrupted.

This is a serious error if the index is corrupt or recovery fails.

Action

If the 913 error is caused by accessing a stored procedure or view that refers to a table in a database that has been dropped, either re-create the database or update the procedure or view to point to the appropriate database.

When a database has been dropped and re-created, it can end up with a different database ID. If a stored procedure still references the old database ID, drop and re-create the stored procedure so that it will reference the new database ID. Refer to drop procedure and create procedure in the *Reference Manual* for information about dropping and re-creating stored procedures.

If the cause of the error is a corrupted nonclustered index on a system table, use one of these options:

- Use the sp_fixindex stored procedure to repair the index.
- Restore master from backup. Refer to Chapter 1, "System Database Recovery" for information about restoring the master database.

Version in which this error is raised

All versions

Error 916

Severity

14

Message text

```
Server user id %d is not a valid user in database '%.*s'
```

Explanation

When a user tries to open a database, Adaptive Server performs the following checks to determine the validity of the user:

1 Looks for the suid (Adaptive Server user ID, copied from syslogins) of the process in sysusers to find the corresponding uid (user ID, unique in this database) and gid (group ID to which this user belongs).

2 Looks for a "guest" row in sysusers (the suid for "guest" is -1).

3 Looks for the user ID in sysalternates.

4 Checks sysusers to determine whether the user has "sso" role if the user does have an suid in the database. This check is related to Adaptive Server auditing.

Error 916 occurs when a user tries to open a database and Adaptive Server finds that the user ID is not valid.

Action

Decide whether you want the user to access the database with:

- A normal login and user ID

- A guest user ID
- An alias to another user ID
- If you want the user to access the database with a normal login and user ID, determine whether the Adaptive Server user ID exists in the current database:

```
1> use database_name
2> go

1> select * from sysusers
2> where suid = Server_user_ID
3> go
```

  where *Server_user_ID* is the Adaptive Server user ID displayed in the error message. If the user is not found, use sp_adduser to add them.

  If the user is found and you are getting Error 916 for that user, run dbcc checkdb on the database to find out whether it is corrupt. If dbcc checkdb returns error messages, corruption has occurred. Follow the directions in this manual for recovering from those errors.

- If you want the user to access the database as guest, determine whether there is a "guest" row in sysusers of the database being opened:

```
1> use database_name
2> go

1> select * from sysusers
2> where suid = -1
3> go
```

  where *database_name* is the name of the database in the error message. If guest is not found, use sp_adduser to add it.

  If guest is found, run dbcc checkdb on the database the user is trying to access to find out whether it is corrupt. If dbcc checkdb returns error messages, corruption has occurred. Follow the directions in this manual for recovering from those errors.

- If you want the user to access the database with an alias to another user, determine whether the user ID exists in sysalternates of the database being opened:

```
1> use database_name
2> go

1> select * from sysalternates
2> where suid = Server_user_ID
3> go
```

If the user ID is not found, use sp_addalias to add an alias for them.

If the user ID is found and you are getting Error 916 for that user, run dbcc checkdb on the database the user is trying to access to find out whether it is corrupt. If dbcc checkdb returns error messages, corruption has occurred. Follow the directions in this manual for recovering from those errors.

- If the database is being accessed due to a referential constraint, make sure the user ID exists in each database.

If Error 916 occurs after the above actions have been taken, call Sybase Technical Support.

Additional information

Refer to the *Reference Manual* for information about sp_addlogin, sp_adduser, and sp_addalias.

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level
- Text of all error messages
- Output from syslogins
- Output from sysusers and sysalternates for the database being opened

Version in which this error is raised

All versions

Error 921

Severity

14

Message text

```
Database '%.*s' has not been recovered yet  - please
wait and try again.
```

Explanation

This error occurs during the database recovery phase of Adaptive Server startup when you try to access a database for which recovery has not yet been completed. It can also occur if the database was marked suspect due to an earlier problem.

If the database has been marked suspect by an earlier problem, this is a serious error and must be corrected if you want to access your database again.

Action

The time it takes to recover a database depends on:

- The number and size of transactions that must be committed or rolled back

- The amount of memory

Determine whether the recovery phase of the database has completed:

- Examine your Adaptive Server error log
- Enter two select * from sysprocesses where spid=1 commands about 20 seconds apart. Find the *spid* of the process doing the recovery (the value of spid for the process doing database recovery is always "1" during Adaptive Server startup) and determine whether *physical_io* is increasing over time. If it is, recovery is still in progress.

If the database is still being recovered, wait until recovery is complete and then try to access the database.

If you are still unable to access the database and the database has been marked suspect because of an earlier problem, determine what the earlier problem was by looking at the Adaptive Server error log. Follow the instructions in this manual for recovering from that error. If no instructions exist, call Sybase Technical Support.

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log, including boot sequence and all errors
- Operating system error log
- select * from master..sysusages output
- select dbid, name, status from sysdatabases output
- Text of all the error messages

Version in which this error is raised

All versions

Error 924

Severity

14

Message text

```
Database '%.*s' is already open and can only  have one
user at a time.
```

Explanation

This error occurs when an attempt is made to access a database that has been set to single-user mode and is being accessed by another user. Only one user at a time can access a database that is set to single-user mode.

Action

1 Verify that the database being accessed is in single-user mode:

```
1> sp_helpdb database_name
2> go
```

If the database is in single-user mode you will see “single user” in the status column along with any other options that are set for the database.

2 Execute sp_who to determine who is accessing the database. Look for the database name under the dbname column and the login name in the loginame column.

3 Contact the individual and arrange for database access or contact your System Administrator (“sa”). The “sa” or database owner (“dbo”) can set the database to multi-user access when the database is not in use by executing the following from the master database:

```
1> sp_dboption database_name, single, false
2> go

1> use database_name
2> go

1> checkpoint
2> go
```

Refer to the *Reference Manual* for information about sp_dboption.

Version in which this error is raised

All versions

Error 925

Severity

19

Message text

```
Maximum number of used databases for each query has been
exceeded. The maximum allowed is %d.
```

Explanation

Adaptive Server has an eight-level-deep stack of open databases for each process. Every time a database is opened, it gets pushed onto the stack, and when it is closed, it is removed from the stack.

There is usually one open database for each process. If you run a query that uses a table in a different database, Adaptive Server opens that database just long enough to use the table and then closes the database.

Sometimes, however, a database must remain open longer. For example, if you execute a stored procedure in a different database than the one you are currently using, Adaptive Server will open that database and keep it open until the stored procedure finishes.

Only eight databases can be open at once. A single database opened multiple times is added to the stack each time it is opened. This can result in Error 925, even though only one database has been used. If you have stored procedures that call other procedures in other databases, you can exceed this limit.

Note The use *database_name* command closes the database that was previously open, so Error 925 would *not* be generated by eight use *database_name* commands in a row in Transact-SQL.

An example of a set of queries that exceeds the eight-level limit follows:

```
1> use db9
2> go
1> create proc pr9
2> as
3> select * from table1
4> go
1> use db8
2> go
1> create proc pr8
2> as
3> exec db9..pr9
4> go
1> use db7
2> go
1> create proc pr7
2> as
3> exec db8..pr8
4> go
1> use db6
2> go
1> create proc pr6
2> as
3> exec db7..pr7
4> go
1> use db5
2> go
1> create proc pr5
2> as
3> exec db6..pr6
```

```
4> go
1> use db4
2> go
1> create proc pr4
2> as
3> exec db5..pr5
4> go
1> use db3
2> go
1> create proc pr3
2> as
3> exec db4..pr4
4> go
1> use db2
2> go
1> create proc pr2
2> as
3> exec db3..pr3
4> go
1> use db1
2> go
1> create proc pr1
2> as
3> exec db2..pr2
4> go
```

Action

Be aware of the eight-level open database limit when writing queries.

If you believe Error 925 is occurring for a reason other than a query that has more than eight database openings, restarting Adaptive Server should clear the problem.

Version in which this error is raised

All versions

Error 926

Severity

```
14
```

Message text

```
Database '%.*s' cannot be opened. An earlier  attempt
at recovery marked it 'suspect'. Check the SQL Server
errorlog for  information as to the cause.
```

Explanation

This error occurs when you attempt to reference a database that has been marked suspect in one of the following circumstances:

- During start-up of Adaptive Server.
- By the System Administrator as a result of certain critical errors.

This is a serious error and must be corrected if you want to access your database again.

Note This error may be caused by a hardware problem.

Error 3414 is a related error which occurs during startup when Adaptive Server cannot complete the recovery of a database. Refer to Error 3414.

Action

Since the 926 error is the result of an earlier error or action, the recommended action is to determine what caused the database to be marked suspect.

In order to determine the cause, check the Adaptive Server error log for error messages for the database in question and try to eliminate those first by using the troubleshooting procedures in this manual.

Depending on why the database was marked suspect, you may choose to remove its suspect flag if you are certain that the critical error which caused the database to be marked suspect has been resolved (for example, if one of the database devices was not available when Adaptive Server was started and you are sure that the device is available now). If you choose to reset the suspect status, refer to Chapter 2, "How to Reset a Database's "suspect" Status".

If you cannot find any procedures recommended for your specific errors, call Sybase Technical Support for assistance.

If the specified database does not contain important data or if you have a known, clean backup of it, you may choose to drop it first, re-create it, and then load the clean database dump into it. Refer to Chapter 2, "How to Drop a Database When drop database Fails" for information on how to drop a database that has been marked suspect.

Before loading the database dump into the newly created database, make sure that the new database and the dumped database have the same data and log mapping, and the same user segment definitions. Refer to Error 2558 for information about how to do this.

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log, including boot sequence and all errors
- Operating system error log

- select * from master..sysusages output
- select dbid, name, status from sysdatabases output
- Text of all the error messages

Version in which this error is raised

All versions

Error 930

Severity

14

Message text

```
Database '%.*s' cannot be opened because either an
earlier system termination left LOAD DATABASE
incomplete or the database is created with 'for load'
option. Load the database or contact a user with System
Administrator (SA) role.
```

Explanation

This error occurs when you attempt to reference a database:

- Whose load database process has been terminated before completion, or
- That was created with the for load option and that has not been loaded yet.

Action

Eliminate the 930 error by determining which of the above situations applies to you and then following the appropriate set of steps below.

Previous Load Database Failed

Since the 930 error is, in this case, the result of an earlier error, the recommended action is to determine what caused the load database process to be terminated before completion.

1. Check the Adaptive Server and Backup Server error logs for error messages about the database in question, and try to eliminate those first by using the troubleshooting procedures in this manual. If you cannot find any procedures recommended for your specific errors in this manual, call Sybase Technical Support for assistance.
2. Examine your operating system error log to determine if hardware errors may have caused the corruption. If errors of this type occurred, solve the problem that caused them.
3. Enter the load database command again.

Incorrect Use of *create database for load*

Load the database before you try to reference it.

If you do not have enough space allocated to the specified database and get Error 3105 when you attempt to load the database, use one of the following methods to increase the size of the database to be loaded:

- Drop the newly created database first and then re-create it with the size specified in the Error 3105 message. If you cannot drop the database with drop database because the database has been marked suspect, and you are getting Error 926, refer to "How to Drop a Database When drop database Fails".

Or:

- Use the command alter database db_name for load.

  To avoid 2558 errors or a remapped database, the database being loaded into must be identical to the dumped database in terms of data, log, and segment mapping. Refer to Error 2558 for more information.

If the specified database is the default database for a certain user, that user will get Errors 930 and 4001 and will be put in the master database when they connect to the Adaptive Server. Use the following command if you want to change that user's default database:

```
1> sp_modifylogin user_name, defdb, database_name
2> go
```

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log
- Backup Server error log
- Operating system error log
- select * from master..sysusages output
- select * from master..sysdatabases output
- Text of all the error messages

Version in which this error is raised

All versions

Error 935

Severity

10

Message text

```
WARNING - the timestamp in database '%.*s' is
approaching the maximum allowed.
```

Explanation

Each database has one global timestamp which is kept in memory in a structure called a dbtable. The timestamp is not a date and time; rather, it is a sequence number which keeps track of modifications of pages within the database. Each time a page in the database is modified, the page gets the current timestamp and the global timestamp gets incremented. The global timestamp "travels" with a database through dumps, loads and upgrades. Timestamps are critical during recovery, to ensure that transactions are properly rolled forward or rolled back. The global timestamp is also used for timestamp columns.

The maximum value for timestamps is very high (0xffff 0xffffffff), so it is very unlikely that under normal circumstances a database will reach this limit. However, use of some unsupported database maintenance commands can cause a higher than expected global timestamp value.

Each time dbtable is created for the database, Adaptive Server checks whether the timestamp value is approaching the maximum by comparing the current timestamp value with a threshold value defined by Adaptive Server. Error 935 is raised if the timestamp value exceeds this threshold.

Action

Error 935 gives you advance warning about a serious impending problem.

Warning! After Error 935 is raised, only a few more transactions are possible in the database. If you do not address the problem, the timestamp value will reach the maximum and the database will no longer be usable.

Database dump/load is *not* an option to address this problem. When you receive Error 935, re-create the database in question as soon as possible using these steps:

1 bcp out all table data.

2 defncopy out all procedures, rules, defaults and triggers.

3 Use the drop database command to drop the database.

4 Use the create database command to re-create the database.

5 Create all tables.

6 bcp in all table data.

7 Create indexes.

8 defncopy in all procedures, rules, defaults and triggers.

9 Check that the user IDs in sysusers in the restored database are consistent with master..syslogins. Mismatched user IDs in databases can lead to permission problems.

Scripts are recommended for reliable re-creation of the database objects.

Additional information

Refer to the glossary for additional information about timestamps and dbtable. Refer to "Developing a Backup and Recovery Plan" in the *System Administration Guide* for complete information about how to safely re-create databases.

Version in which this error is raised

All versions

Error 940

Severity

26

Message text

```
Dbtable in wrong state for operation: %S_DBT.
```

Explanation

Adaptive Server tracks the status of each open database using a memory structure called a dbtable, which contains state information such as the number of open transactions in the database, the number of open objects, keep counts that ensure proper use of each dbtable, and so on. Error 940 is raised when a database operation is attempted which is incompatible with the current state of its dbtable.

Error 940 is due to an internal Adaptive Server problem. It results in a stack trace and prevents further work in the server.

The error is raised with the following states:

State	Meaning
1	Database is being dropped but the dbtable is not marked for dropping.
2	dbtable was marked for dropping by a task other than the one dropping the database.
3	Database is being dropped but the dbtable's keep count is not zero. This state is only raised by diagserver.
4	Database is being dropped but is not marked active.
5	dbtable is being locked but the database is not active.
6	While decrementing the dbtable keep count, it was found to be less than zero.
7	Illegal attempt by a process to increment the keep count.

State	Meaning
8	Similar to state 7.
10	Invalid pointer to dbtable structure.
11	While terminating use of a dbtable, the keep count was found to be less than zero.
12	Attempted to unlock a dbtable which is not currently locked.
13	Attempted to cleanup a dbtable not marked for cleanup.
14	A non-system task attempted to update the keep count.
15	Attempted to update a newly created dbtable which is not marked as created. (Raised by diagserver.)
16	Similar to state 14
17	Attempted to free a dbtable not marked for cleanup. (Raised by diagserver.)
18	Attempted to free a dbtable whose keep count is not 1.
19	Retrieved a dbtable from the free list, but the dbtable is not marked as free.
20	Illegal attempt to mark a dbtable active (raised by diagserver).
21	Illegal attempt by a process that does not own a dbtable to keep it.
22	Attempted to uninstall a dbtable that is not in the appropriate state ((raised by diagserver).
23	Attempted to uninstall a dbtable that is not on the active chain.
25	Database is not in 'drop' state as expected.
26	dbtable was found to be free although it is not marked free.
27	dbtable was found to be free although its keep count is non-zero.
29	dbtable is found to be active although it is not marked active.
30	More than one dbtable is found for a database.

Action

Shut down and restart the server. If the problem persists, call Sybase Technical Support. You may need to upgrade to a version in which the problem is resolved.

Additional information

Have the following information ready when you call Technical Support:

- Server version and SWR rollup level
- Server error log

Version in which this error is raised

11.5.1 and earlier

Error 941

Severity 20

Message text `Illegal database context operation.`

Explanation

The Adaptive Server Database Context Manager controls access to databases. Any process opening or closing a database must do so through this subsystem, which maintains a list of open databases, the server user ID (suid) accessing the database, and other information in memory. The manager also accesses the Process Status Structure (PSS), which is a memory structure associated with each process describing the state of the process and usage of various resources used by the process.

Error 941 is raised when the Database Context Manager encounters a problem in fulfilling a database context change request.

Error 941 occurs with the following states:

State	Meaning
1	A request is made to close the current database and return the context to the previous database for the requestor. However there is no previous database on the context list.
2	A request is made to close the last database open for the requestor. However the context list indicates that this is not the last open database.
3	A database close request does not include a context for the action (such as returning to the previous database).
4	A request is made to initialize the context list. However the list is not empty and cannot be initialized.

Error 941 is due to corruption of the PSS or other memory structures.

Action

Shut down and restart Adaptive Server to clear any memory corruption. If the problem persists, increase the value of the number of open databases parameter.

Version in which this error is raised

All versions

Error 945

Severity 21

Message text

Version 11.0.3 and Later

```
Unable to locate page %ld for database id %d in the in-
memory copy of Sysusages (spid=%ld). This is an internal
error. Please contact Sybase Technical Support.
```

Version 11.0.2 and Earlier

```
Function %s was asked to locate page %ld in database
%ld. The page number is invalid for the database.
```

Note This error may be caused by a hardware problem.

Explanation

When the Adaptive Server Open Database Manager requests a page (for allocating extents and a number of other operations), it receives a page number and a pointer to a disk map array. It then determines the index of the map element that contains the target page and returns that index. When Adaptive Server cannot find the requested page number in the database disk map array, Error 945 occurs.

Possible causes of Error 945 are:

- Memory corruption.
- Bad allocation page pointers (sysusages table corruption).

Action

1 Shut down and restart Adaptive Server. If the corruption is in memory, this should clear the 945 error. If Error 945 persists, go on to step 2.

2 Check your record of sysusages information for the database in the error message against what is currently in master..sysusages. (If you do not have a record of your sysusages information, try sp_help database_name to get a subset of the sysusages information.) To determine the sysusages information currently held by Adaptive Server, type:

```
1> select * from sysusages where dbid = database_ID
2> go
```

where *database_ID* is the number from the error message output.

3 If the information you have on record for that database does not match what is currently in sysusages, refer to Chapter 2, "Restoring Device and Database Information in the System Catalog" for information about restoring a corrupted sysusages table.

4 If the information you have on record for that database matches what is currently in sysusages, the database specified in the error message is corrupt. Restore the database from known, clean backups.

Version in which this error is raised	All versions

Error 949

Severity	14
Message text	Database '%.*s' is unavailable. It is being bound to a named cache.
Explanation	When a database is bound to a named cache, all objects in that database are bound to that cache unless they are explicitly bound to a different cache. Binding a database to a named cache requires an exclusive lock on the database. Error 949 occurs when you try to access an object in a database that is in the process of being bound to a named cache.
Action	Try again later, once the bind has completed.
Additional information	Refer to the *Adaptive Server Enterprise Reference Manual* for details about sp_bindcache.
Version in which this error is raised	All versions

Error 950

Severity	14
Message text	Database '%.*s' is currently offline. Please wait and try your command again later.
Explanation	The Adaptive Server Database Context Manager controls access to databases. Any process opening or closing a database must do so through this subsystem. Error 950 is raised when the Database Context Manager is unable to honor a database access request because the database is offline, and is not available for general use. During a load database, Adaptive Server takes the database being loaded offline, and the database remains in this state until it is brought back online. During an Adaptive Server installation upgrade, a database is marked offline when loading the transaction log that is to be upgraded. You may see 950 errors:

- If the load completed and the upgrade was successful, but the database was not brought back online.
- If the upgrade failed, leaving the database offline.

Action

If a load is in progress, complete the load sequence (including any loads of transaction log dumps), then use online database to make the database available for use.

If this error is raised due to an upgrade failure, the corrective action depends on several factors:

- The original SQL Server/Adaptive Server version, and the version to which you are upgrading.
- The phase of the upgrade process where failure occurred.
- Whether the database(s) being upgraded has tables or other objects marked for replication. If the transaction log is not drained because log records are waiting to be processed by the Log Transfer Manager, the upgrade process for that database will fail, leaving the database offline.

Refer to "Recovering From a Failed Installation or Upgrade" in the Installation and Configuration Guide for your platform to determine how to proceed.

Additional information

If you use Replication Server, you must disable replication before attempting to upgrade. Refer to the installation guide for further information. For example:

- If upgrading to version 11.5, see "Upgrading a SQL Server with Replicated Databases" in your Adaptive Server installation guide.
- If upgrading to version 11.0.x, see "Replication Server: Preparing to Upgrade SQL Server" in your SQL Server installation guide.

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level
- Text of all error messages

Version in which this error is raised

All versions

Page Manager Errors

This section contains error messages for the Adaptive Server Page Manager.

Error 1105

Severity

17

Message text

Version 11.5

```
Can't allocate space for object '%.*s' in database
'%.*s' because '%.*s' segment is full/has no free
extents. If you ran out of space in syslogs, dump the
transaction log. Otherwise, use ALTER DATABASE or
sp_extendsegment to increase size of the segment.
```

Version 11.0.x and Earlier

```
Can't allocate space for object '%.*s' in database
'%.*s' because the '%.*s' segment is full. If you ran
out of space in syslogs, dump the transaction log.
Otherwise, use ALTER DATABASE or sp_extendsegment to
increase the size of the segment.
```

Explanation

This error occurs when Adaptive Server is unable to allocate space in a database. The state of the error indicates what type of space could not be allocated:

Table 1-4: Allocation error: type of space

Type of Space That Cannot Be Allocated			
State 1	**State 2**	**State 3**	**State 4**
Data space	Extent allocation space	Log space	Log space

1105 errors are written to the error log only when they occur during the checkpoint process. When a user process encounters this error, the message is returned to the client application without being written to Adaptive Server's error log.

To correct the error, you must know:

- The Server version level.
- Whether the error occurred during run time or during recovery:

- Run time is a period during which Adaptive Server performs normal daily activities, such as receiving a query from a client application, processing a transaction, or returning results from a transaction.
- An Adaptive Server is in recovery when doing activities such as restarting, performing the processing that occurs after a database dump is loaded, and performing the processing that occurs after a transaction dump is loaded.

- The state of the error: 1, 2, 3, or 4.

Action

Use the sections below to correct the error, depending on whether the error was at run time or during recovery.

Runtime 1105 Errors

If the error occurred during runtime, use the next section to correct the error. Refer to "Recovery 1105 Errors" if the error occurred during recovery.

Runtime 1105 Errors: State 1 or State 2

The data segment is full on the indicated data base. To obtain more data space, do one or both of the following:

- Use alter database or sp_extendsegment to increase the size of the data segment.
- Drop objects from the database or delete rows from a table.

Runtime 1105 Errors: State 3

The log segment is full on the indicated database. To clear space in the log, perform the following steps:

1 Determine approximately how many pages the transaction log occupies using the commands:

```
1> use database_name
2> go
1> select data_pgs (8, doampg)
2> from sysindexes where id=8
3> go
```

where “8” is the ID for syslogs. The result is the number of data pages (2K pages on most operating systems) that the transaction log occupies.

Note The query results may be inaccurate by as many as 16 pages, but using the query is much faster than counting the rows in the transaction log.

2 Dump the inactive portion of the transaction log using the dump transaction command. If this command fails with the 1105 error, retry the command using the with truncate_only option.

If this command also fails with the 1105 error, run dump transaction with no_log.

Note Using the dump transaction with no_log command can result in an 813 error. Therefore, do not use dump transaction with no_log unless it is absolutely necessary (when dump transaction with truncate_only will not truncate the log).

3 Repeat step 1. If the transaction log occupies significantly fewer pages, continue with step 4..

If, however, there are still a large number of pages in the syslogs table, an outstanding transaction is probably preventing the log from being cleared. If this is the case, choose one of the following options:

- Refer to Chapter 2, “How to Detect and Clear Long-Running Transactions”.
- Restart Adaptive Server and repeat step 2. When Adaptive Server starts and this database is recovered, the outstanding transaction is rolled back, allowing the log to be cleared by a subsequent dump transaction. Depending on the size of the outstanding transaction, rolling back the transaction during startup recovery may be significant.

4 If the dump transaction command was executed with either the no_log or the truncate_only options in step 2, dump the database now, because these options prevent subsequent changes recorded in the log from being used to recover from a media failure. You must do a dump database now to ensure the recoverability of subsequent database modifications using transaction logs.

Note This database dump is not required if your backup and recovery procedures do not include saving the transaction logs for media failure recovery.

Do not assume that the occurrence of 1105 errors automatically means your transaction log is too small. If the data and the log are on the same segment, either the State 1 and State 2, or the State 3 actions may free enough space without increasing the size of the transaction log.

If you are concerned that your transaction log is not large enough for your Adaptive Server, refer to "Creating and Managing User Databases" in the *System Administration Guide* and read the section "Estimating the Transaction Log Size" before deciding to increase the log size. If appropriate, refer to the *ASE Reference Manual* for more information about the alter database command used to increase log size. Remember that once they have been increased, log and database sizes cannot be decreased.

The above procedures may not be effective in resolving the 1105 errors due to excessive data growth in mixed data and log segments; you may need to move the data to a different device in this situation.

Runtime 1105 Errors: State 4

The log space is full on the indicated database. LCT (last chance threshold) was trying to empty the log but ran out of space trying to log information. Perform the same procedure as in "Runtime 1105 Errors: State 3".

If dbcc checktable on syslogs shows space available in the transaction log, check your sysusages system table. Error 1105 may occur if the unreservedpgs column for this database has NULL values in it when it should have numeric values. Update that column to 0 (instead of NULL).

Runtime 1105 Errors: Replicated Databases

You may see 1105 errors when the log is full (States 3, 4) on a database that is replicated or has been loaded from a replicated database. This occurs because Adaptive Server is unable to truncate the log past the LTM truncation point, even with truncate log on checkpoint set.

To correct this problem, instruct the server to ignore the LTM truncation point for the database:

```
1> dbcc settrunc (ltm, ignore)
2> go

1> dump transaction with truncate_only
2> go
```

The server will now be able to truncate the log.

Runtime 1105 Errors: *tempdb*

1105 errors occur on tempdb while Adaptive Server is processing transactions, never during recovery.

The tempdb transaction log behaves just like a user database transaction log with the sp_configure option trunc log on chkpt set on. But data management is more difficult with tempdb because it is more difficult to get information about space use: worktables cannot be referenced by the user, and the log is truncated approximately every 60 seconds. The default size of tempdb is 2MB when Adaptive Server is installed.

Note Read all relevant sections about transaction log management in the *System Administration Guide* and the *Reference Manual* before making any changes at your site.

Possible causes of the 1105 error in tempdb, and their symptoms and solutions include:

Table 1-5: Causes of Error 1105 in tempdb

Symptom	Possible Causes	Solution
Too many transactions fill up tempdb transaction log.	Doing too much work in tempdb (or temporary tables) fills up the log.	Design your application for more efficient use of tempdb.
	Processing more transactions per 60 seconds than can be logged in tempdb with current space.	Increase size of tempdb's transaction log (refer to alter database in the*Reference Manual* for details) or perform manual dump transaction with truncate only commands on tempdb.
One transaction is too big.	Your transaction may create a Cartesian product which fills the log. There may be too many steps between the begin transaction and the commit transaction. A single statement such as insert, delete or update may involve so many rows that the transaction log fills up before the statement can finish.	Rewrite the transaction to avoid Cartesian products and long individual queries. Read about transaction log management in the *System Administration Guide* and the *Reference Manual* for information and suggested strategies.

Recovery 1105 Errors

If the 1105 error occurred during run time, use the section "Runtime 1105 Errors" to correct the error. If the error occurred during recovery, use the procedure below which corresponds to the database on which the 1105 error occurred:

- A user database
- The master database
- The model database

Recovery 1105 Errors on a User Database

Use the following procedure to correct 1105 errors on a user database which occur during recovery:

1 Manually set the database's status to no chkpt on recovery and single-user:

Warning! Because the database was marked suspect on the original recovery attempt, this procedure also resets some internal status bits to allow the database to recover normally. Therefore, *do not use this procedure under any other circumstances,* or you may create additional serious problems.

```
1> sp_configure "allow updates", 1
2> go

1> begin transaction
2> go

1> update master..sysdatabases
2> set status = status & ~256
3> where name = "database_name"
4> go

1> update master..sysdatabases
2> set status = status | 4112
3> where name = "database_name"
4> go
```

2 Check that each of the above update commands affected only one row. If more than one row was affected, issue a rollback transaction. Otherwise, commit the transaction and shut down Adaptive Server:

```
1> commit transaction
2> go
1> shutdown
2> go
```

3 Restart Adaptive Server.

4 Dump the transaction log with the no_log option and reset the database status:

```
1> use master
2> go
1> dump tran database_name with no_log
2> go

1> sp_dboption database_name, "no chkpt", false
2> go

1> sp_dboption database_name, single, false
2> go

1> use database_name
```

```
2> go

1> checkpoint
2> go

1> sp_configure "allow updates", 0
2> go
```

Recovery 1105 Errors on the *master* Database

If an 1105 error occurs on the master database during recovery and Adaptive Server will not run, correct the problem with the following procedure:

1 Create a runserver file that includes the 3607 trace flag. This trace flag allows Adaptive Server to restart without going through normal recovery, and should only be used in this procedure. Refer to Chapter 2, "How to Start Adaptive Server with Trace Flags" for instructions.

2 Start Adaptive Server with the 3607 trace flag using the runserver file created in step 1. The recovery process will perform the minimum actions required to allow access to each database, including master.

3 Once Adaptive Server has restarted and the master database is accessible, dump the transaction log with the no_log option and shut down the Adaptive Server:

```
1> dump tran master with no_log
2> go
1> shutdown with nowait
2> go
```

Note It is mandatory that you use shutdown with nowait here. Failure to use the with nowait option will suicide every database's transaction log.

4 Restart Adaptive Server with the original runserver file. The database should recover normally. Now dump the master database:

```
1> dump database master
2> to logical_dump_device
3> go
```

Recovery 1105 Errors on the *model* Database

If an 1105 error occurs on the model database during recovery, Adaptive Server may not start. The error occurs because the tempdb database is required to start Adaptive Server, but it cannot be built due to the problem with model. To correct the problem and restart Adaptive Server follow these steps:

1 Create a new runserver file that includes the 3608 trace flag. This trace flag causes Adaptive Server to recover only the master database. Refer to Chapter 2, "How to Start Adaptive Server with Trace Flags" for instructions.

2 Start Adaptive Server with the 3608 trace flag using the runserver file created in step 1.

3 Manually set model status to no chkpt on recovery and single-user, as follows:

```
1> sp_configure "allow updates", 1
2> go
1> begin transaction
2> go
1> update master..sysdatabases
2> set status = status & ~256
3> where name = "model"
4> go
1> update master..sysdatabases
2> set status = status | 4112
3> where name = "model"
4> go
```

4 Check that each of the above update commands affected only one row. If more than one row was affected, issue a rollback transaction. Otherwise, commit the transaction and shut down Adaptive Server:

```
1> commit transaction
2> go
1> shutdown with nowait
2> go
```

5 Restart Adaptive Server, using your regular runserver file.

6 Execute the following command to truncate the inactive portion of the transaction log in model:

```
1> dump tran model with no_log
2> go
```

7 Reset the database options on model with the following command:

```
1> sp_dboption model, "no chkpt", false
2> go
1> sp_dboption model, single, false
2> go
1> use model
2> go
1> checkpoint
```

```
2> go
```

8 Disable updates to the system catalog and shut down Adaptive Server:

```
1> sp_configure "allow updates", 0
2> go
1> shutdown with nowait
2> go
```

Note It is mandatory that you use shutdown with nowait here. Failure to use the with nowait option will suicide every database's transaction log.

9 Restart Adaptive Server.

Version in which this error is raised

All versions

Error 1108

Severity

21

Message text

```
Cannot deallocate extent %ld, database %d. Object id
%ld, index id %d, status %d in extent does not match
object id %ld, index id %d, status %d in object being
deallocated.  Run DBCC CHECKALLOC.
```

Explanation

Error 1108 occurs when Adaptive Server is deallocating an extent, if the extent information on the allocation page does not match the information for the object being deallocated. This can be caused by corruption of the allocation page or memory corruption.

A software problem can cause this error to occur:

- During rollback of a transaction
- When rows for an object are being deleted or updated
- When an object is being dropped

This is a serious error as it indicates corruption has occurred.

Error 1108 is only raised by diagserver.

Action

1 Run dbcc checkalloc on the database named in the message. Resolve any errors reported by dbcc checkalloc using information about those errors in this manual.

2 If errors still exist, restore the database from a clean backup or call Sybase Technical Support.

Additional information

If you need assistance from Sybase Technical Support, have the following information available when you call:

- Server version and SWR version level
- Server error log
- Text of all error messages
- Operating system error log

Version in which this error is raised

All versions

Error 1120

Severity

20

Message text

```
Attempt to call pg_readalloc on page %ld which is not
an allocation page.
```

Explanation

Allocation pages are used to track space in the database. Each allocation page for an object has an entry in the Object Allocation Map (OAM) for that table or index. The OAM entry for each allocation page stores the number of used and unused data/index pages on that allocation page.

Error 1120 occurs when you have an invalid page number for an allocation page in the allocation section of an OAM. In this case, Adaptive Server will not have accurate information about which pages have been allocated.

This is a very serious error as Adaptive Server cannot correctly insert or update data when it occurs. Action should be taken immediately or further database corruption will occur.

Action

Use the page number in the error message to identify the object (and thus the database) that is affected by the error. Refer to Chapter 2, "How to Find an Object Name from a Page Number" for details.

To recover from Error 1120, restore your database either using bcp or from clean backups. Refer to Chapter 2, "How to Rescue Data from a Corrupted Table" for information about using bcp to save your data.

Additional information

For information about OAM pages, refer to "Checking Database Consistency" in the *System Administration Guide* and see the section "Understanding Page and Object Allocation Concepts".

If you need assistance from Sybase Technical Support, have the following information available when you call:

- Server version and SWR version level
- Server error log
- Text of all error messages
- Operating system error log
- Output of dbcc page for the corrupted page
- Output of dbcc checkalloc for the database identified by dbcc page

Version in which this error is raised

All versions

Error 1124

Severity

20

Message text

```
Failed to get lock on allocation page %ld.
```

Explanation

Adaptive Server uses allocation pages to track space in the database. Allocation pages are updated when an object's space requirements change.

Adaptive Server raises Error 1124 due to the following sequence of events:

- A server process acquires a lock on an object's allocation page in order to allocate pages for the object.
- An exception occurs during allocation. This necessitates "undoing" the allocation steps.
- The server is unable to acquire a lock necessary to abort the allocation, resulting in Error 1124.

This error may be accompanied by a kernel error ("Current process infected with %d"); Error 945 ("Unable to locate page %ld for database id %d in the in-memory copy of Sysusages (spid=%ld). This is an internal error. Please contact Sybase Technical Support."); Error 1205 ("Your server command (family id #%d, process id #%d) was deadlocked with another process and has been chosen as deadlock victim. Re-run your command.") or other errors.

Error 1124 is due to an Adaptive Server problem when rolling back space allocation.

Action

When Error 1124 is accompanied by a 945 error on tempdb (dbid = 2), which is not resolved by shutting down and restarting the server, there may be insufficient space on tempdb or a problem with tempdb storage. You can resolve these problems by following the instructions in "How to Reduce the Size of tempdb" in Chapter 2.

In all other instances of the 1124 error, call Sybase Technical Support.

Additional information

Have the following information ready before calling Sybase Technical Support:

- Server version and SWR rollup level
- Server error log

Version in which this error is raised

11.5.x and earlier

Error 1127

Severity

22

Message text

```
Internal Error:  The OAM has no room for new inserts.
Check OAM for object %ld, index %d.
```

Explanation

The Object Allocation Map (OAM) tracks information about the allocation and deallocation of storage space for tables and indexes. At least one OAM page exists for each table (except syslogs and sysgams) and for the table's indexes. The OAM page has an entry for each allocation page used by the object. The first OAM page for an object has a maximum of 240 entries, and any subsequent OAM pages have a maximum of 250 entries per page.

When a new page is allocated for the table or index:

- The appropriate OAM page is updated.
- If necessary, a new OAM entry is created for the allocation page which controls the page being added.
- If necessary, a new OAM page is created and linked in to the current OAM page chain for the object. This occurs when there are no available slots for a new OAM entry in any existing OAM page for the object.
- The total number of entries in the object's OAM (corresponding to the total number of allocation pages with extents for this object) is adjusted during table updates and also during database recovery.

Error 1127 is raised:

- If the total number of entries for all OAM pages for the object exceeds its upper limit (overflow in OAM page header). This can happen if the table grows and occupies space on many allocation pages.
- If the total number of entries is corrupted. This is due to an Adaptive Server problem.

Action

There are different ways to correct this error, depending on the conditions under which it occurred and the nature of the OAM problem. Other errors may also be raised at the same time.

If this error is raised when you attempt to start Adaptive Server and the Server hangs, the database may be marked suspect; contact Sybase Technical Support for assistance.

If the error occurs during a table update (for example during a bulk copy operation), follow these steps:

1 If the object is a system table (a system table's object ID is less than 100) and the index ID is *not* 0, refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index.

   If the object is a system table and the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption but it may be necessary to restore the database from clean backups.

   If the object encountering the error is *not* a system table, continue with step 2.

2 If this is a user table, the problem may be due to an overflow in the OAM page header (although this is unlikely), or corruption of the page header. Run the following command to display OAM page information (you must have sybase_ts_role granted to you to run this and also run dbcc traceon(3604) first):

```
1> dbcc listoam(database_name, object_ID, index_ID)
2> go
```

   For example:

```
1> dbcc listoam(pubs2, 104001776, 0)
2> go
-------------------------------------------------
Objid:   104001776     indid: 0
OAM pg cnt:    1       Entry cnt:    1
Rows:         18       Rows Per pg:  6
Used pgs:      4       Unused pgs:   4
OAM status bits set:   PG_OAMPG, PG_OAMSORT
```

```
OAM pg #      520 has the following entries
                 (allocpg:used/unused):
                 512: 4/ 4
------------------------------------------------
DBCC execution completed. If DBCC printed error
messages, contact a user with System Administrator
(SA) role.
```

Call Sybase Technical Support with this information.

Additional information

For more information about OAM pages, refer to the section ”Understanding Page and Object Allocation Concepts” in “Checking Database Consistency” in the *System Administration Guide*.

Before calling Sybase Technical Support, have the following information ready:

- Server version and SWR version level
- Text of all error messages
- Server error log
- Output of dbcc listoam command for the table

Version in which this error is raised

All versions

Error 1129

Severity

22

Message text

```
The requested update to the OAM for object %ld, dbid
%ld, index %d, oam page %ld, would cause a negative page
count. allocation page: %ld, oamarray used: %ld,
oamarray unused: %ld, deltaused: %ld, deltaunused: %ld,
spid: %ld. Contact your System Admin.
```

Explanation

Most tables and all indexes have at least one Object Allocation Map (OAM) page. A single OAM page holds information about up to 250 allocation pages (248 if the object is an identity column). Each allocation page for an object has an entry in the OAM pages for that table or index. The OAM entry for each allocation page stores the number of used and unused data and index pages on that allocation page.

When a new page is needed or a page is deallocated, the OAM entry has to be updated. The number of used and unused pages on the allocation page is not allowed to be less than 0. Error 1129 occurs if the number becomes less than 0.

Action

1 Look at the error message to determine the object ID and the index ID. If the object encountering the error is *not* a system table (a system table's object ID is less than 100), continue with step 2.

   If the object is a system table and the index ID is *not* 0, refer to "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index.

   If the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption but it may be necessary to restore the database from clean backups.

2 For user tables, if the index ID is 0 or 255, continue with step 3.

   If the index ID is *not* 0 or 255, translate it into an index name:

```
1> use database_name
2> go
```

```
1> select name from sysindexes
2> where id = object_ID and indid = index_ID
3> go
```

   To ensure that the information needed to re-create the index is available, run sp_helpindex on the index prior to dropping it.

   Drop the index.

   Re-create the index. This clears the corruption in most cases.

   Run dbcc checktable on the table to verify that the corruption is gone.

3 If the index ID is 0 or 255, do one of the following:

   - Restore the database from clean backups.
   - Refer to Chapter 2, "How to Rescue Data from a Corrupted Table".

Additional information

For more information about OAM pages, refer to the section "Understanding Page and Object Allocation Concepts" within "Checking Database Consistency" in the *System Administration Guide*.

If you need to call Sybase Technical Support, have the following information ready:

- Server version and SWR version level
- Text of all error messages

- Server error log
- dbcc listoam(database_name, object_ID, index_ID) output (you must have *sybase_ts_role* granted to you to run this and also run dbcc traceon (3604) first)

Version in which this error is raised

All versions

Error 1131

Severity

22

Message text

```
The OAM for object %ld, index %d, oam page %ld,
database=%d, is currently allocated to another object.
```

Explanation

The Object Allocation Map (OAM) tracks information about the allocation and deallocation of storage space for tables and indexes. At least one OAM page exists for each table (except syslogs and sysgams) and for the table's indexes. Each allocation page used by the object has an entry in the OAM page, showing the number of used and unused data/index pages on that allocation page.

Error 1131 is raised when Adaptive Server attempts to validate an OAM page belonging to an object, and determines that the page is actually allocated to a different object.

Error 1131 occurs with the following states:

State	Meaning
1	While updating the used and unused counts for an OAM entry, the OAM page is found to be allocated to a different object. This error is raised only by *diagserver*.
3	While checking the validity of an OAM page, it is found that the page is allocated to a different object. This state of Error 1131 is raised in versions 11.5.x and earlier.

Action

1131 errors are often seen when you use the system functions data_pgs, reserved_pgs, used_pgs, or rowcnt in a very active database, or a highly volatile database such as tempdb, where objects are being created and dropped frequently. The error can occur:

- if you use and invalid object id in the function;
- if the page is allocated to another object, or is uninitialized.

Use trace flag 1116 to suppress the 1131 error when using these system functions in an active or volatile database environment.

Additional information

For help with trace flags, see Chapter 2, "How to Start Adaptive Server with Trace Flags".

Version in which this error is raised

All versions

Error 1133

Severity

22

Message text

```
Page %ld was expected to be an OAM page for  %ld and it
is not.
```

Explanation

Every Object Allocation Map (OAM) page has a unique number in the database. Error 1133 occurs when an attempt by Adaptive Server to retrieve an OAM page by specifying the page number fails because there is no OAM page in the current database having that page number.

Error 1133 can occur when checkalloc is run if what is expected to be the OAM page in the OAM page linkage is not an OAM page. The linkage is accessed via the *doampg* or *ioampg* column for the object in sysindexes. (*doampg* is the sysindexes pointer to the *data* pages' OAM. *ioampg* is the sysindexes pointer to the *index* pages' OAM.)

This error can also occur when, due to a problem with Adaptive Server, the OAM page number that is requested is greater than the maximum page number for the database.

In these cases, Error 1133 is a serious error and may result in reading the wrong OAM page into data cache and corrupting a buffer in the process.

Action

1 Use the procedure described in Chapter 2, "How to Find an Object Name from a Page Number" to translate the page number displayed in the error message into an object ID and an index ID.

2 If the corruption is on a user table's index (object ID is 100 or greater, and the index ID is between 1 and 250), dropping and re-creating the index should clear the corruption.

   Use sp_helpindex to list all indexes on the table and then rebuild all the nonclustered indexes using one of two methods:

   - If the index ID is between 1 and 250, drop and re-create each nonclustered index on the table.

Or:

- If a clustered index also exists on the table, drop and re-create it, which causes all nonclustered indexes to be automatically rebuilt. If your table is large, you may not have the space to do this (a rule of thumb is that 150 percent of your table size must be available).

3 If the corruption is not on an index page, run dbcc tablealloc or dbcc checkalloc on the corrupted object. The output should report OAM allocation errors. Try to fix the reported allocation errors first by referring to the appropriate writeups in this manual.

If the 1133 errors persist, call Sybase Technical Support for assistance.

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log
- Text of all the error messages
- Operating system error log
- Output from dbcc page for the corrupted page
- Output of dbcc checkalloc for the database identified by dbcc page

Version in which this error is raised

All versions

Error 1141

Severity

26

Message text

```
Unexpected value returned to page manager routine: %ld.
```

Explanation

Adaptive Server uses Object Allocation Map (OAM) pages to track space allocation for Adaptive Server objects. Error 1141 occurs when Adaptive Server tries to update an OAM page following page allocation or deallocation, but the update fails.

Error 1141 may be raised:

- when the database transaction log is full, so that the necessary log records cannot be written to record page allocations.
- due to an Adaptive Server problem.

The error is raised in the following states, which indicate the specific page manager activity where the OAM update failed. Pages in a given state may involve Allpages Locked (APL) tables, Data-only Locked (DOL) tables, temporary work tables or sort tables, or combinations thereof.

State	Meaning
1	There was an unexpected error when determining whether to update the page counts.
2	OAM update failed when allocating extents; also when deallocating one or more pages in an extent.
3	Update failed when doing post-commit cleanup following deallocation.
4	Update failed when allocating an extent for an APL table.
6	Update failed when deallocating an extent for an APL table.
7	Update failed when deallocating an extent for an APL table, and there were no unused pages in extent.
8	Update failed when performing OAM reallocation. Applies to both APL and DOL tables.
9	Update failed when deallocating an extent for sort pages.
10	Update failed when committing large-scale page deallocations (for example extents with no pages allocated).

When Error 1141 occurs, a stack trace is written to the error log and the current process is terminated.

Action

Other errors may sometimes accompany the 1141 error. For example, an 1105 error can also be raised when the tempdb transaction log becomes full. If other errors are reported, correct them using the information about those errors elsewhere in this manual. Then retry the transaction.

To check if the 1141 error resulted from a full transaction log, use the procedures outlined in "Estimating the Transaction Log Size" in the *System Administration Guide*.

If the 1141 error persists after you have checked for other errors and corrected any log space problems, call Sybase Technical Support.

Additional information

Have the following information ready before calling Technical Support:

- Adaptive Server version and SWR Rollup level
- Server error log

Version in which this error is raised

All versions

Error 1142

Severity

22

Message text

Version 11.0.2 and Later

```
Invalid OAM Page %ld. Found pstat=0x%x, object=%ld,
database=%d, indid=%d.
```

Version 11.0.1 and Earlier

```
Invalid OAM Page %ld. Found pstat=%d,
object=%S_OBJID,database=%S_DBID,indid=%d.
```

Explanation

Depending on what caused it, Error 1142 can be either a serious error due to actual corruption in the database or a result of misusing a system function (data_pgs, reserved_pgs, rowcnt, or used_pgs). Error 1142 can occur during run time as well as during recovery (during Adaptive Server startup, load database, or load transaction).

Incorrect System Function Use

The error can occur from incorrect use of any of the following system functions that retrieve information about allocation pages:

- data_pgs
- reserved_pgs
- rowcnt
- used_pgs

For example, if you incorrectly type:

```
1> select reserved_pgs (id, root) from
2> master..sysindexes
3> where id = object_ID("table_name")
4> go
```

you will get Error 1142, because the reserved_pgs system function attempted to retrieve an invalid OAM page. (Replacing root in the above example with a correct OAM page, for example doampg, would give expected results.)

Incorrect OAM Page Entries

If you get Error 1142 and you are sure that you did not use system functions across databases, this error is a sign of serious corruption in your database. You can verify this by running dbcc tablealloc on the object or dbcc checkalloc on the database as discussed under Action, “Incorrect OAM Page Entry in sysindexes”, Step 3.

Adaptive Server uses OAM pages to track how space is allocated for Adaptive Server objects. Error 1142 occurs when Adaptive Server tries to retrieve an OAM page and the status of that page indicates that it is not an OAM page. It occurs with the following states:

State	Meaning
1	The status of the first OAM page indicated that it was not a valid OAM page.
2	Adaptive Server encountered an invalid OAM page during recovery or during an undo or a redo of a truncate table command.
3	Adaptive Server encountered an invalid OAM page while preparing to read and install an OAM page in the buffer cache.
4	Adaptive Server encountered an invalid OAM page during deallocation of an object page.

Action

Recovery from this error depends on why the error occurred. Follow the instructions below, depending on what caused the error.

Incorrect System Function Use

When you want to use any of the system functions specified in Explanation, "Incorrect System Function Use" across databases, select the database to be accessed with a use database command and then run the query for the current database.

Recovery

If Error 1142 occurred during the threshold accounting phase of recovery, the database is marked "suspect" and the object referenced in the error message output is corrupted. To determine whether the error occurred during this phase of recovery, look in the error log to determine whether the error occurred after either of the following types of messages:

- Messages like the following appear when there is nothing to recover:

```
00:95/07/11 15:42:56.91 server  Recovery dbid 5 ckpt (4612,2)
00:95/07/11 15:42:56.92 server  Recovery no active transactions before
ckpt.
```

- Messages like the following appear after transaction recovery completes:

```
00:95/07/11 15:42:58.11 server  Recovery dbid 7 ckpt (322,26) oldest
tran=(322,25)
00:95/07/11 15:42:59.14 server  469 transactions rolled forward/back.
```

If Error 1142 occurred after messages like these, perform the following steps to recover from the 1142 error. Otherwise, go to "Incorrect OAM Page Entry in sysindexes".

1 Turn off threshold accounting for the database affected by the 1142 error so that it can complete recovery:

```
1> use master
2> go
1> sp_dboption database_name,
2> "no free space acctg", true
3> go
```

where *database_name* is the name of the database in the error message.

2 For the affected database, reset the suspect status using one of the methods supplied in Chapter 2, "How to Reset a Database's "suspect" Status". Shut down and restart the server to complete recovery.

3 The object named in the error message is still corrupted. Follow the instructions in Chapter 2, "How to Rescue Data from a Corrupted Table" to recover the data.

Note Be sure to perform step 3 before going to step 4. Otherwise, the next time you restart Adaptive Server, the database that got the 1142 error will get it again.

4 Turn threshold accounting back on for the database affected by the 1142 error:

```
1> use master
2> go
1> sp_dboption database_name,
2> "no free space acctg", false
3> go

1> use database_name
2> go

1> checkpoint
2> go
```

where *database_name* is the name of the database in the error message.

Incorrect OAM Page Entry in *sysindexes*

In this case, Error 1142 is a serious error and may result in reading the wrong OAM page into data cache and corrupting a buffer in the process.

Follow these steps to clear Error 1142:

1 Use the procedure described in Chapter 2, "How to Find an Object Name from a Page Number" to get dbcc page output for the page specified in the 1142 output. If the *page status bits* row has the value 0x8000, the page is a valid OAM page. If the value 0x8000 does not appear in the dbcc page output for the *page status bits* row, the page is not a valid OAM page and the object is corrupted. If it appears that the object is not corrupted, call Sybase Technical Support to determine why the 1142 error occurred.

2 If the corruption is on a user table's index (object ID is 100 or greater and index ID is between 1 and 250), run dbcc checktable to verify that the data page linkage is good. If checktable runs cleanly, dropping and re-creating the index should clear the corruption.

   Use sp_helpindex to list all indexes on the table, and then rebuild all the affected nonclustered indexes using one of two methods:

   - If index ID is between 1 and 250, drop and re-create each nonclustered index on the table.

   Or:

   - If a clustered index also exists on the table, drop and re-create it, which causes all nonclustered indexes to be automatically rebuilt. If your table is large, you may not have the space to do this (a rule of thumb is that 150 percent of your table size must be available).

3 If the corruption is not on an index page, run dbcc tablealloc or dbcc checkalloc on the corrupted object. The output should report OAM allocation errors. Try to fix the reported allocation errors by referring to the information in this manual.

If the 1142 errors persist, call Technical Support for assistance.

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log
- Text of all the error messages
- Operating system error log
- Output from dbcc page for the corrupted page
- Output from dbcc tablealloc or dbcc checkalloc for the corrupted object

Version in which this error is raised

All versions

Error 1143

Severity

22

Message text

Version 11.0.2 and Later

```
Invalid OAM Page %ld. Expected object=%ld. Found
object=%ld, database=%d, pstat=0x%x, indid=%d.
```

Version 11.0.1 and Earlier

```
Invalid OAM Page %ld. Expected object=%S_OBJID. Found
Object=%S_OBJID,database=%S_DBID,pstat=%d, indid=%d.
```

Explanation

This error occurs when the status on the page is not for an Object Allocation Map (OAM) page or the page is an OAM page and the object ID on the page does *not* match the object that the OAM page supposedly belongs to.

Error 1143 is a serious error and may result in reading the wrong OAM page into data cache and corrupting a buffer in the process.

Action

1 Make a note of the object ID and index ID in the error message output.

2 If the corruption is on a user table's index (object ID is 100 or greater and index ID is between 1 and 250), run dbcc checktable to verify that the data page linkage is good. If checktable runs cleanly, dropping and re-creating the index should clear the corruption.

Use sp_helpindex to list all indexes on the table and then rebuild all the affected nonclustered indexes using one of two methods:

- If index ID is between 1 and 250, drop and re-create each affected nonclustered index on the table.

Or:

- If a clustered index also exists on the table, drop and re-create it, which causes all nonclustered indexes to be automatically rebuilt. If your table is large, you may not have the space to do this (a rule of thumb is that 150 percent of your table size must be available).

3 If the corruption is not on an index page, run dbcc tablealloc and/or dbcc checkalloc on the corrupted object. The output should report OAM allocation errors. Try to fix the reported errors first by referring to the appropriate writeups in this manual.

If the 1143 errors persist, call Sybase Technical Support for assistance.

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log
- Text of all the error messages
- Operating system error log
- Output from dbcc page for the corrupted page
- Output from dbcc tablealloc or dbcc checkalloc for the object identified by dbcc page

Version in which this error is raised

All versions

Lock Manager Errors

This section contains error messages for the Adaptive Server Lock Manager.

Error 1203

Severity

20

Message text

```
Caller of lock manager is incorrectly trying to unlock
an unlocked object. spid=%d locktype=%d dbid=%d object
id or page number=%ld.
```

Explanation

Before the Lock Manager tries to unlock an object, it checks to make sure the object is in fact locked. When an object is not locked during this check, Error 1203 occurs.

This error is usually due to an Adaptive Server problem which causes an incorrect session descriptor pointer or logical page number to be passed to the Lock Manager.

Error 1203 can occur during an insert, update, or select, particularly on a table with an index.

Action

Call Sybase Technical Support.

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log
- Output of sp_lock
- Text of all error messages

Version in which this error is raised

All versions

Error 1204

Severity

19

Message text

```
SQL Server has run out of LOCKS. Re-run your command
when there are fewer active users, or contact a user
with System Administrator (SA) role to reconfigure SQL
Server with more LOCKS.
```

Explanation

This error occurs when Adaptive Server runs out of locks . The number of locks available is controlled by the Adaptive Server configuration parameter number of locks.

Following is an example of the output from sp_configure, and a brief description of the output related to locks:

```
1> sp_configure "number of locks"
2> go
Parameter Name        Default     Memory Used Config Value Run Value
--------------------- ----------- ----------- ------------ -----------
number of locks              5000           0         5000        5000
```

- The Default column contains the default value for number of locks, 5000.
- The Memory Used column indicates the amount of memory used by the configured locks. Each lock requires 72 bytes of memory (11.5.x and earlier) or 120 bytes (11.9.x and later).
- The Config Value column contains the value to which the number of locks configuration parameter has been set with sp_configure. If the value has not been explicitly configured, there is a 0 in this column and the default value is used.
- The Run Value column contains the value Adaptive Server is using.

Action

Rerun the command when there are fewer active users, or ask the Sybase System Administrator (“sa”) to reconfigure Adaptive Server for more locks. Increasing the number of available locks impacts performance and memory requirements.

Warning! Before making any changes to Adaptive Server, refer to the *Performance and Tuning Guide* and “Configuring Memory” in the *System Administration Guide* for information about how Adaptive Server uses memory resources.

To increase the number of locks available, complete the following steps:

1 Determine the number of locks currently configured for Adaptive Server:

```
1> use master
2> go
1> sp_configure "number of locks"
```

```
2> go
```

Warning! If you increase the number of locks available by too many locks, Adaptive Server may run out of memory and be unable to restart. Refer to Chapter 1, "Adaptive Server Does Not Start After Altering Configuration" if you run out of memory and cannot restart Adaptive Server.

2 Choose the number of locks you want to configure and issue the following command:

```
1> sp_configure "number of locks", new_value
2> go
```

3 Restart Adaptive Server to operate with the new values.

Additional information

Refer to the *Performance and Tuning Guide* for information about locks and other Adaptive Server resources.

Refer to "number of locks" in the *System Administration Guide* for information about the number of locks configuration parameter.

Version in which this error is raised

All versions

Error 1205

Severity

13

Message text

Version 11.5 and Later

```
Your server command (family id #%d, process id #%d) was
deadlocked with another process and has been chosen as
deadlock victim. Re-run your command.
```

Version 11.0.x and Earlier

```
Your server command (process id #%d) was deadlocked with
another process and has been chosen as deadlock victim.
Re-run your command.
```

Explanation

This error occurs when a process tries to acquire a lock on an object that is locked by a second process when the second process is waiting for a lock on an object that has been locked by the first process. This situation is a deadlock, and can involve more than two processes.

Adaptive Server detects this situation, rolls back the transaction that has accumulated the least amount of CPU time, and notifies the application program of this action with Error 1205. This allows the other users' processes to move forward.

Deadlocks are caused by a number of situations, including:

- Transactions modify tables in different orders. There is a greater chance of deadlock between two transactions if one is processing in the sequence A - B - C while the other runs C - B - A.
- Transactions access tables via a nonclustered index. If the optimizer chooses a different nonclustered index for the same table for two different queries, a nonclustered index is not in the physical data sequence and the two processes are acquiring page locks in a random order. Thus, there is a greater chance that one process will lock a page that the other needs.
- Transactions that use the keyword holdlock or use the set isolation level command to hold shared locks. When holdlock is appended to a select transaction it holds the shared lock for the remainder of the transaction. This increases the risk of deadlock.
- Transactions that require a long time to run. The longer a transaction runs, the more likely it is that another user will require a resource held by the transaction.

Action

Restart the transaction that has been rolled back.

To minimize future occurrences of deadlocks, use any of the following procedures that apply to your site.

Application Error Handling

Each application should have deadlock handling routines. Refer to the dbsetuserdata entry in the *Open Client DB-Library Reference Supplement* for a sample deadlock handling routine.

Use Well-Constructed Transactions

Using transactions constructed to avoid deadlocks greatly reduces their occurrence. Some techniques for writing transactions that avoid deadlock include:

- Access tables in the same order in each transaction. Use coding conventions that require all transactions that access several tables to process them in the same order.

- Access tables via a clustered index when possible. If it is not possible to change a nonclustered index to a clustered index to minimize deadlocks, then trap the deadlock error in the application and provide appropriate recovery routines.
- If you are using holdlock, decide whether you really need to be using it. Use holdlock only when you require repeatable reads within a transaction.
- Avoid long-running transactions. Some ways to avoid long-running transactions are:
  - Never allow user interaction within a transaction.
  - Separate logical units of work into transactions. For example, acquiring a sequential key from a key table for use in an insert statement can be separated into transactions similar to the following (in the first set of commands, you acquire the key; in the second set of commands, you use the key to do the insert):

```
1> declare @key int
2> begin transaction
3> update key_table set key = key + 1
4> select @key = key
5> commit transaction
6> go

1> insert mytable
2> values (@key, ...)
3> go
```

  - Check whether client applications, including third party tools, allow users to accidentally create long running transactions.

Additional information

Refer to “Lock Manager” in the *System Administration Guide*.

Version in which this error is raised

All versions

Error 1243

Severity

26

Message text

```
Couldn't find object to be released.
```

Explanation

When Adaptive Server no longer needs a lock, it releases it. Error 1243 occurs when Adaptive Server cannot find the object for the lock to be released. It occurs with the following states:

State	Meaning
1	Adaptive Server uses address locks to lock memory addresses for buffers containing OAM and allocation pages. If, during the release of an address lock, Adaptive Server fails to find one of the buffers, Error 1243 occurs with State 1.
2	Adaptive Server uses logical locks for database pages and tables. The server searches for the logical lock before releasing it. If, during the release of a logical lock, the server fails to find the requested object on the lock chain, Error 1243 is raised with State 2.

Action

Shut down and restart Adaptive Server.

Note Since the unreleased lock may be an exclusive lock, there may be some objects you cannot access until you restart Adaptive Server.

If Error 1243 occurs again, call Sybase Technical Support.

Version in which this error is raised

All versions

Error 1249

Severity

20

Message text

Version 11.5 and later

```
Process became runnable before lock it was waiting for
was released.
```

Version 11.0.x and earlier

```
Process (spid=%d) became runnable before lock it was
waiting for was released. This is an unexpected lock
manager error. Please contact Technical Support and
provide them the information logged in the errorlog.
```

Explanation

The Lock Manager maintains queues of lock requests to enable server tasks to obtain read, write and update locks as needed. A task sleeps while waiting for a lock request to be granted. Error 1249 is raised when the task awakes but the Lock Manager detects an abnormal condition in the lock queue for that lock. Error 1249 is raised in the following states:

State	Meaning
1	The lock request was granted, but the lock is not at its expected position in the queue.
2	The lock received a signal to requeue, but it is not at its expected position in the queue.

Error 1249 is caused by an Adaptive Server problem and breaks the connection to the server. In a common scenario leading to this error, separate sessions open cursors on the same table, and the user aborts one session in the middle of an update; attempting the same transaction again raises the 1249 error.

Action

Call Sybase Technical Support.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level
- Server error log
- Operating System error log

Version in which this error is raised

All versions

Error 1265

Severity

```
20
```

Message text

```
A lock manager routine expected a lock record as an
argument but was passed something else.
```

Explanation

This error occurs when the Adaptive Server Lock Manager is:

- Releasing a lock
- Transferring a lock
- Waiting for a lock

and the lock record structure is not valid.

This error is usually due to an Adaptive Server problem.

Action

Shutting down and restarting Adaptive Server will usually clear the 1265 error. If it does not clear the error or if the error occurs repeatedly, call Sybase Technical Support.

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log
- Output of sp_lock at the time the 1265 error occurred
- Text of all error messages

Version in which this error is raised

All versions

Error 1279

Severity

16

Message text

```
SQL Server has run out of locks on engine %d. Re-run
your command when there are fewer active users, or
contact a user with System Administrator (SA) role to
reconfigure maximum engine freelocks or number of locks.
```

Explanation

An Adaptive Server engine is a process running an Adaptive Server that communicates with other Adaptive Server processes via shared memory. An Adaptive Server running on a multiprocessor machine can have one or more engines.

Adaptive Server uses two configuration parameters to control lock use:

- number of locks controls lock use Server-wide
- max engine freelocks controls lock use for each Adaptive Server engine

Adaptive Server maintains a list of Server-wide free locks and a list of per-engine free locks. When locks are needed for an engine, Adaptive Server moves a block of free locks from the Server-wide free lock list to the free lock cache for that engine.

Error 1279 occurs when Adaptive Server runs out of locks on an engine and is unable to move any more locks from the Server-wide free list but there is an another engine with some free locks. (Error 1204 occurs when Adaptive Server completely runs out of locks.) The error message includes the number of the engine that has run out of locks.

Action

Run your command again later when there are fewer active users or increase the number of locks available to the engine. If you want to increase the resources, you can do one of the following:

- Increase the total number of locks available Server-wide by using sp_configure to increase the value of the number of locks configuration parameter.
- Decrease the maximum number of locks available for each engine by using sp_configure to decrease the value of the max engine freelocks configuration parameter. By decreasing the value of this parameter, you make more locks available to the engine for which you are getting the error.

Note Increasing the value of the number of locks parameter increases the memory resources used by Adaptive Server. Make sure your Adaptive Server memory is configured appropriately if you change this value.

Additional information

Refer to "Lock Manager" in the *System Administration Guide* for information about Adaptive Server engines and the Adaptive Server Lock Manager.

Refer to the *Reference Manual* for information about sp_configure and configuration parameters.

Refer to the *Performance and Tuning Guide* for information about memory management for Adaptive Server.

Version in which this error is raised

All versions

Sort Manager Errors

This section contains error messages for the Adaptive Server Sort Manager.

Error 1501

Severity 20

Message text Sort failure

Explanation

This error occurs during index creation when Adaptive Server fails to complete a sort operation. There are a number of different states for this error, each indicating a different reason for the failure:

State	Meaning
2	When creating a high level index structure, Adaptive Server is unable to build a parent node in the ancestor list.
3	When creating a high level index structure, Adaptive Server is unable to build a parent node from the root of the tree.
4	When creating a high level index structure, Adaptive Server is unable to build a parent node from the new root of the ancestor tree.
6	The sever is adding the results of an internal sort of a special leaflet node to an external merge in progress, and a single internal merge step has a merge order greater than 8.
7	The server is adding the results of an internal leaf merge node and the previous merge step caused the merge order to be greater than 8.
8	Adaptive Server is starting a non-sort of already sorted data to create a clustered index, and it cannot find a page to an index that was supposed to have been built for the sorted data.
9	The server is unable to allocate and initialize space for managing the sort's buffers.
10	The server is building indexes and the length of a row in an overflow page is either zero or greater than the size of a page.

When this error occurs, no index is created.

The configuration parameters number of sort buffers and sort page count (available in 11.0.x only) control how Adaptive Server uses memory during sorting. number of sort buffers controls the sort buffer size and sort page count controls the sort page count. In version 11.0.x, the default value for both parameters is 0; in later versions, the default value of number of sort buffers is 500. Changing the values of number of sort buffers and sort page count can cause Error 1501.

Error 1501 can also be caused by an Adaptive Server problem.

Action

Version 11.0.x

1 If you have not changed the values of number of sort buffers or sort page count, call Sybase Technical Support.

2 Determine the current values of number of sort buffers and sort page count:

```
1> sp_configure "number of sort buffers"
2> go
1> sp_configure "sort page count"
2> go
```

3 Change the values, perhaps back to 0:

```
1> sp_configure "number of sort buffers", new_value
2> go

1> sp_configure "sort page count", new_value
2> go
```

If you cannot resolve this error by adjusting the number of sort buffers and sort page count values (for example, setting them back to 0), call Sybase Technical Support.

Version 11.5 and later

1 Determine the current value of number of sort buffers:

```
1> sp_configure "number of sort buffers"
2> go
```

2 Change the value, perhaps back to the default value of 500 (recommended except when creating indexes in parallel):

```
1> sp_configure "number of sort buffers", new_value
2> go
```

Additional information

Refer to “create index” in the *Reference Manual* for information about creating indexes.

Refer to “Setting Configuration Parameters” in the *System Administration Guide* for information about configuration parameters.

Version in which this error is raised: All versions

Error 1505

Severity: 14

Message text:

```
Create unique index aborted on duplicate key. Primary
key is '%S_KEY'.
```

Explanation

This error occurs when you try to create a unique index on a column or columns containing at least two rows with the same index value. Each time a unique index is created, Adaptive Server checks for duplicate index values (if data already exists), and displays this message if it finds any.

When this error occurs, no index is created.

A command of the following type generates a 1505 error if there are duplicate values on col1, col2, and col3:

```
create unique index index_name
on table_name(col1, col2, col3)
```

Action

Decide whether you want to allow or prevent duplicate keys in the table where the index will be created. Use one of the following procedures:

- If you decide to allow duplicate keys, do not use the keyword unique when you create the index.
- If you decide to prevent duplicate keys, delete rows that have duplicate index values.

If you choose to delete duplicate keys, use the following procedure. For a brief period of time some rows from the table exist only in a temporary table. If a power failure or other interruption occurs, you could lose the temporary table and the rows in it. Therefore, before you begin this procedure, make a backup copy of the table or modify the procedure to use a table you create for this purpose instead of a temporary table.

1 Find the rows that caused the error:

```
1> select * from table_name
2> group by col1, col2, col3 having count(*) > 1
3> go
```

The query result contains all the rows causing the 1505 error.

2 Put one instance of each distinct row from step 1 into a temporary table:

```
1> select distinct * into #temp_table_name
2> from table_name
3> group by col1, col2, col3 having count(*) > 1
4> go
```

3 Some of the rows may contain identical values in the column or columns to be indexed and different values in other columns. Use the query below to find these rows:

```
1> select * from #temp_table_name
2> group by col1, col2, col3 having count(*) > 1
3> go
```

Examine the individual rows to decide which rows should be deleted. Delete these rows from the temporary table now. The rows you delete in this step will be deleted from the original table in step 4.

4 Delete all the rows which caused the 1505 error from the original table:

```
1> delete table_name from #temp_table_name
2> where table_name.col1 = #temp_table_name.col1
3> and table_name.col2 = #temp_table_name.col2
4> and table_name.col3 = #temp_table_name.col3
5> go
```

5 Each row in the temporary table should have unique values for the columns in question. Add these rows back into the original table:

```
1> insert into table_name
2> select * from #temp_table_name
3> go
```

6 Now you can create the unique index on the table.

Additional information: Refer to “create index” in the *Reference Manual* for information about indexes.

Version in which this error is raised: All versions

Error 1508

Severity: 14

Message text:

```
Create index aborted on duplicate rows. Primary key is
'S_KEY'.
```

Explanation: This error occurs when you try to create a clustered index on a table that contains duplicate rows.

Action

Decide whether you want to allow or prevent duplicate rows in the table where the clustered index will be created. Use one of the following procedures.

Allow Duplicate Rows

Create the index with the allow_dup_row option:

```
1> create clustered index index_name on table_name
2> with allow_dup_row
3> go
```

Subsequent update and insert commands can create duplicate rows after the allow_dup_row option is set.

Prevent Duplicate Rows

Create the index with the ignore_dup_row option:

```
1> create clustered index index_name on table_name
2> with ignore_dup_row
3> go
```

Existing duplicate rows will be deleted from the table as the index is created. Subsequent attempts to enter a duplicate row with insert or update are ignored and the insert or update is cancelled with an informational message. If the insert or update is part of a larger transaction, the rest of the transaction will continue as usual.

Note ignore_dup_row and allow_dup_row are mutually exclusive. Using them together in the same create index statement causes Error 1916 .

Version in which this error is raised

All versions

Error 1509

Severity

20

Message text

Row compare failure

Explanation

This error occurs when Adaptive Server compares two rows during a sort and the compare fails because a row was being compared against itself.

Error 1509 can occur during normal processing or when you try to create an index on a table and the operation fails because of data corruption or an Adaptive Server problem.

This is a serious error as it indicates possible corruption.

Action

1 Run dbcc checktable and dbcc tablealloc on the table being accessed at the time of the error to check for corruption.

2 If there is no corruption, call Sybase Technical Support.

3 If there is corruption, refer to the writeups in this manual for other error messages that appear. Run dbcc checktable and dbcc tablealloc again to determine whether the corruption has been cleared.

4 If corruption still exists after eliminating all error messages you can resolve, copy your data out and back in using the procedure in "How to Rescue Data from a Corrupted Table" or restore your database from clean backups using the procedure below.

> **Warning!** Some data might be lost on this page if you recover your table using bcp or select into (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

If you choose to restore your database from clean backups, do the following:

1 Drop the database. If the drop fails, follow the instructions in Chapter 2, "How to Drop a Database When drop database Fails".

2 Create a database for load. Make sure the database you create has sizes as least as large as those in sysusages for the original database (and that all other sysusages values match the original values). For more information about creating a database for load, refer to "create database" in the *Reference Manual*.

3 Load the database from backup. Refer to "load database" in the *Reference Manual*.

4 Run the online database command to make the database available for use.

Additional information

If you need to call Sybase Technical Support, have the following documentation ready:

- Server release and SWR rollup level
- Server error log output
- Text of all error messages
- dbcc checktable and dbcc tablealloc output

- sp_configure output

Version in which this error is raised: All versions

Error 1510

Severity: 17

Message text: `Sort failed: Out of space in database '%.*s'`

Explanation

Adaptive Server uses a clustered index to sort rows so that their physical order is the same as their logical (indexed) order. When creating a clustered index, Adaptive Server needs a considerable amount of space, 120 to 150 percent of the size of the table, because it needs to sort the data in the table first and then create the clustered index on the sorted data. This space should be available in the database or on the segment where you want to create your index.

Error 1510 occurs when you try to create an index and one of the following is true:

- There is not enough space in the database to create the index. Usually this error occurs with an 1105 error.
- No more locks are available.

Unfortunately, there is no straightforward way to identify which problem caused the error.

Action

Select the appropriate action from the sections below.

Not Enough Space in the Database

Use either of the following two strategies to correct this error:

- Use the instructions from Error 1505 to increase the space available in your database.
- Move the table where the index is to be created to a different, larger, user-defined segment, and then create the index. Refer to “Creating and Using Segments” in the *System Administration Guide* for information about segments.

  a Determine whether you have enough space available on the new segment by typing:

```
1> sp_helpsegment segment_name
2> go
```

b Verify that the database where you want to create your index has enough space allocated on the device or devices that contain the segment from the above query:

```
1> sp_helpdb database_name
2> go
```

No More Locks Available

If the 1510 error occurs because Adaptive Server has no more locks available, use sp_configure to increase the number of locks:

```
1> sp_configure "number of locks", new_value
2> go
```

Shut down and restart Adaptive Server to initialize the new configuration parameter value.

Refer to the *Performance and Tuning Guide* and "Setting Configuration Parameters" and "Configuring Memory" in the *System Administration Guide* for information about configuration parameters and about how Adaptive Server memory is allocated.

Additional information

Refer to "create index" in the *Reference Manual*.

Version in which this error is raised

All versions

Error 1514

Severity

21

Message text

```
Page allocated to sort found to be busy. Page number
%ld.  %S_BUF.
```

Explanation

During index creation, Adaptive Server allocates and uncaches pages to sort the table's rows. The sort routine double checks whether a page allocated for the sort is in the data cache. If it is in the cache, the page cannot be used, the sort operation fails and Error 1514 is raised.

Failure to sort means that the index creation fails; consequently 4313 errors often accompany Error 1514.

Action

If Error 1514 is raised in an isolated case during index creation on a specific table, run dbcc checktable on the table. If errors are reported, try to fix them using the sections in this manual that describe those errors. Next use one of these options to recover from the 1514 problem:

- Retry index creation.
- Shutdown and restart Adaptive Server, and retry index creation.
- Select the table data into a new table, drop the old table, and rename the new table to the old table name. Create the desired index on the table.
- Bulk copy the affected table out, drop and re-create the table, and bulk copy back in. Create the desired index on the table.

If the problem persists, or is not isolated to one table, contact Sybase Technical Support for assistance.

Version in which this error is raised

All versions

Error 1520

Severity

18

Message text

```
Sort failed for table '%.*s' in database '%.*s' because
internal count of used pages is incorrect. Run DBCC
TABLEALLOC to correct the value.
```

Explanation

The create index command allocates a work area based on an estimate in the OAM page of the number of data pages in the table. If this value is too low, the work area is too small, the sort fails, and Error 1520 is raised.

Action

Use the following command to correct this error:

```
1> dbcc tablealloc (table_name)
2> go
```

The dbcc command updates the page count. You can now create the index with the create index command.

Version in which this error is raised

All versions

Error 1530

Severity

16

Message text

```
Create index with sorted_data was aborted because of row
out of order. Primary key of first out of order row is
'%S_KEY'
```

Explanation

The sorted_data option speeds creation of an index when the data in the table is already in sorted order. The speed increase becomes significant on tables larger than 1GB.

Error 1530 occurs when you try to create an index with the sorted_data option on a column or columns containing data that is not in sorted order.

When this error occurs, no index is created.

Note The sorted_data option speeds index creation only for clustered indexes or unique nonclustered indexes. Creating a nonunique, nonclustered index with the sorted_data option may succeed, but there is no improvement in performance.

Action

To create an index after getting a 1530 error, do one of the following:

- Sort the data (using your standard operating system procedures), bcp it into the table, and then create the index with the sorted_data option.
- Create the index without the sorted_data option.

Additional information

Refer to “create index” in the *Reference Manual* for information about indexes.

Version in which this error is raised

All versions

Error 1531

Severity

```
16
```

Message text

```
The sorted_data option cannot be used for a nonclustered
index if the keys are not unique within the table.
Create index was aborted because of duplicate keys.
Primary key is '%S_KEY'.
```

Explanation

This error occurs when you try to create a nonunique, nonclustered index with the sorted_data option on a column or columns containing rows with the same index value. If this situation exists, a statement of the following type will fail:

```
1> create index index_name
2> on table_name(col1, col2, col3)
3> with_sorted_data
```

```
4> go
```

When Error 1531 occurs, no index is created.

Action

To avoid the 1531 error, do one of the following:

- Do not use the sorted_data option, or
- Delete rows with duplicate index values.

Delete Rows With Duplicate Index Values

If you decide to delete rows with duplicate index values, use the following procedure.

Note For a brief period of time, some rows from the table will exist only in a temporary table, and if a power failure or other interruption occurs, you could lose the temporary table and the rows in it. To protect yourself against data loss, you can make a backup copy of the table before you begin the procedure. Alternatively, you can modify the procedure to use a specific table instead of a temporary table.

1 Find the rows that caused the 1531 error:

```
1> select * from table_name
2> group by col1, col2, col3
3> having count (*) > 1
4> go
```

2 Put one instance of each distinct row from step 1 into a temporary table:

```
1> select distinct * into #temp_table_name
2> from table_name
3> group by col1, col2, col3 having count(*) > 1
4> go
```

3 Some of the rows may contain identical values in the column or columns to be indexed and different values in other columns. Use the following query to find these rows:

```
1> select * from #temp_table_name
2> group by col1, col2, col3 having count(*) > 1
3> go
```

4 Examine the individual rows to decide which rows should be deleted. Delete these rows from the temporary table now. The rows you delete in this step will be deleted from the original table in step 5.

5 Delete all the rows which caused the 1531 error from the original table:

```
1> delete table_name from #temp_table_name
2> where table_name.col1 = #temp_table_name.col1
3> and table_name.col2 = #temp_table_name.col2
4> and table_name.col3 = #temp_table_name.col3
5> go
```

6 Now each row in the temporary table should have unique values for the columns in question. Add these rows back into the original table:

```
1> insert into table_name
2> select * from #temp_table_name
3> go
```

7 Create the index with the sorted_data option.

Additional information

Refer to “create index” in the *Reference Manual* for information about indexes.

Version in which this error is raised

All versions

Initialization Errors

This section contains error messages for the Adaptive Server initialization routines.

Error 1601

Severity 21

Message text There are not enough 'user connections' available to start a new process. Retry when there are fewer active users, or ask your System Administrator to reconfigure SQL Server with more user connections.

Explanation

Error 1601 occurs when all user connections are in use and a request is made to create another Adaptive Server process.

Error 1601 is caused by one of the following:

- Adaptive Server is not configured with enough user connections to support all requests.
- On operating systems communicating over DECnet, there are not enough logical links available.

Action

Perform the actions from the appropriate sections below.

Too Few User Connections (All Operating Systems)

Use sp_configure to increase the value of the number of user connections configuration parameter:

1 Determine the current number of user connections configured:

```
1> sp_configure "number of user connections"
2> go
```

2 Set the new value for user connections:

```
1> sp_configure "number of user connections",
2> new_value
3> go
```

3 Restart Adaptive Server to activate the changes.

Depending on the exact configuration of your Adaptive Server, each user connection typically requires between 40K and 60K of memory (86K for OpenVMS), whether or not it is in use. Gradually increase the number of user connections until these errors no longer occur. When you change the value of the number of user connections configuration parameter, you may need to increase the total memory configuration parameter. If you have increased the stack size or default network packet size configuration parameters, each user connection will require more memory. Refer to the *Performance and Tuning Guide* and "Configuring Memory" in the *System Administration Guide* for information about how Adaptive Server uses memory.

If you have an OpenVMS system, recalculate the quotas using the worksheet Sybase provides, and make any needed Sybase quota file and operating system adjustments.

For more information on determining the number of user connections your Adaptive Server needs, refer to "number of user connections" in the *System Administration Guide*.

Note This error often appears in the Adaptive Server error log without the 1601 error number: "kernel: no pss structure available for new process." Take the above action whether or not the error number is displayed.

DECnet Protocol

Make sure there are enough logical links available. For more information about DECnet and logical links, refer to the Adaptive Server installation and configuration guide.

Version in which this error is raised: All versions

Error 1602

Severity: 21

Message text: Unable to initialize network %d

Note This error may be caused by a hardware problem.

Explanation

This error occurs during start-up when Adaptive Server is unable to initialize a network connection specified in your network addressing (*interfaces* or *sql.ini*) file.

Note Throughout this writeup, your network addressing file is referred to as your "interfaces file".

Error 1602 can occur as a result of any of the following problems in your *interfaces* file:

- Duplicate entries
- Incorrect network address
- Other incorrect information such as incorrect host name or illegal port number
- For NT:
  - The WINSOC libraries or the protocol on which you want to run (that is, IPX/SPX) are not loaded.
  - You try to bring up the Adaptive Server on a bad default Server name (NT has default values).
- For NetWare, if the socket (port number) is already used or if the protocol is not loaded (TCP/IP is optional in NetWare)

Error 1602 can also occur as a result of changes in your network which make entries in your *interfaces* file unavailable to Adaptive Server.

Action

Check your *interfaces* file to make sure everything is correct. If any information is incorrect or duplicate entries exist, use one of the following to make corrections to your interfaces file:

- sybinit
- sqledit (for OS/2 Adaptive Servers)
- sqledit or sybinit (for NT Adaptive Servers)

If you do not find any incorrect information in your *interfaces* file, ask your network administrator to verify (at the operating system level) the entries in the *interfaces* file to make sure all network connections are working properly.

Additional information

For information about modifying the *interfaces* file for your platform, refer to the Adaptive Server installation and configuration guide.

If you need to call Sybase Technical Support, have your error log output and a copy of your *interfaces* file ready before you call.

Version in which this error is raised

All versions

Error 1603

Severity

21

Message text

Process priority %d invalid or no process slots available

Explanation

When starting a new process, Adaptive Server checks that:

- the process is being started at the correct priority. Valid priorities range from 0 (highest) to 7 (lowest).
- there are enough connections (slots) available.

Error 1603 occurs when the new process is requested at an invalid priority, or when not enough connections are available to start the process.

Error 1603 occurs with the following states:

State	Meaning
1	Unable to create a new, internal server process.
2	Unable to create a process to handle a new user connection (for example, an isql connection) to the server.

Action

Check the following parameters, and increase the values as needed:

```
number of user connections
number of remote connections
```

Also check any applications currently accessing the server. The number of connections needed by an application depends entirely on how the application is programmed. You may need to shut down and correct any applications that are using excessive connections.

Process Priority

Check the value of process priority reported in the message text. If the value does not fall in the valid range between 0 and 7, check the priority at which applications are accessing the server.

Additional information

number of user connections and number of remote connections are static parameters. You must restart the server after you change their values.

Refer to “Setting Configuration Parameters” in the *System Administration Guide* for details about parameters that control new connections.

Version in which this error is raised

All versions.

Error 1605

Severity

21

Message text

Failed to open virtual socket for new connections

Note This error may be caused by a hardware problem.

Explanation

The Adaptive Server Network Handler listens on the ports specified in your network addressing (*interfaces* or *sql.ini*) file. When a connection comes in, the Network Handler opens another socket to redirect that client, clearing the listener port for the next connection, hence the term “virtual” socket. The connection appears, from the outside, to be connected to the interfaces port, when it is actually connected at a different address to which there is a pointer.

Error 1605 occurs when a failure occurs in the process of opening the new socket. There are a number of things that can cause this:

- Exceeding the number of configured user connections
- Exceeding the number of configured sockets
- Network configuration problems such as using the same IP address for two different clients

Error 1605 error can be raised at the following times:

- At startup time, when the network is initialized.
- At creation of a process, when a new connection request comes in and Adaptive Server is unable to handle it.
- At the forced termination of a process, when a new connection request comes in and Adaptive Server is unable to handle it.

The information below includes several internal configuration parameters. You will not be changing these directly, but they are included to help explain how to calculate values for sp_configure parameters that relate to connection handling. Following are the kernel and configuration parameters used to calculate *cnsocket* (the maximum number of virtual sockets):

- *kimaxfd* – the maximum number of configured per process file descriptors that a process can use at a given time. The error log message "Using ## file descriptors" shows what this value is at boot time.
- number of remote sites – this configuration parameter specifies the maximum number of remote sites.
- number of user connections – this Adaptive Server configuration parameter specifies the maximum number of user connections allowed.
- *cnsocket* – the maximum number of virtual sockets, calculated as:

```
cmaxnetworks (maximum # of user connections configured)
+ cusrconnections (the number of user connections configured)
+ cfgsitebuf (the number of remote sites configured)
+ XTRA_SOCKETS (defined as 2 in 11.0.x)
```

During start-up, if the value for *cnsocket* is greater than the value for *kimaxfd*, then the value of *cnsocket* will be set to the value of *kimaxfd*, and the message:

```
kistartup: number of virtual sockets set to ###
```

is printed to the log where *###* is the value of *kimaxfd*/*cnsocket*.

Adaptive Server recalculates the value for *cnsocket* when you use sp_configure to change related parameter values.

Action

Check the error that precedes the 1605 error. If that error is one of the following:

```
nopen: No virtual sockets available
nopen: no room in socket table
```

then the value for the Server configuration block parameter *cnsocket* has been exceeded. To solve the problem in this case, increase user or network connections and possibly increase the maximum number of file descriptors per process:

1 Determine the current value for *cnsocket*.

   Use dbcc resource (this command requires sybase_ts_role and dbcc tracecon(3604) to determine the current configuration for sockets.

2 Determine the current value for the maximum number of network connections:

```
1> select @@max_connections
2> go
```

   If the value of *cnsocket* is greater than the value of @@*max_connections*, a problem may occur. Step 5 provides details.

3 Determine the current value for number of devices (the number of database devices that Adaptive Server can use, excluding dump devices):

```
1> sp_configure "number of devices"
2> go
```

4 Determine an appropriate value for @@*max_connections*, based on the number of users expected to be logged on at any given time, and the number of devices. Remember that each device takes up a user connection. *@@max_connections* is the kernel resource value *rmaxuconn* and is calculated using the following information:

- *ksnmstprts* – number of Server listener ports configured (in your interfaces file, additional Server listener ports are listed as clones of the "master" line for the Adaptive Server, one additional line for each interface you want Adaptive Server to listen on).
- *rmaxuconn* – maximum value for user connections. It is calculated by: *@@max_connections* = *kimaxfd* - *ksnmstprts* - 4 (5 on OpenVMS). The pool of four is reserved for:

  stdout
  error log
  console
  dump/load device
  (debug on OpenVMS)

5 Compare the appropriate value for sockets to the current value for sockets, as determined in step 1. (Remember that the maximum number of virtual sockets = number of user connections + number of remote sites + 3.) If the current value is too low, use sp_configure to increase the value of number of user connections and/or number of remote sites so that the calculated value for sockets will be correct. This will probably correct the error. If the error still occurs, call Sybase Technical Support.

Note Under UNIX, you may need to raise the maximum number of file descriptors at the operating system level for the Adaptive Server if the new value for *cnsocket* exceeds the maximum number of per process file descriptors times the number of online engines for 11.0.x and later versions.
In an OpenVMS environment, if you increase the value for number of user connections, you may also need to increase the value of the OpenVMS operating system parameter *channelcnt* and recalculate your system quotas.
Refer to *Managing and Monitoring Sybase Adaptive Server Enterprise Using Sybase Central* for information about managing the maximum number of user connections available on an NT Server.

Additional information

Refer to “Setting Configuration Parameters” in the *System Administration Guide* for information about configuration parameters. Refer to the *Reference Manual* for information about sp_configure.

For parameter information specific to your platform, refer to the Adaptive Server installation and configuration guide.

If you need to call Technical Support, have the following information available:

- Server release and SWR rollup level
- Server error log
- Operating system version and error log
- Output of dbcc resource and sp_configure
- Text of all error messages

Version in which this error is raised

All versions

Error 1608

Severity 18

Message text A client process exited abnormally, or a network error was encountered. Unless other errors occurred, continue processing normally.

Explanation

This error occurs when a client process stops without informing Adaptive Server and a subsequent attempt by Adaptive Server to send results to this client fails.

Error 1608 is usually not serious and may be viewed as notification that a client process no longer exists. Some of the possible reasons a client process can disappear are:

- The client application was killed or ended abnormally
- The client machine was powered off or was restarted
- There is excessive network traffic
- The network has been shut down

Note Using Ctrl-c (on non-PC operating systems) to break out of an application should not cause this error.

Action

Do not be concerned if this error occurs only sporadically. However, if the error occurs frequently and continuously, or frequently for short periods of time, it may be a result of excessive network traffic or network problems.

If you suspect network problems, have your network administrator run diagnostics on the network between the client and the Adaptive Server machines to isolate the problem.

Version in which this error is raised All versions

Error 1613

Severity 21

Message text Could not close network %d connection for server process %d.

Explanation

This error occurs when Adaptive Server is unable to close a network connection for a Server process.

Error 1613 can be caused by:

- System parameters being set too low.
- Network problems.

This is not a serious problem unless it occurs frequently.

Note This error may be caused by a hardware problem.

Action

Check the values for your system parameters that affect network use:

- Under OpenVMS, these might include *ASTLM*, *FILLM*, *BYTLM*, and *CHANNELCNT*.
- Under UNIX, these might include the limit on the maximum number of file descriptors (limits the number of sockets that can be open at one time) and the per process limit on file descriptors.

If you suspect network problems, have your network administrator run diagnostics on the network between the client and the Adaptive Server machines to isolate the problem.

Version in which this error is raised

All versions

Error 1621

Severity

18

Message text

Type '%c' not allowed before login.

Explanation

Adaptive Server uses TDS (Tabular Data Stream) packets to communicate with clients. A token is attached to every TDS packet to differentiate between different types of packets. During login, Adaptive Server expects a client to supply either the TDS_LOGIN or TDS_ECHO token.

Error 1621 occurs when a client sends a token other than TDS_LOGIN or TDS_ECHO at login. Possible causes include:

- Non-printing characters in the login script.
- Another process trying to access the port used by the server.
- An incompatible database driver; for example, an unsupported ODBC driver.

Action

To see if other processes are interfering with the port, use netstat (UNIX) or Task Manager (NT) and check port status. You can also start the server using a different port number to remove any possible interference.

Check client applications to make sure logins are using a supported database driver. On NT, you can use the Drivers tab under ODBC Data Source Administrator to see the available drivers.

Check the TDS login packets being sent to Adaptive Server. To identify the affected logins, you can monitor packets:

- by using trace flags 3605 and 4001 (3605 sends the output to the error log)

  Warning! Use of trace flag 4001 may generate a large number of messages in the error log.

- with the ribo utility (Version 12.0).

Correct your login scripts to remove non-printing characters or control characters, and retry the connection.

Additional information

Refer to "How to Start Adaptive Server with Trace Flags" in Chapter 2 for details on using trace flags.

Version in which this error is raised

All versions

Error 1622

Severity

18

Message text

Type '%c' not implemented.

Explanation

Adaptive Server uses TDS (Tabular Data Stream) packets during communication with clients to differentiate between different types of packets. A token is attached to every TDS packet. Error 1622 occurs when a client sends an invalid TDS token.

Valid packet header tokens include:

- TDS_LANGUAGE
- TDS_RPC
- TDS_ECHO
- TDS_LOGOUT

In the message output, “%c” is the token name causing the 1622 error.

Error 1622 occurs with the following states:

State	Meaning
1	Error 1622 occurs with State 1 if the packet header cannot be converted to one of the tokens listed above.
2	Error 1622 occurs with State 2 if the header token is correct but the data token is bad.

Action

Check your network to make sure there are no unexpected characters being sent to Adaptive Server because of noise on the line or hardware problems. If you have confirmed that unexpected characters are not occurring, call Sybase Technical Support.

Version in which this error is raised

All versions

Error 1623

Severity

21

Message text

Warning: dataserver -p option fails. Please check for state '%d' to see where it fails.

Explanation

This error occurs when Adaptive Server is unable to create a new password for a user with “sso” role in response to the dataserver command with the -p option.

There are three possible states:

State	Meaning
1	Login passed with the -p option does not exist in syslogins.
2	Login exists but does not have the “sso” (system security officer) role.
3	Call to the routine that sets the passwords failed. This routine generates its own errors to help you diagnose the problem.

Action

The actions for recovering from Error 1623 are listed below, by state.

States 1 and 2

If you did not remove “sso” role from the “sa” login, use “sa” in your dataserver command in your runserver file. If you did remove “sso” role from the “sa” login, choose the appropriate option below.

- If you can, log into Adaptive Server as any user.

Determine which login names have “sso” role:

```
1> use master
2> go
1> select l.name, s.name from sysloginroles slr,
2> syslogins l,syssrvroles s
3> where slr.srid = 1  and slr.suid = l.suid and
4> slr.srid = s.srid
5> go
```

- If you cannot log into Adaptive Server as any user, and you have a historical dump of the master database that includes a user with the “sso” role and for which you know the password, create a new master device, load that dump and determine whether you can log into it. Refer to Chapter 1, “master Database Is Corrupt and Adaptive Server Starts” and read “Valid Dump of the master Database Exists”for instructions on how to re-create your master database.
- If you cannot log into Adaptive Server as any user and do not have a historical dump of the master database, refer to Chapter 1, “master Database Is Corrupt and Adaptive Server Starts” and read “Valid Dump of the master Database Does Not Exist” for instructions on how to re-create your master database.

State 3

Other error messages will be listed. Follow the directions in this manual for recovering from those errors. If no directions exist, call Sybase Technical Support.

Version in which this error is raised

All versions

Create Utilities Errors

This section contains error messages for Adaptive Server create operations (create table, create database, and so on).

Error 1702

Severity

16

Message text

```
Create table failed because column '%.*s' in table
'%.*s' exceeds the maximum of 250 columns.
```

Explanation

This error occurs when you try to create a table with more than 250 columns. The maximum number of definable columns per table is 250.

Note Column names must follow the rules for identifiers. They must be unique within a given table, but you can use the same column name in different tables in the same database.

Action

Reduce the number of columns in the table you are trying to create to 250 or less. Refer to “create table” in the *Reference Manual* for details.

Version in which this error is raised

All versions

Error 1732

Severity

20

Message text

```
Failed to reserve/initialize allocation hints space in
table '%.*s'.
```

Explanation

Each table, index, and text chain has one or more Object Allocation Map (OAM) pages allocated to the table or index. These OAM pages store pointers to each allocation unit that contains pages for the object. The object's first OAM page stores allocation hints, indicating which OAM page has information about allocation units with free space.

Error 1732 is raised when Adaptive Server is allocating pages for a new table, but is unable to create allocation hints space for the table.

Error 1732 occurs with the following states:

State	Meaning
1	Unable to create allocation hints space for the table.
2	Unable to initialize allocation hints space for the table.
3	Unable to create allocation hints space for a text page chain.
4	Unable to initialize allocation hints space for a text page chain.

The table named in the message may be a user table, or a temporary table the server creates when executing a stored procedure or other query. Error 1732 can be due to an Adaptive Server problem.

Action

This is not a serious error. Reconnect to the server and retry the query which raised the error; it should run normally.

Contact Sybase Technical Support if the problem persists.

Additional information

Have the following information ready before calling Sybase Technical Support:

- Server version and SWR rollup level
- Text of all error messages.

Version in which this error is raised

All versions

Error 1739

Severity

10

Message text

```
Warning: Parameter 'exp_row_size' is not applicable to
table '%.*s' in allpages locking scheme
```

Explanation

As of Version 11.9, Adaptive Server provides two new locking schemes:

- Datapages locking, which locks only the data pages
- Datarows locking, which locks only the data rows.

Since neither scheme locks index pages, they are often referred to together as the data-only locking scheme.

The pre-11.9 locking scheme, known as allpages locking, continues to be available with 11.9. This scheme locks the data and index pages affected by queries, and is the default locking scheme.

exp_row_size is a space management parameter that you can use in 11.9, when a data-only locked table contains variable length columns and the row size is expected to grow with updates.

Error 1739 is raised:

- when you attempt to specify an expected row size using the exp_row_size parameter in a create table or alter table command for an allpages-locked table. This is just an informational message and the command is executed.
- when you attempt to specify an expected row size using the exp_row_size parameter for a table with fixed-length rows. The command fails.

Action

Retry the command if necessary. Do not specify exp_row_size on an allpages-locked table. Use this parameter only for tables that use the data-only-locked scheme and contain variable length rows.

Version in which this error is raised

11.9 and later

Error 1740

Severity

10

Message text

```
Warning: Parameter 'max_rows_per_page' is not
applicable to table '%.*s' in 'datapages' or 'datarows'
locking scheme
```

Explanation

As of Version 11.9, Adaptive Server provides two new locking schemes:

- Datapages locking, which locks only the data pages
- Datarows locking, which locks only the data rows.

Since neither scheme locks index pages, they are often referred to together as the data-only locking scheme.

The pre-11.9 locking scheme, known as allpages locking, continues to be available with 11.9. This scheme locks the data and index pages affected by queries, and is the default locking scheme.

max_rows_per_page is a parameter you can use in the create table, create index, or alter table commands to limit the number of rows Adaptive Server can put on each page. This parameter can help reduce lock contention on frequently accessed tables that use the allpages locking scheme.

Error 1740 is an informational message and is raised when you attempt to specify the max_rows_per_page parameter for a table that uses data-only locking.

Action

Do not specify max_rows_per_page on a data-only-locked table. Use this parameter only for tables that use the allpages-locked scheme.

Version in which this error is raised

11.9 and later

Error 1803

Severity

17

Message text

```
CREATE DATABASE failed. Could not allocate enough disk
space for a new database on the disks named in the
command. Total space allocated must be at least %d
Mbytes (%ld 2048-byte pages) to accommodate copy of
Model Database.
```

Explanation

When you issue a create database command, Adaptive Server:

- Assigns space to the database on the specified database devices. If you use the default keyword or if you omit the on clause altogether, Adaptive Server puts the database on one or more of the default database devices specified in master..sysdevices.
- Makes a copy of the model database in the new database space, creating the new database's system tables. The new database thus inherits all the changes you have made to the model database, including the addition of user names and objects.

Error 1803 is raised when Adaptive Server is unable to allocate the space required for the model database on the specified device(s).

Action

Make sure that the available space on the database device you specified (or the available space on the default device) will accommodate the model database.

To check the size of model, use the following command:

```
1> sp_helpdb model
2> go
```

To check the space available on the device, first use the following commands to check the space already allocated to other databases on the device:

```
1> use master
2> go
```

```
1> select dbid, size, phyname "physical device"
2> from sysusages, sysdevices
3> where name = 'device_name'
4> and vstart between low and high
5> compute sum(size)
6> go
```

For example:

```
1> select dbid, size, phyname "physical device"
2> from sysusages, sysdevices
3> where name = 'sd5f'
4> and vstart between low and high
5> compute sum(size)
6> go

dbid   size    physical device
------ ------- ---------------
15       17920 /dev/rsd5f
16       20480 /dev/rsd5f
17        7680 /dev/rsd5f
18       20480 /dev/rsd5f
21        5120 /dev/rsd5f
      sum
      =======
        71680
```

Subtract the sum from the total space on your physical device to determine the available space on the database device.

Additional information

Refer to “model Database” in the *System Administration Guide* and “create database” in the *Reference Manual* for more information.

Version in which this error is raised

All versions

Error 1808

Severity

21

Message text

```
Crdb_disk: Getnext SCAN_NOINDEX on
sysdevices.status=DEFAULT failed to find default rows
```

Explanation

This error occurs during execution of a create or alter database command when you do not specify a device and no disk device is defined as the default disk for an Adaptive Server.

By default, the master device is defined as a default disk. If you have used the stored procedure sp_diskdefault to turn off default for the master device:

```
1> sp_diskdefault master, defaultoff
2> go
```

and you have not used sp_diskdefault to specify another default disk device:

```
1> sp_diskdefault  dev_name, defaulton
2> go
```

(where *dev_name* is the new default disk device) then you will not have any devices defined as default devices and Error 1808 will occur when you run create or alter database.

Action

Select one of the following options:

- Always use the on dev_name clause with create and alter database statements.
- Follow these steps:
  a Use one of the methods below to determine whether you have any disk devices defined as default.
  b If you do not have any disk devices defined as default, use sp_diskdefault to specify a device as a default disk device.

Method 1

Type:

```
1> sp_helpdevice
2> go
```

If you see a line such as:

```
device_name  physical_name  description
------------------------------------------------------------------------
master       d_master       special,default disk,physical disk,17.00MB
```

where “default disk” is listed, then that device will be used by create or alter database when the on dev_name clause is omitted for allocation for the database. If no entries have “default disk” in the description field, then you do not have any disk devices defined as default.

Method 2

You can determine whether you have any disk devices defined as default by checking whether the first bit of the status column in sysdevices is turned on for any devices:

```
                    1> select status, name from sysdevices
                    2> where status & 1 = 1
                    3> go
 status name
------ ------------------------------
     3 master
```

Additional information — Refer to the *Reference Manual* for information about sp_diskdefault.

Version in which this error is raised — All versions

Error 1809

Severity — 14

Message text —

```
CREATE DATABASE must be preceded by a 'USE master'
command. Check with your DBO <or a user with System
Administrator (SA) role> if you do not have permission
to USE master.
```

Explanation — This error occurs when you try to create a new database without being in the master database.

Action — Before issuing a create database command, be sure you are in master:

```
1> use master
2> go
```

If you do not have the necessary permissions to create a database, ask your Sybase System Administrator to grant you permission to use the create database command.

Version in which this error is raised — All versions

Error 1810

Severity — 16

Message text —

```
CREATE DATABASE failed because of incorrect size
parameter(s). Total number of megabytes specified must
be at least %d megabytes so that the Model Database can
be copied to the new database.
```

Explanation

When a create database command is issued, Adaptive Server makes a copy of the model database, which contains the system tables needed by each database. Error 1810 occurs when you try to create a database of a size smaller than the size of the model database. The default size of model is 2MB.

Action

Make sure the size you plan to use in your create database command is at least as large as the model database (combine the data space size and the log space size).

To check the size of model , use the following command:

```
1> sp_helpdb model
2> go
```

Additional information

Refer to “model Database” in the *System Administration Guide* and “create database” in the *Reference Manual* for more information.

Version in which this error is raised

All versions

Error 1811

Severity

16

Message text

```
'%.*s' is the wrong type of device for CREATE or ALTER
database. Please check sysdevices. The CREATE or ALTER
is aborted.
```

Explanation

This error occurs when a dump device is specified in a create or alter database command. You cannot create or alter a database on a dump device.

Action

Change the command to specify a device with a valid type.

1 Obtain a list of all of the devices recognized by your Adaptive Server as follows:

```
1> sp_helpdevice
2> go
```

2 Check the description field of the sp_helpdevice output to identify which devices have the correct type. Any device with a value in the description field other than “dump device” can be used.

Additional information

Refer to “Initializing Database Devices” in the *System Administration Guide* for information about devices.

Version in which this error is raised

All versions

Error 1813

Severity 16

Message text Cannot open new database '%.*s'. CREATE DATABASE is aborted.

Explanation During a create database command, Adaptive Server opens the new database at several stages, first after inserting the rows into the system catalogs and later to initialize the log segment. Error 1813 occurs when Adaptive Server is unable to open the new database or initialize the log segment for the new database during a create database command.

This error can be caused by the following:

- Uncleared allocation buffers. This can occur after a database has been dropped and then a create database command is issued.
- Corrupt system catalogs such as sysdatabases.
- An Adaptive Server problem.

Action If Error 1813 occurs just after a database has been dropped, shut down and restart Adaptive Server to clear the allocation buffers and then re-enter your create database command.

If this does not clear the problem or you have not just dropped a database, call Sybase Technical Support.

Additional information Have the following information available when you call Sybase Technical Support:

- Adaptive Server release and SWR Rollup level
- Adaptive Server error log
- Text of all error messages
- dbcc checkcatalog output for master

Version in which this error is raised All versions

Error 1820

Severity 20

Message text This command adds %S_MSG space to disk '%.*s', which previously contained only %S_MSG. You must specify WITH

```
OVERRIDE to force this allocation.
```

Explanation

When a database is being extended or a new database created, Adaptive Server allocates space on the specified device (if no device is named, it uses the default device). Error 1820 is raised when allocating space for the database would convert the device from a dedicated device - one that stores only data or log - into a mixed-use device that contains both data and log on the same segment.

The error may be detected:

- when you use the alter database or create database commands
- when proxy databases are being created for the high availability (HA) feature (version 12.0 and higher).

Error 1820 is raised with the following states:

State	Meaning
1	Cannot allocate log space on a device dedicated to data. Device was named in the create database/alter database command.
2	Cannot allocate data space on a device dedicated to the log. Device was named in the create database/alter database command.
3	Cannot allocate log space on a device dedicated to data. Device was not explicitly specified. This state is also raised during HA proxy database creation (12.0 and higher).
4	Cannot allocate data space on a device dedicated to the log. Device was not explicitly specified. This state is also raised during HA proxy database creation (12.0 and higher).

Action

You can avoid the 1820 error by using the with override clause when creating or altering a database.

Warning! Using the with override clause to resolve the 1820 error is not advisable. It will mix data and transaction log on the same device, making it difficult to achieve up-to-the-minute recoverability for your database.

Additional information

For details about the management and recovery of the transaction log, see "Placing the Transaction Log on a Separate Device" in the *System Administration Guide*.

Version in which this error is raised

11.0.3 and later

Error 1902

Severity

16

Message text

```
Cannot create more than one clustered index on table
'%.*s'. Drop the existing clustered index '%.*s' before
creating another
```

Explanation

Adaptive Server uses a clustered index to sort rows so that their physical order is the same as their logical (indexed) order. The bottom or leaf level of a clustered index contains the actual data pages of the table.

In a nonclustered index, the physical order of the rows is not the same as the indexed order. The leaf level of a nonclustered index contains pointers to rows on data pages.

Many nonclustered indexes are allowed on a table, but only one clustered index per table is allowed. Error 1902 occurs when you try to create a second clustered index on the table named in the error message.

Note Using the primary key clause in a create table statement creates a unique clustered index.

Action

1 List the existing indexes on the table:

```
1> use database_name
2> go

1> sp_helpindex object_name
2> go
```

where *database_name* is the database where the table named in the error message resides and *object_name* is the name of the table in the error message. The index_description column of the output shows whether the index is clustered or nonclustered.

2 If a clustered index already exists, drop it.

3 Create a new clustered index.

Version in which this error is raised

All versions

Error 1903

Severity

16

Message text

```
%d is the maximum allowable size of an index. Composite
index specified is %d bytes.
```

Explanation

A composite index is any index that uses from 2 to 16 columns. Error 1903 occurs when the sum of the lengths of all the columns used in a composite index exceeds the allowable limits.

Following is an example of a composite index using objects from the pubs2 sample database:

```
1> create index example
2> on authors (phone, state, postalcode)
3> go
```

phone is defined as char(12), state as char(2), and postalcode as char(10). Since each char takes 1 byte of storage, the total length of this composite index is 24 bytes.

Action

Examine the lengths of the columns in your composite index and make sure that the lengths total:

- no more than 255 bytes (11.0.x)
- no more than 600 bytes (11.5 and higher).

Refer to “System Datatypes” and “User-Defined Datatypes” in the *Adaptive Server Enterprise Transact-SQL User's Guide* for information about how to calculate the storage size for different Adaptive Server datatypes.

Additional information

Refer to “create index” in the *Reference Manual* for information about indexes.

Version in which this error is raised

All versions

Error 1904

Severity

```
16
```

Message text

```
Cannot specify more than %d column names for index key
list. %d specified.
```

Explanation

A composite index is any index that uses more than one column. Error 1904 occurs when:

- more than 16 columns are combined in a single composite index (11.5 and earlier)
- more than 31 columns are combined in a single composite index (11.9.2)

Action

Change your create index statement to include no more than the maximum number of columns allowed.

Additional information

Refer to "create index" in the *Reference Manual* and composite indexes in the *Transact-SQL User's Guide* for more information.

Version in which this error is raised

All versions

Error 1916

Severity

16

Message text

```
CREATE INDEX options %s and %s are mutually exclusive.
```

Explanation

This error occurs when you try to create an index with two mutually exclusive create index options. When Error 1916 occurs, no index is created.

The following options are not compatible with each other:

- ignore_dup_row
- allow_dup_row
- ignore_dup_key

The following table shows which options can be used with which type of indexes:

Table 1-6: Compatibility chart for index types and index options

Type of Index	Index Option		
	ignore_dup_key	ignore_dup_row	allow_dup_row
Nonunique Nonclustered	Don't use	Don't use	Don't use
Nonunique Clustered	Don't use	Okay	Okay
Unique Nonclustered	Okay	Don't use	Don't use
Unique Clustered	Okay	Don't use	Don't use

Action

- Use only one of the options.
- Use options appropriate for the type of index being created.

Additional information

Refer to "create index" in the *Reference Manual* for information about indexes.

Version in which this error is raised	All versions

Error 1928

Severity	16
Message text	`Cannot create clustered index on '%.*s' because it is partitioned.`
Explanation	By default, Adaptive Server stores a heap table's data in one doubly linked chain of database pages. Adaptive Server inserts all new rows into the last page of the chain. A transaction holds an exclusive lock on the last page while inserting new rows, which can block other, concurrent transactions from inserting into the table. The partition clause of the alter table command allows you to partition user tables that do not have a clustered index. Partitioning creates additional page chains on the table, each with its own last page. This reduces page contention for concurrent inserts, and can also reduce I/O contention if the table exists on user-defined segments and is distributed over multiple physical devices. Since clustered indexes are not allowed on partitioned tables, Error 1928 occurs when you try to create a clustered index on a partitioned table.

Action	If you want to create a clustered index on the table, use the unpartition clause of the alter table command to concatenate all partitions:
	
	`1> alter table table_name unpartition`
	`2> go`
	
	Then create your clustered index.

Additional information	Refer to the *Reference Manual* for information about the alter table command.
Version in which this error is raised	All versions

Error 2110

Severity	20
Message text	`The proccreate() function is trying to put too many lock`

```
requests in the lock_requests[] array.
```

Explanation

When you create and modify triggers (and other stored objects such as procedures, views, and rules), Adaptive Server uses the proccreate function to update the appropriate system tables. As an example, Adaptive Server uses this function to write the text of a trigger into the syscomments table. Depending on the action requested, proccreate must acquire locks on the objects being created and any objects affected by the created object.

For example, since a table can only have one trigger for each operation (create, **update** or delete), a new trigger on a table for the same operation overwrites any previous trigger for that operation. proccreate must be able to lock the new trigger as well as the existing trigger.

Error 2110 is raised when you attempt to create a new trigger or other stored object on a table, or when you attempt to overwrite an existing trigger, and proccreate is unable to acquire the necessary locks on the objects in question.

State	Meaning
1	Could not acquire a lock for a trigger or other object being created.
2	Could not lock an existing delete trigger on this object.
3	Could not lock an existing insert trigger on this object.
4	Could not lock an existing update trigger on this object.
5	Could not lock the table to which the trigger is being attached.

Error 2110 breaks your connection to Adaptive Server.

Action

Take the following steps to correct the problem.

1 Retry the object creation.

2 If the error was raised with State 5, and the target table is newly created, check your trigger creation code to be sure that the table's CREATE statement is committed before attempting trigger creation.

3 If the above steps do not resolve the problem, Error 2110 may be due to an Adaptive Server problem. Contact Sybase Technical Support for assistance.

Additional information

When calling Technical support, have the following information available:

- Server release and SWR Rollup level
- Server error log
- Text of all error messages
- Text of the trigger that raises the error

Version in which this error is raised	All versions

High Availability Utility Errors

This section contains error message writeups for Adaptive Server's High Availability utility.

Error 2243

Severity

16

Message text

```
QUIESCE DATABASE failed with error %ld. See the error
log for more detail.
```

Explanation

Popular techniques for backing up an Adaptive Server database, without the need for a full dump and load cycle, include unlinking a mirror device and using file system backup utilities outside the server. In Version 12.0, quiesce database allows you to unlink mirrors or back up devices at the operating system level without having to first shut down the server. quiesce database hold establishes a quiesce point and prevents updates in the specified database(s). quiesce database release allows transactions to resume once the device is unlinked or the external backup completes. quiesce database hold creates a child process, known as a quiesce db agent, to carry out the tasks necessary to establish the quiesce point.

Error 2243 is raised when you issue quiesce database hold, but the quiesce db agent detects a condition which prevents it from creating the quiesce point. Possible causes include:

- dump database was not executed at least once on the database prior to issuing quiesce database hold
- dump transaction has been executed on the database since the last dump database;
- the timeout limit was reached while the agent waited for transactions in PREPARE state to finish.

Action

Note the error number specified in the 2243 message text. Locate this error in the server error log to determine the cause of the problem. Take the appropriate action, such as issuing a dump database, to resolve the error.

Version in which this error is raised

12.0 and later

Character Set Conversion Errors

This section contains error messages for Adaptive Server character set conversion.

Error 2401

Severity

11

Message text

```
Character set conversion is not available between client
character set '%.*s' and server character set '%.*s'.
```

Explanation

Character set conversion is the change of the encoding scheme of a set of characters on the way into or out of Adaptive Server. Conversion is used when Adaptive Server and a client communicating with it use different character sets. For example, if Adaptive Server uses character set ISO 8859-1 (iso_1) and a client uses Code Page 850, character set conversion must be turned on so that both Adaptive Server and client interpret data passing back and forth in the same way.

Character set conversion is only available for certain character sets. Valid conversion matches are listed in the *System Administration Guide*. Error 2401 occurs when an invalid character set conversion is requested (for example, ascii_8 with ISO 8859-1). When this error is raised, no conversion is performed.

Error 2401 can occur in the following situations:

- When a client requests a connection, Adaptive Server checks whether it can convert from the client's character set to its own character set. If it cannot do the requested conversion because conversion is not available between the two character sets, Error 2401 is raised.
- If you use the set char_convert command to start conversion between the Adaptive Server character set and a different character set and conversion for those character sets is not available, Error 2401 is raised.

Action

Refer to “Configuring Client/Server Character Set Conversions” in the *System Administration Guide*. and check the table of valid conversion matches under “Conversion Paths Supported”. Make sure you are not requesting conversion involving an incompatible character set.

Additional information

Refer to “Configuring Client/Server Character Set Conversions” in the *System Administration Guide* for details about character set conversion.

Version in which this error is raised: All versions

Error 2402

Severity: 16

Message text:

```
Error converting client characters into server's
character set. Some character(s) could not be converted.
```

Explanation

This error occurs during insertion of data (insert or bcp) when Adaptive Server fails to convert a character to the required character set.

Error 2402 usually occurs for one of the following reasons:

- The character exists in the client character set but it does not exist in the Adaptive Server character set.
- The character exists in both the client and the Adaptive Server character set, but is represented by a different number of bytes in the client character set than in the Adaptive Server character set.

This error occurs during normal processing and it prevents query execution.

Action

The following options are available for recovering from Error 2402.

Change Your Data

Modify the incoming data so that it contains characters recognizable by Adaptive Server.

Turn Off Character Set Conversion

If the error occurs while you are using isql, bcp, or defncopy, you can use the -J (UNIX and PC) or /clientcharset (OpenVMS) command-line option with no character set name to set the client's character set to NULL. If you use this command-line option without specifying a character set name, no conversion takes place and no error message is sent. As a result, some characters sent by the client to the Adaptive Server may not be interpreted correctly by the Adaptive Server and vice versa. (If only 7-bit characters are being handled, no incorrect interpretation will take place.)

Otherwise, you can turn off conversion so that characters are sent and received unchanged with the following command:

```
1> set char_convert off
2> go
```

Turn Off Character Set Conversion Error Reporting

You can turn off the printing of error messages with the following command:

```
1> set char_convert on with no_error
2> go
```

Bytes which cannot be converted are replaced with an ASCII question mark ("?").

Additional information

Refer to "Converting Character Sets Between Adaptive Server and Clients" in the *System Administration Guide* for details about character set conversion.

Version in which this error is raised

All versions

Error 2409

Severity

11

Message text

```
Cannot find the requested character set in Syscharsets:
name = '%.*s'.
```

Explanation

Information about the character sets that are defined for use by Adaptive Server is stored in syscharsets. Error 2409 occurs when Adaptive Server cannot find a character set needed for character set conversion.

Error 2409 can occur in the following situations:

- When a client requests a connection, Adaptive Server checks whether it can convert from the client's character set to its own character set. If Adaptive Server cannot do the requested conversion because a character set is not found, Error 2409 is raised.
- If you use the set char_convert command to start conversion between the Adaptive Server character set and a different character set and the character set you specify does not exist in syscharsets, Error 2409 is raised.

Action

Make sure all necessary character sets are loaded, including the client's character set (as shown in the error message output):

```
1> use master
2> go
1> select csid, id, name from syscharsets
2> go
 csid id  name
 ---- --- ------------------------------
```

```
0   0 ascii_8
0   1 iso_1
1  50 bin_iso_1
```

If you are missing needed character sets, follow the directions in the Adaptive Server installation and configuration guide to load them using the sybinit utility.

Standalone Utilities

By default, for standalone utilities, the character set used for the client is the default for the platform. This may not be the correct character set. If the 2409 error is occurring while you are using isql, bcp, or defncopy, use the -J (UNIX and NT) or /clientcharset (OpenVMS) command-line option to specify the client character set.

Open Client Applications

For Open Client applications such as DB-Library, you can specify the character set for the client through the login packet. Refer to the reference manual for the application and the *Open Client/Server Supplement* for your platform for details.

Additional information

Refer to "Configuring Client/Server Character Set Conversions" in the *System Administration Guide* for details about character set conversion.

Version in which this error is raised

All versions

dbcc Errors

This section contains error messages for dbcc (database consistency checker) commands.

Messages that ordinarily have severities greater than 16 will show severity 16 when they are raised by dbcc checktable and dbcc checkalloc so that checks will continue with the next object.

Error 2501

Severity

16

Message text

```
Table named %.*s not found; check sysobjects.
```

Explanation

This error occurs when a dbcc command is run if Adaptive Server is unable to find the table name (or ID) in sysobjects.

Error 2501 can occur when running any of the following dbcc commands:

- dbcc checktable
- dbcc checkdb
- dbcc tablealloc
- dbcc reindex
- dbcc fix_text

This error can be serious as it indicates possible corruption.

Action

If this error occurs on a command for which you have specified a table name, check to make sure you have spelled the table name correctly:

```
1> select * from sysobjects
2> go
```

If this error occurs because the entry for an object has been removed from sysobjects but the object is still referenced by some other system tables, call Sybase Technical Support. They will help you try to remove the references to that object from the other system tables.

Additional information

Before calling Technical Support, have the following information available:

- Adaptive Server version and SWR version level
- Adaptive Server error log
- Output of dbcc checkalloc and dbcc checkdb

- Text of all error messages

Version in which this error is raised

All versions

Error 2502

Severity

16

Message text

```
Table Corrupt: A page is linked in more than one chain;
check this page:page number=%ld allocation status=%d
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server detects an inconsistency in the page chains associated with an object. For example, a page may be connected in more than one chain. If this error is detected by a dbcc command, a 605 error may also occur.

Warning! If Error 605 occurs with this error, use the instructions presented under Error 605 or call Technical Support.

This is a serious error that must be corrected immediately.

Action

1 Note the page number in the error message. Use the procedure in Chapter 2, "How to Find an Object Name from a Page Number" to determine which table is associated with the corrupt page chain.

2 If the object with the error is *not* a system table (a system table's object ID is less than 100), continue with step 3.

If the object with the error is a system table and the index ID is *not* 0, refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index.

If the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore from clean backups.

3 If the object with the error is *not* a system table, follow one of the procedures below to correct the problem, depending on what kind of object is involved:

- If the page belongs to an index, drop and re-create the index.

- If the page does not belong to an index, recover your database from a known clean backup or recover the tables not affected by this error via bcp. Refer to Chapter 2, "How to Rescue Data from a Corrupted Table" for details.

4 Examine your operating system error log and the Adaptive Server error log to determine if hardware errors may have caused the corruption.

Additional information

Before calling Technical Support, have the following information available:

- Adaptive Server version and SWR version level
- Server error log
- Text of all error messages
- Output of dbcc page

Version in which this error is raised

All versions

Error 2503

Severity

16

Message text

```
Table Corrupt: Page linkage is not consistent; check the
following pages: (current page#=%ld; page# pointing to
this page=%ld; previous page# indicated in this
page=%ld)
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server detects an inconsistency in the page linkage of one of the page chains associated with a table. There is one doubly linked page chain for the table's data as well as one for each index level. For example, a page's "next" pointer points to a page whose "previous" pointer points to a different page.

This is a serious error that must be corrected immediately.

Action

1 Note the first page number in the error message. Use the procedure in Chapter 2, "How to Find an Object Name from a Page Number" to identify which object is associated with the "current page".

2 If the object with the error is *not* a system table (a system table's object ID is less than 100), continue with step 3.

   If the object with the error is a system table and the index ID is *not* 0, refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index.

   If the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore from clean backups.

3 If the object is *not* a system table, perform one of the actions below, depending upon what kind of object was involved in the error.

Current Page Belongs to a Nonclustered Index

If the "current page" belongs to a nonclustered index (index ID is between 2 and 250), run dbcc checktable to verify that the data page linkage is good. If checktable runs cleanly, drop and re-create the index.

Current Page Belongs to a Clustered Index

If the "current page" belongs to a clustered index (index ID is 1), drop all the indexes on that table, re-create the clustered index first, and then re-create the remaining nonclustered indexes.

Current Page Does Not Belong to an Index

If the "current page" does not belong to an index at all (index ID is 0 or 255) and there were no errors other than the 2503 error, perform the following steps:

1 Use bcp or select into to retrieve your data from the affected table. Refer to Chapter 2, "How to Rescue Data from a Corrupted Table" for instructions.

2 If errors still exist, restore from clean backups or call Sybase Technical Support.

3 Examine your operating system error log and the Adaptive Server error log to determine if hardware errors may have caused the corruption.

Additional information

For information on locating device fragments, refer to Chapter 2, "How to Determine Which Physical Devices a Database is On".

Version in which this error is raised

All versions

Error 2506

Severity

16

Message text

```
Table Corrupt: The values in adjust table should be in
ascending order starting from the end of the table
(page#=%ld row#=%d); check adjust table in this row
```

Note This error may be caused by a hardware problem.

Explanation

An adjust table is stored at the end of every data or index row which has variable length columns. It is used to improve performance of Adaptive Server for rows with lengths greater than 256 bytes.

Error 2506 occurs when dbcc checkdb or dbcc checktable detects inconsistencies in the values stored in the adjust table of the row specified by the error message.

This error can occur on a data page or an index page.

Action

If you have clean backups of your database, recovery from backups may be your best option. If you do not have clean backups of your database, refer to Chapter 2, "How to Identify and Fix a Corrupted Table" to correct the error.

Warning! Some data on this page might be lost if you recover your table using bcp or select into (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

Before dumping your database, make sure it is working correctly. More specifically, the following commands should be run prior to each dump:

1 dbcc checkdb.

2 dbcc checkalloc or dbcc checkalloc with the fix option. (Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about running these commands in multi-user mode and how to prevent spurious allocation errors from dbcc commands.)

Hardware Errors

Hardware errors on your database devices can cause Error 2506. To help rule out this possibility, do the following:

- Check the Adaptive Server error log to determine whether there are other indications of hardware problems, such as kernel messages reporting I/O errors.
- Check the hardware error log or diagnostics utilities for I/O errors. Refer to Chapter 2, "Checking the Operating System Error Log".

If problems persist, refer to the error documentation in this manual or contact Sybase Technical Support.

Additional information

Refer to bcp in the Adaptive Server utility programs manual.

Refer to "select" in the *Reference Manual* for information about the select into command.

Have the following information available before calling Sybase Technical Support:

- Adaptive Server version and SWR version level
- Adaptive Server error log
- Operating system error log
- Output of dbcc checkdb and dbcc checkalloc

Version in which this error is raised

All versions

Error 2507

Severity

16

Message text

```
Table Corrupt: Offset table is incorrect (page#=%ld
row#=%d offset in offset table is=%d correct offset
is=%d)
```

Note This error may be caused by a hardware problem.

Explanation

An offset table is stored at the end of every data page and indicates where rows are located on that page. Adaptive Server uses the offset table to quickly access the rows in each page.

Error 2507 occurs when dbcc checkdb or dbcc checktable detects inconsistencies in the offset table of a row on the page specified by the error message.

Action

Refer to Chapter 2, "How to Identify and Fix a Corrupted Table" to correct the error.

Warning! Some data on this page might be lost if you recover your table using bcp or select into (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

Before dumping your database, make sure it is working correctly. More specifically, the following commands should be run prior to each dump:

1 dbcc checkdb.

2 dbcc checkalloc or dbcc checkalloc with the fix option. (Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about running these commands in multi-user mode and how to prevent spurious allocation errors from dbcc commands.)

Hardware Errors

In addition to the above possible causes, hardware errors on your database devices can cause Error 2507. To rule out this possibility, do the following:

- Check the Adaptive Server error log to determine whether there are other indications of hardware problems, such as kernel messages reporting I/O errors.
- Check the hardware error log or diagnostics utilities for I/O errors.

If problems persist, refer to the error documentation in this manual or contact Sybase Technical Support.

Additional information

Refer to bcp in the Adaptive Server utility programs manual.

Refer to "select" in the *Reference Manual* for information about the select into command.

Have the following information available before calling Sybase Technical Support.

- Adaptive Server version and SWR version level
- Adaptive Server error log
- Operating system error log
- Text of all error messages

- Output from dbcc checkdb, dbcc checktable, and dbcc checkalloc

Version in which this error is raised

All versions

Error 2509

Severity

16

Message text

```
Table Corrupt: The row number and offset of each row in
the page should have a matching entry in row number
table; check this page (page#=%ld row#=%d offset in row
number table=%d)
```

Note This error may be caused by a hardware problem.

Explanation

The offsets for data or index rows are stored at the end of every page (in the row number table) and indicate where a certain row is located on that page.

Error 2509 occurs when the dbcc checkdb or dbcc checktable command detects that a row does not have an entry matching its offset (location) on the page in the row offset table.

Attempts to delete the offending row will result in Error 631, while attempts to select the offending row may be successful.

This problem is probably a result of a problem within Adaptive Server but may also be caused by one of the following:

- Hardware failure.
- Sybase System Administration problems.
- UNIX System Administration problems.

Action

First, make sure that you ruled out any of the above-mentioned causes of this error by referring to the appropriate sections in Chapter 2, "Encyclopedia of Tasks".

After you eliminated other more serious errors on this table, follow these steps to correct the 2509 error:

1 Follow the instructions in Chapter 2, "How to Find an Object Name from a Page Number" to identify which table and index correspond to the page number from the error message text.

2 If the object with the error is *not* a system table (object ID is more than 100), continue with step 3.

If the object with the error is a system table and the index ID is *not* 0, refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index.

If the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption or it may be necessary to restore from clean backups.

3 If the object with the error is a user table, use one of the following three methods to clear the 2509 error:

- Create a clustered index on the corrupted table. Creating a clustered index will copy the whole table onto new data pages, and will overwrite the row number table on each page. If a clustered index already exists on the table identified in step 1, drop the clustered index and re-create it.

  Warning! If you have other serious errors on this table (Errors 614, 2506, 2507, or 2524), eliminate these errors first, before you create a clustered index. Refer to the documentation in this manual for how to eliminate other errors on this table.

- Select the whole table into a new table, and drop the old table.
- Bulk copy the affected table out, drop and re-create the table, and bulk copy back in.

  Warning! Some data on this page might be lost if you recover your table using bcp or select into (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

For more information about how to copy a table in a new table or file, refer to Chapter 2, "How to Rescue Data from a Corrupted Table".

Additional information

Refer to create index and select in the *Reference Manual* and in the *Transact-SQL User's Guide*.

Refer to bcp in the Adaptive Server utility programs manual.

Version in which this error is raised

All versions

Error 2510

Severity 16

Message text

Version 11.0.2 and Later

```
Key mismatch between index page %ld (row %d) and data
page %ld in database '%.*s'. Drop and re-create index
id %d of table '%.*s'.
```

Version 11.0.1 and Earlier

```
Key mismatch between index page and data page in
database '%.*s'. Drop and re-create the index. (index
page %ld, row %d, data page %ld)
```

Note This error may be caused by a hardware problem.

Explanation

This message indicates that an index is inconsistent with the table data that it represents.

Action

Use the following procedure to recover from this error:

1 Record the index page number and the data page number from the error text. Refer to Chapter 2, "How to Find an Object Name from a Page Number" to identify which table and index correspond to the data and index page number. Record the object ID and index ID.

2 If the object with the error is *not* a system table (its object ID is 100 or greater), go to step 3.

If the object with the error is a system table, refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index.

3 If the object with the error is a user table, find the name of the index involved by executing the following query:

```
1> select name from sysindexes
2> where indid = index_id
3> and id = object_id
4> go
```

Drop and re-create the index. (Refer to the *Adaptive Server Reference Manual* for information about the drop index and create index commands.)

Run dbcc checktable and dbcc tablealloc on the affected table to verify that all problems have been corrected at the table level.

If problems persist, refer to the error documentation in this manual or contact Sybase Technical Support. Have the output from the appropriate dbcc commands available when you call.

Version in which this error is raised

All versions

Error 2511

Severity

16

Message text

Version 11.0.2 and Later

```
Keys of index id %d for table '%.*s' in %S_MSG %ld should
be in ascending order. Drop and re-create the index in
database '%.*s'.
```

Version 11.0.1 and Earlier

```
Keys of index id %d for table '%.*s' in %S_MSG should
be in ascending order. Drop and re-create the index.
(index page %ld)
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when an index is not ordered correctly.

There are three possible values of ”%S_MSG” from the error message text:

- Data page (only on clustered indexes)
- Index page
- Leaf page

Action

1 Note the index ID and table name displayed in the error message output.

2 If the object with the error is *not* a system table (object ID is more than 100), continue with step 3.

If the object with the error is a system table, refer to Chapter 2, “How to Fix a Corrupted Index on System Tables” for instructions on how to repair the system table index.

3 If the object with the error is a user table, find the name of the index involved by executing the following query:

```
1> select name from sysindexes
```

```
2> where indid = index_ID
3> and id = object_ID
4> go
```

Drop and re-create the index. To verify that all problems have been resolved at the table level, run the dbcc checktable and dbcc tablealloc commands on the affected table.

If problems persist, refer to the error documentation in this manual or contact Sybase Technical Support. Have the output from the appropriate dbcc commands available for their review.

Version in which this error is raised

All versions

Error 2513

Severity

16

Message text

```
Table Corrupt: Object id %ld (object name = %S_OBJID)
does not match between %.*s and %.*s
```

Explanation

This error occurs when dbcc checkcatalog finds that an object exists in one system table but not in another system table where Adaptive Server expected to find it. This is a serious error as it indicates corruption.

Error 2513 occurs with the following states:

State	Meaning
1	The object was found in syscolumns but not in either sysobjects or systypes.
2	The object (a view) was found in sysobjects but not in syscolumns.
3	The object (a view) was found in sysobjects but not in syscomments.
4	The object (a view) was found in sysobjects but not in sysprocedures.
5	The object (a stored procedure) was found in sysobjects but not in syscomments.
6	The object (a stored procedure) was found in sysobjects but not in sysprocedures.
7	The object (a rule or default) was found in sysobjects but not in syscomments.
8	The object (a rule or default) was found in sysobjects but not in sysprocedures.
9	The object (a table) was found in sysobjects but not in syscolumns.
10	The object (a table) was found in sysobjects but not in sysindexes.
11	The object (a table) was found in sysindexes but not in sysobjects.

Action

Call Sybase Technical Support when this error occurs. They may be able to help you delete the object that is causing the error. However, because other objects may reference that object, deleting it cleanly may be difficult. If this is the case, recovering from backups may be a better choice.

Version in which this error is raised

All versions

Error 2514

Severity

16

Message text

```
Table Corrupt: Type id %ld (type name = %.*s) does not
match between %.*s and %.*s
```

Explanation

syscolumns contains a row for every column in a table, and systypes contains a row for every type in a table. Error 2514 occurs when dbcc checkcatalog detects a type mismatch between the systypes and the syscolumns system tables.

Error 2514 is most often caused by a select into across databases when at least one column from the original table is a user-defined datatype. It can also be caused by cross-database views and applications that map data across databases when a user-defined datatype involved in the operation has conflicting definitions in the two databases.

Action

The 2514 error often prints out a usertype value that seems out of place: 0 or a large number. If your 2514 error does not match this description, call Technical Support for further assistance.

Use the following procedure to find the incorrect entry in syscolumns and replace it with the correct value. Note that “c” and “t” are actually typed in; they are not variables. This method allows Adaptive Server to compare all tables to search for the mismatch.

1 Identify the rows containing incorrect entries:

```
1> select c.name, c.usertype, c.type
2> from syscolumns c
3> where not exists
4> (select * from systypes t
5> where t.usertype = c.usertype)
6> go
 name                          usertype     type
```

```
------------------      ---------   -------
PartNumber                     0        56
```

More than one row may be returned.

2 Search for the correct value for each incorrect syscolumns usertype by querying the systypes usertype column, using the value reported in type. Do this for each row reported in step 1 above:

```
1> select t.name, t.usertype, t.type
2> from systypes t
3> where t.type in ( 56 )
4> and t.usertype < 100
```

Results of the query should resemble the following:

```
name                 usertype         type
---------------    ------------     ---------
int                          7            56
```

The value returned under usertype is the correct value which should appear in the usertype column of syscolumns, 7 in this example.

3 Change the incorrect row value in the systypes column usertype with the correct value you obtained in step 2:

a Enable updates to system tables (substitute your site's values for “7,” “0,” “56” and “PartNumber” in this query):

```
1> sp_configure "allow updates", 1
2> go
1> begin transaction
2> update syscolumns set usertype = 7
3> where usertype = 0 and type = 56
4> and name = "PartNumber"
5> go
```

b To commit the transaction, type:

```
1> commit transaction
2> go
```

c Repeat the steps above for any other incorrect rows, substituting the required values for usertype, type, and name in the query.

d When you have updated all the incorrect rows, disable updates to system tables:

```
1> sp_configure "allow updates", 0
2> go
1> checkpoint
```

```
2> go
```

4 Run dbcc checkcatalog again to verify that the problem is corrected.

Version in which this error is raised

All versions

Error 2517

Severity

16

Message text

```
Table Corrupt: Procedure id %ld (procedure name =
%S_OBJID) does not match between %.*s and %.*s
```

Explanation

This error occurs when dbcc checkcatalog finds that an entry for a compiled object (for example, a stored procedure or trigger) exists in the table *sysprocedures* and that entry does *not* exist in the sysobjects table.

Action

1 If the error occurred on an object in sysprocedures in master, enable updates to system tables:

```
1> use master
2> go
1> sp_configure "allow updates", 1
2> go
```

2 Determine how many rows contain the object entry:

```
1> use database_name
2> go
```

```
1> select * from sysprocedures
2> where id = procedure_ID
3> go
```

where database_name is the name of the database that contains the table with the object and *procedure_ID* is the ID of the procedure in the 2517 error.

3 If you need to recover the stored procedure, get the text of the stored procedure:

```
1> sp_helptext stored_proc_name
2> go
```

4 Remove the object entry from sysprocedures:

```
1> begin transaction
2> delete sysprocedures
```

```
3> where id = object_ID
4> go
```

where *object_ID* is the ID of the procedure in the 2517 error.

5 Check the results carefully, and commit the transaction only if the update affected the expected number of rows. If it affected more than the expected number of rows, roll back the transaction. To commit the transaction, type:

```
1> commit transaction
2> go
```

6 Disable updates to system tables:

```
1> sp_configure "allow updates", 0
2> go
1> checkpoint
2> go
```

7 Run dbcc checkcatalog again to verify that the problem is corrected. If the object entry still exists in other tables, you may get other error messages (such as Error 2513). If you do get other error messages, follow the directions in this manual for handling those errors.

8 If you wanted the object that has been deleted, re-create it using the information obtained in step 3.

Version in which this error is raised

All versions

Error 2520

Severity

16

Message text

```
Database named %.*s not found; check sysdatabases
```

Explanation

This error occurs when Adaptive Server is unable to find the database you specify when running dbcc checkcatalog, dbcc checkdb, or dbcc dbrepair.

Some reasons for this error are:

- You have mistyped the name of a database or the database you specified does not exist.
- Corrupt indexes exist on sysdatabases.

Action

Determine whether the database you are specifying exists:

```
1> sp_helpdb
```

```
2> go
```

Search the output to determine whether you have misspelled the database name or whether the dbid for the database you specified exists but the name has been removed or changed.

If the database you are specifying exists and you are getting a 2520 error, run dbcc checkdb on the master database to find out whether it is corrupt. If dbcc checkdb returns error messages, corruption has occurred. Whether or not you believe corruption has occurred, call Sybase Technical Support. You may be able to rebuild the indexes for master or you may have to restore master from a backup.

Before calling Technical Support, have the following information available:

- Adaptive Server version and SWR version level
- Output of dbcc checkdb
- Text of all error messages

Version in which this error is raised

All versions

Error 2521

Severity

16

Message text

```
Table Corrupt: Page is linked but not allocated. Run
DBCC TABLEALLOC to correct the problem. (alloc
page#=%ld, extent id=%ld, logical page#=%ld, object id
in extent=%ld, index id in extent=%ld, object
name=%S_OBJID)
```

Note This error may be caused by a hardware problem.

Explanation

This error is serious, especially if it occurs on a table's data pages. It means that a page is currently in use by a table or index but it has not been marked as allocated. The same page could be allocated again, resulting in a loss of whatever data resides on the page.

Pages encountering the 2521 error will not be included in a database dump. This is because database dumps are performed by reading allocation pages and not by traversing page chains. Therefore, this error should be corrected before dumping the database.

Note The instructions below are for fixing 2521 errors once they have occurred. Two easy-to-use strategies exist for detecting this error sooner in the future. Refer to Chapter 2, "Detecting Allocation Errors as Early as Possible" for information about these strategies.

Action

Occasionally dbcc checkalloc reports this error when no real error condition exists. You can either check to see if these errors are real, or continue with this section and take action to correct them, whether or not they reflect a real allocation error.

Because the process used to determine whether the errors are real can be time-consuming, you may want to go directly to "Error Resolution" below.

Verifying That the Error Is Real

Run dbcc checkalloc in single-user mode if you suspect the 2521 error messages are incorrect. If the error is in master, use the section Chapter 2, "How to Start Adaptive Server in Single-User Mode" for instructions about how to invoke Adaptive Server in single-user mode. Refer to "dbcc" in the *Reference Manual* for information about dbcc checkalloc.

Error Resolution

If many of these errors are occurring, it is possible to clear them all at once by using the dbcc checkalloc and dbcc checkalloc with fix option commands. Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about using dbcc checkalloc.

If the text of the error message includes a real object name, not a number, then the error is on an existing object to which the system catalog has correct references. Continue now to "Identify Table: User or System Table".

If a number appears instead of the object name, then that object only partially exists and the error must be corrected using the procedure described in Chapter 2, "How to Fix and Prevent Allocation Errors".

Identify Table: User or System Table

Look at the value for "object id in extent" in the error message. If it is 100 or greater, go to "Action for User Tables". If the "object id in extent" is below 100, it is a system table and requires a different procedure described in the section "Action for System Tables".

Action for User Tables

If the "object id in extent" in the error message is 100 or greater, follow these steps to correct the error:

1 Check the value of the "index id in extent" in the error message to determine whether it is a table (value is 0) or an index (with a value between 0 and 255).

2 Run dbcc tablealloc or dbcc indexalloc, depending on whether the object named in the 2521 error message is a table or an index. Before you run the appropriate command, keep the following in mind:

- dbcc tablealloc can correct this problem on a table or an index, but if the problem is on an index, you can avoid affecting the entire table by using the dbcc indexalloc command. If the table is large or heavily used, it may be more practical to use dbcc indexalloc.
- These commands correct the error only when run in the full or optimized mode. Do not specify the nofix option, or the 2521 error will not be corrected.
- You can use either the "object name" or "object id in extent" values from the error message in the commands above where the argument "object_name" appears.

Use the command appropriate for your situation:

For Tables (index id in extent = 0)	For Indexes (0 < index id in extent < 255)
1> dbcc tablealloc (*object_name*) 2> go	1> dbcc indexalloc (*object_name*, 2> *index_id_in_extent*) 3> go

Refer to "dbcc" in the *Reference Manual* and "Checking Database Consistency" in the *System Administration Guide* for information about the dbcc tablealloc and dbcc indexalloc commands.

Action for System Tables

If the "object id in extent" in the error message is less than 100, follow these steps to correct the error:

1 Put the affected database in single-user mode:

- If the database is master, use the procedure in Chapter 2, "How to Start Adaptive Server in Single-User Mode", and then continue with this procedure.
- If the database is not master, use the sp_dboption stored procedure to put the affected database in single-user mode:

```
1> use master
2> go
1> sp_dboption database_name, single, true
2> go

1> use database_name
2> go

1> checkpoint
2> go
```

2 Check the value of the "index id in extent" in the error message to determine whether it is a table (value is 0) or an index (value is greater than 0).

3 Run dbcc tablealloc or dbcc indexalloc, depending on whether the object named in the 2521 error message is a table or an index. Before you run the appropriate command, keep the following in mind:

- dbcc tablealloc command above can correct either a table or an index, but if the problem is on an index, you can avoid affecting the entire table by using the dbcc indexalloc command. If you need to minimize the amount of time the table is unavailable, it may be more practical to use dbcc indexalloc.
- These commands correct the error only when run in the full or optimized mode, with the fix option specified. This is necessary because the default value is nofix when these commands are run on system tables.
- You can use either the "object name in extent" or "object id" values from the error message in the commands above where the argument "object_name" appears.

Use the command appropriate for your situation:

For Tables (index id in extent = 0)	For Indexes (0 < index id in extent < 255)
1> dbcc tablealloc *(object_name*, 2> full, fix) 3> go	1> dbcc indexalloc (*object_name*, 2> *index_id_in_extent*, full, fix) 3> go

4 Turn off single-user mode in the database:

- If the database is master, use the procedure in Chapter 2, "Returning Adaptive Server to Multiuser Mode".
- If the database is not master, use the following procedure:

```
1> use master
2> go
1> sp_dboption database_name, single, false
2> go

1> use database_name
2> go

1> checkpoint
2> go
```

Refer to "dbcc" in the *Reference Manual* and "Checking Database Consistency" in the *System Administration Guide* for information about the dbcc tablealloc and dbcc indexalloc commands.

Version in which this error is raised: All versions

Error 2524

Severity: 16

Message text:

```
Table Corrupt: Row length is inconsistent between the
computed row length and the recorded row length on page;
check the following page and row: pageno=%ld row#=%d
computed row length=%d row length on page=%ld
```

Note This error may be caused by a hardware problem.

Explanation

Adaptive Server computes the length of each row by using the values stored in the offset and adjust tables of each row. Errors 2524 occurs when dbcc checkdb or dbcc checktable detects that the actual length of a row does not match the computed length of that row.

This error can occur on a data page as well as on an index page.

Action

If you have clean backups of your database, recover from backups. If you do not have clean backups of your database, refer to Chapter 2, "How to Identify and Fix a Corrupted Table" to correct the error.

Warning! Some data on this page might be lost if you recover your table using bcp or select into (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

Before dumping your database, make sure it is working correctly by running the following commands prior to each dump:

1 dbcc checkdb.

2 dbcc checkalloc or dbcc checkalloc with the fix option. (Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about running these commands in multi-user mode and how to prevent spurious allocation errors from dbcc commands.)

Hardware Errors

Error 2524 is probably caused by a problem within Adaptive Server but may also be caused by hardware errors on your database devices. To help rule out this possibility, do the following:

- Check the Adaptive Server error log to determine whether there are other indications of hardware problems, such as kernel messages reporting I/O errors.
- Check the operating system error log or diagnostics utilities for I/O errors. Refer to Chapter 2, "Checking the Operating System Error Log".

If problems persist, refer to the error documentation in this manual or contact Sybase Technical Support.

Additional information

Refer to bcp in the Adaptive Server utility programs manual.

Refer to "select" in the *Reference Manual* for information about the select into command.

Have the following information available before you call Sybase Technical Support:

- Adaptive Server version and SWR version level
- Adaptive Server error log
- Operating system error log
- Output of dbcc checkdb and dbcc checkalloc
- Text of all error messages

Version in which this error is raised

All versions

Error 2525

Severity

16

Message text

```
Table Corrupt: Object id wrong; tables: alloc page %ld
extent id=%ld l page#=%ld objid in ext=%ld (name =
%S_OBJID) objid in page=%ld (name = %S_OBJID) objid in
sysindexes=%ld (name = %S_OBJID)
```

Explanation

This error occurs when dbcc checkalloc detects one of the following problems:

- The object ID on a page does not match the object ID on the allocation structure (extent) associated with that page. This is the case when the "objid in ext" matches the "objid in page" from the error message.
- The object ID on a page does not match the object ID in the system table sysindexes . This is the case when the "objid in ext" does not match the "objid in sysindexes" from the error message.

Action

Compare the "objid in ext" to "objid in page" from the error message. If they do not match, go to "Object IDs Do Not Match". If they are equal, go to "Object IDs Match".

Warning! If you drop an object before you correct the 2525 error condition, the drop operation could cause pages belonging to another object to be deallocated. This can cause 1108 errors.

Object IDs Do Not Match

There are two possible situations in this scenario. Either the object exists and may be corrupted or it no longer exists. To determine which scenario you have, do the following:

```
1> use database_name
2> go

1> select object_name(objid_in_extent)
2> go
```

where *database_name* is the database where the corruption is reported and *objid_in_extent* is from the error message.

If the result of the query is a table that currently exists or if either object (object ID on the page or the object ID of the allocation structure associated with that page) is a system table (object ID is less than 100), call Sybase Technical Support. If you have dial-in facilities, Technical Support may be able to repair this corruption without losing any data. However, restoring from known, clean backups may be necessary.

If the query does not return any rows, this means the object that is marked as owning the extent no longer exists.

If the table is a user table, you may be able to correct this error. Back up the data in the table either by using bcp to copy out the table whose object ID appears in "objid in page" or by selecting it into a new table. Refer to Chapter 2, "How to Rescue Data from a Corrupted Table" for instructions.

Once a copy of the table has been made, you can do one of the following:

1 Create a clustered index on the table. If you already have a clustered index, drop it and then re-create it.

2 Drop the corrupt table, re-create it, and move the data back into the table using a bcp in, an insert with a subselect from the holding table, or select into. Alternatively, the holding table itself can be renamed once the original table has been dropped.

Object IDs Match

There are two possible scenarios:

- The pointer to a distribution page for a particular index may be invalid.
- There may be a mismatch between the index ID on the allocation page and the index ID on the data page.

- Use the following commands to determine which index has the problem. This query compares the "id" column in the output to the "objid in sysindexes" from the error message, and the "distribution" column to the "l page #" from the error message.

```
1> use database_name
2> go

1> select indid, name, object_name(id)
2> from sysindexes
3> where id = objid_in_sysindexes and
4> distribution = l_page_#
5> go
```

  If this query does not return any rows, contact Sybase Technical Support.

- If the query does return a row, drop and re-create the index specified by this query to correct the problem. (Refer to the *Adaptive Server Reference Manual* for more information about drop index and create index.)
- Execute dbcc tablealloc to make sure that the problem is corrected. If the problem still exists, contact Sybase Technical Support.

Additional information

If the object IDs match, dbcc extentdump and dbcc page can also be used to investigate the problem.

Version in which this error is raised

All versions

Error 2526

Severity

16

Message text

```
Incorrect DBCC command: Please see the SQL Server
Reference Manual for DBCC commands and options.
```

Explanation

The following situations can cause this error:

- An attempt was made to execute an invalid dbcc command.
- An incorrect number of parameters has been passed to a dbcc command.
- Some other syntax error was made on a dbcc command.

Action

Make sure you are using the correct name and syntax for dbcc commands. To determine the correct syntax, refer to "dbcc" in the *Reference Manual*.

Version in which this error is raised

All versions

Error 2529

Severity 16

Message text

```
Table Corrupt: Attempted to get page %ld, object %ld;
got page %ld, object %ld.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when dbcc checkdb, dbcc checktable, dbcc checkalloc, or dbcc tablealloc detects an inconsistency between the location of a page on disk and the page number stored in the page header. In the best case, only the value of the page number is wrong. In the worst case, the entire page is corrupt.

This is a serious error that must be corrected immediately.

If the page number in the error is evenly divisible by 256, the error has occurred on an allocation page. All objects referred to by that allocation unit (that page and the 255 pages following it) may be affected (up to 32 objects). Otherwise, the error occurred on a system or user table.

Warning! Attempts to drop an object with a 2529 error may result in 821 and 813 errors and additional database corruption.

Action

1 Note the object ID displayed in the error message.

2 Run dbcc checkalloc and dbcc checkdb on the database affected to determine the full extent of the corruption.

3 The 2529 error often results from hardware errors, or system administration problems such as inadvertent use of a raw partition for two separate purposes. Check on and correct any such problems before attempting to restore the database:

- Refer to Chapter 2, "Correct Use of Raw Partitions" for information about using raw partitions.
- Check the Adaptive Server error log to determine whether there are other indications of hardware problems, such as kernel messages reporting I/O errors.
- Check the operating system error log or diagnostics utilities for I/O errors.

4 If the 2529 error is on a system table (object ID is less than 100), call Sybase Technical Support. If the object ID is 99, another option is to restore from backups.

5 If the 2529 error is on a user table (object ID is greater than or equal to 100), you may be able to select the data into another table and drop the original table. Refer to Chapter 2, "How to Rescue Data from a Corrupted Table" for instructions.

After you drop the original table, you may get an 821 error. In this case, shut down Adaptive Server and restart it. Try a select from the old corrupt table to make sure it is gone, then run dbcc checkalloc and dbcc checkdb to confirm that there is no more corruption.

In some cases, a 2529 error is first reported as a 2503 error, and becomes a 2529 error after Adaptive Server is restarted. Although the procedure described above may clear the 2529 error, it will not remove the main cause of your corruption.

Additional information

Before calling Technical Support, have the following information available:

- Adaptive Server version and SWR version level
- Adaptive Server error log
- Operating system error log
- Output of dbcc checkdb and dbcc checkalloc
- Text of all the error messages

Version in which this error is raised

All versions

Error 2540

Severity

16

Message text

```
Table Corrupt: Page is allocated but not linked; check
the following pages and ids: allocation pg#=%ld extent
id=%ld logical pg#=%ld object id on extent=%ld (object
name = %S_OBJID) indid on extent=%ld
```

Explanation

This error occurs when dbcc checkalloc determines that a page is marked as allocated to an object but that page is not being used. There is no corruption or data loss associated with this error.

Each 2540 error means the loss of one blank data page. A few 2540 errors are no cause for concern. However, if many of these errors occur, the amount of "lost" disk space could be significant.

Note The instructions below are for fixing 2540 errors once they have occurred. Two easy-to-use strategies exist for detecting this error sooner in the future. Refer to Chapter 2, "Detecting Allocation Errors as Early as Possible" for information about these strategies.

Action

Occasionally dbcc checkalloc reports this error when no real error condition exists. You can either check to determine whether the error is real, or continue with this section and take action to correct it, whether or not it reflects a real allocation error.

Because the process used to discover whether or not the error is real can be time-consuming, you may want to go directly to "Error Resolution" now.

Verifying That the Error Is Real

Run dbcc checkalloc in single-user mode if you suspect the 2540 error messages are incorrect. If the error is in master, use the section Chapter 2, "How to Start Adaptive Server in Single-User Mode" for instructions about how to invoke Adaptive Server in single-user mode. Refer to "dbcc" in the *Reference Manual* for information about dbcc checkalloc.

Error Resolution

If many of these errors are occurring, it is possible to clear them all at once by using the dbcc checkalloc and dbcc checkalloc with fix option commands. Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about using dbcc checkalloc.

If the text of the error message includes a real object name, not a number, then the error is on an existing object which the system catalog has correct references to, and you should continue now to "Identify Table: User or System Table".

If a number appears instead of the object name, then that object only partially exists and the error must be corrected using the procedure described in Chapter 2, "How to Fix and Prevent Allocation Errors".

Identify Table: User or System Table

Look at the value for "object id on extent" in the error message. If it is 100 or greater, continue with "Action for User Tables". If the "object id on extent" is below 100, it is a system table and requires a different procedure as described in the section "Action for System Tables".

Action for User Tables

If the "object id on extent" in the error message is 100 or greater, follow these steps to correct the error:

1 Check the value of the "indid on extent" in the error message to determine whether it is a table (value = 0) or an index (value > 0).

2 Run dbcc tablealloc or dbcc indexalloc, depending on whether the object named in the 2540 error message is a table or an index. Before you run the appropriate command, keep the following in mind:

- dbcc tablealloc corrects this problem on a table or an index, but if the problem is on an index, you can avoid affecting the entire table by using dbcc indexalloc. If the table is large or heavily used, it may be more practical to use dbcc indexalloc.
- These commands can correct the error only when run in the full or optimized mode, and with the nofix option not specified, the default for user tables.
- You can use the object name or object ID in the following commands where the argument "object_name" appears.

Use the command appropriate for your situation:

For Tables (index id in extent = 0)	For Indexes (0 < index id in extent < 255)
1> dbcc tablealloc (*object_name*) 2> go	1> dbcc indexalloc (*object_name*, 2> *indid_on_extent*) 3> go

Refer to "dbcc" in the *Reference Manual* and "Checking Database Consistency" in the *System Administration Guide* for information about dbcc tablealloc and dbcc indexalloc.

Action for System Tables

If the "object id on extent" in the error message is less than 100, follow these steps to correct the error:

1 Put the affected database in single-user mode:

- If the database is master, use the procedure in Chapter 2, "How to Start Adaptive Server in Single-User Mode", and then go to step 2.
- If the database is not master, use the sp_dboption stored procedure to put the affected database in single-user mode:

```
1> use master
2> go
1> sp_dboption database_name, single, true
2> go

1> use database_name
2> go

1> checkpoint
2> go
```

2 Check the value of the "indid on extent" in the error message to determine whether it is a table (value is 0) or an index (value is greater than 0).

3 Run dbcc tablealloc or dbcc indexalloc, depending on whether the object named in the 2540 error message is a table or an index. Then execute the appropriate command. Before you run the appropriate command, keep the following in mind:

- dbcc tablealloc corrects either a table or an index, but if the problem is on an index, you can avoid affecting the entire table by using dbcc indexalloc. If you need to minimize the amount of time the table is unavailable, it may be most practical to use dbcc indexalloc.
- These commands correct the error only when run in the full or optimized mode, with the fix option specified, because the default value is nofix on system tables.
- You can use the "object name" or "object id on extent" in the commands above where the argument *object_name* appears.

Use the command appropriate for your situation:

For Tables (index id in extent = 0)	For Indexes (0 < index id in extent < 255)
1> dbcc tablealloc (*object_name*, 2> full, fix) 3> go	1> dbcc indexalloc (*object_name*, 2> *indid_on_extent*, full, fix) 3> go

4 Turn off single-user mode in the database:

- If the database is master, refer to Chapter 2, "Returning Adaptive Server to Multiuser Mode".

- If the database is not master, use the following procedure:

```
1> use master
2> go
1> sp_dboption database_name, single, false
2> go

1> use database_name
2> go

1> checkpoint
2> go
```

Refer to "dbcc" in the *Reference Manual* and "Checking Database Consistency" in the *System Administration Guide* for information about dbcc tablealloc and dbcc indexalloc.

Version in which this error is raised

All versions

Error 2546

Severity

16

Message text

```
Table Corrupt: Extent id %ld on allocation pg# %ld has
objid %ld and used bit on, but reference bit off.
```

Explanation

This error is reported by dbcc checkalloc if an allocation structure or extent is not linked to the other extents for the object referenced by "objid", but does reference a page which is linked in the object's page chain. Each occurrence of this error can represent a loss of up to eight pages on disk, or 16K. The pages cannot be used until the error is corrected. Error 2546 can lead to data corruption and can result in various run-time failures. If no other errors are occurring, you can wait until nonpeak hours to correct the problem.

Note The instructions below are for fixing 2546 errors once they have occurred. Two easy-to-use strategies exist for detecting this error sooner in the future. Refer to Chapter 2, "Detecting Allocation Errors as Early as Possible" for information about these strategies.

Action

Occasionally dbcc checkalloc reports this error when no real error condition exists. You can either check to determine whether the error is real, or continue with this section and take action to correct it, whether or not it reflects a real allocation error.

Because the process used to discover whether or not the error is real can be time-consuming, you may want to go directly to "Error Resolution" now.

Verifying That the Error Is Real

Run dbcc checkalloc in single-user mode if you suspect the 2546 error messages are incorrect. If the error is in master, use Chapter 2, "How to Start Adaptive Server in Single-User Mode" for instructions about how to invoke Adaptive Server in single-user mode. Refer to "dbcc" in the *Reference Manual* for information about dbcc checkalloc.

Error Resolution

If many of these errors are occurring, it is possible to clear them all at once by using the dbcc checkalloc and dbcc checkalloc with fix option commands. Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about using dbcc checkalloc.

Execute the following query to make sure the object exists and is correctly referred to in the system catalog:

```
1> use database_name
2> go

1> select object_name ( objid_from_error_msg)
2> go
```

If an object name is returned, then the error is on an existing, correctly referenced object. If this is the case, go now to the section "Identify Table: User or System Table".

If a number, or something other than an object name is returned, use Chapter 2, "How to Fix and Prevent Allocation Errors".

Identify Table: User or System Table

Look at the value for "objid" in the error message. If it is 100 or greater, continue with the next section "Action for User Tables". If the "objid" is below 100, it is a system table and requires a different procedure as described in the section "Action for System Tables".

Action for User Tables

If the "objid" in the error message is 100 or greater, follow these steps to correct the error:

1 Run the dbcc page command described in Chapter 2, "How to Find an Object Name from a Page Number" to obtain the value for the "indid". Substitute the value for "extent id" in the 2546 error message in place of the "page_number" as described in the procedure.

2 Run dbcc tablealloc or dbcc indexalloc, depending on whether the value for "indid" indicates it is a table or an index. Before you run the appropriate command, keep the following in mind:

- dbcc tablealloc will correct this problem on a table or an index, but if the problem is on an index, you can avoid affecting the entire table by using dbcc indexalloc. If the table is large or heavily used, it may be more practical to use dbcc indexalloc.
- These commands will correct the error only when run in the full or optimized mode as long as the nofix option is not specified, the default for user tables.
- You can use the object name, if you know it, or the "objid" value from the error message in the commands above where the argument *objid* appears.

Use the command appropriate for your situation:

For Tables (index id in extent = 0)	For Indexes (0 < index id in extent < 255)
1> dbcc tablealloc (*objid*) 2> go	1> dbcc indexalloc (*objid*, *indid*) 2> go

Refer to "dbcc" in the *Reference Manual* and "Checking Database Consistency" in the *System Administration Guide* for information about dbcc tablealloc and dbcc indexalloc.

Action for System Tables

If the "objid" in the error message is less than 100, follow these steps to correct the error:

1 Run the dbcc page command described in Chapter 2, "How to Find an Object Name from a Page Number" to obtain the value for the "indid." Substitute the value for "extent id" in the 2546 error message in place of the "page_number" as described in the procedure.

2 Put the affected database in single-user mode:

- If the database is master, use the procedure in Chapter 2, "How to Start Adaptive Server in Single-User Mode", and then go to step 3.
- If the database is not master, use the sp_dboption stored procedure to put the affected database in single-user mode:

```
1> use master
2> go
1> sp_dboption database_name, single, true
2> go

1> use database_name
2> go

1> checkpoint
2> go
```

3 Run dbcc tablealloc or dbcc indexalloc, depending on whether the value of the "indid" from step 1 above indicates that it is a table or an index, then execute the appropriate command. Before you run the appropriate command, keep these facts in mind:

- dbcc tablealloc will correct either a table or an index, but if the problem is on an index, you can avoid affecting the entire table by using dbcc indexalloc. If you need to minimize the amount of time the table is unavailable, it may be more practical to use dbcc indexalloc.
- These commands will correct the error only when run in the full or optimized mode with the fix option specified, because the default for system tables is nofix.
- You can use the object name if you know it, or "object id" value from the error message in the commands above where the argument *objid* appears.

Use the command appropriate for your situation:

For Tables (indid = 0)	For Indexes (0 < indid < 255)
1> dbcc tablealloc (*objid*, 2> full, fix) 3> go	1> dbcc indexalloc (*objid*, 2> *indid*, full, fix) 3> go

4 Turn off single-user mode in the database:

- If the database is master, use Chapter 2, "Returning Adaptive Server to Multiuser Mode".
- If the database is not master, use the following procedure:

```
1> use master
```

```
2> go
1> sp_dboption database_name, single, false
    2> go

1> use database_name
2> go

1> checkpoint
2> go
```

Refer to “dbcc” in the *Reference Manual* and “Checking Database Consistency” in the *System Administration Guide* for information about dbcc tablealloc and dbcc indexalloc.

Version in which this error is raised

All versions

Error 2550

Severity

16

Message text

```
Missing segment in sysusages segmap.
```

Explanation

This error is reported by dbcc checkcatalog if there is a row in sysusages (in the master database) that has a value of 0 in the segmap column. A value of 0 means that the corresponding space will not be used for further space allocation (data, log, or user-defined objects). The error is not fatal but it does indicate that any free space on that section of the database will not be used.

Causes of this error might be:

- All of the segments have been removed from a logical device (via **sp_dropsegment**).
- A direct update has been made to the segmap column of sysusages where the value has been changed to 0.

Action

If you do not want any new space allocated on the affected segment, no action is necessary.

Otherwise, run sp_addsegment or sp_extendsegment on the Sybase logical device that has the problem. If the segment does not already exist in syssegments, use sp_addsegment. If the segment exists, use sp_extendsegment.

To add a new segment:

```
1> use database_name
```

```
2> go

1> sp_addsegment segname, database_name, devname
2> go
```

where:

- *segname* is the name of the new segment.
- *database_name* is the name of the database where the segment is to be defined.
- *devname* is the name of the database device where *segname* will be located. (A database device may have more than one segment associated with it. The space on the new segment will only be used if you create tables or indexes on that new segment.)

To extend an existing segment on the device that has a segmap of 0:

```
1> use database_name
2> go

1> sp_extendsegment segname, database_name, devname
2> go
```

where:

- *segname* is the name of the existing segment. You can extend the segment called default. If you do, the command would be:

  ```
  1> use database_name
  2> go

  1> sp_extendsegment "default", database_name,
  2> devname
  3> go
  ```

- *database_name* is the name of the database where the segment is to be extended.
- *devname* is the name of the database device that has the segmap 0.

To check that the problem has been resolved, type:

```
1> select * from sysusages where
2> dbid = db_id("database_name") and
3> segmap = 0
4> go
```

where database_name is the name of the database where the segment you added or extended resides. If no rows are returned, the problem has been resolved.

Additional information

For further information about sp_addsegment and sp_extendsegment, refer to the *Reference Manual*.

For further information about segments, refer to "Creating and Using Segments" in the *System Administration Guide*.

Version in which this error is raised

All versions

Error 2558

Severity

16

Message text

```
Extent not within segment: Object %ld, indid %d includes
extents on allocation page %ld which is not in segment
%d.
```

Explanation

A segment is a label used to point to one or more database devices. Segments map the future allocations of different types of data to different devices.

Error 2558 occurs when dbcc checkalloc, dbcc tablealloc, or dbcc indexalloc discovers that a portion of a database resides on a segment which was not defined to contain that type of data. For example, if dbcc checkalloc discovers user data on a segment designated to hold transaction log data, Error 2558 will occur.

Note Error 2558 is only raised when the 2513 trace flag is turned on.

When a database gets this error, it still functions. However, 1105 errors may occur as a side effect of the 2558 error.

In addition, if you have log on your data segment, and you lose the data device for some reason, you cannot get a good transaction dump with no_truncate. This means you cannot get up-to-the-minute recovery.

Sybase recommends that you do *not* turn on trace flag 2513 unless you have been getting 1105 errors or have recently run sp_logdevice, sp_placeobject, or alter database commands. When trace flag 2513 is on, dbcc checkalloc, dbcc tablealloc, and dbcc indexalloc commands take significantly longer to run and a lot of additional output may be created.

To turn on trace flag 2513, follow these commands:

```
1> dbcc traceon (2513)
```

```
2> go
```

To turn off trace flag 2513, follow these commands:

```
1> dbcc traceoff (2513)
2> go
```

About Creating Segments

By default, when Adaptive Server creates a database, it creates the following three segments:

- The *system* segment allows system catalog tables to be placed on devices that contain this segment for a particular database.
- The *default* segment allows user tables or indexes to be placed on devices that contain this segment without specifying the on clause (for example, create table test (i int) on segment_1) for a particular database.
- The *log* segment allows the transaction log to be placed on devices that contain this segment for a particular database.

Note A device can contain more than one type of segment, although a segment can contain only one type of data (user data, transaction log data, and so on). Thus, different data can be on the same device, but not on the same segment.

For example, this command creates the default, system and log segments for the database example:

```
1> create database example
2> on device_1 = 10
3> log on device_2 = 2
4> go
```

The system and default segments are mapped to device_1 and the log segment is mapped to device_2. For the example database, all future space allocations for system or user objects will occur on device_1 and all transaction log records will be placed on device_2. Other databases may use other portions of device_1 or device_2 and these devices may have different segment mappings for that database.

A 2558 error would result in the example database if system or user objects were located on device_2 or transaction log records on device_1.

What Causes a 2558 Error?

Following is a summary of how and when Error 2558 is raised. For details, refer to the sections listed in the "See..." column.

Causes of Error 2558	See...
Cause: Loading into a database that has different data and log mapping than the dumped database.	"Different Data and Log Mapping"
Cause: When you use sp_placeobject, the old allocation area is still in use[a].	"sp_placeobject"
Cause: When you use sp_dropsegment, existing objects are still mapped to the dropped segment.	"sp_dropsegment"
a. Error 2558 is only raised in this case when trace flag 2513 is turned on.	

Different Data and Log Mapping

If this error occurred after you re-created and loaded a database from a dump, there were probably incorrect or different entries in the master..sysusages table. A database loaded from a dump must be created in exactly the same way as the database that was dumped. If it is not, different entries in master..sysusages may cause 2558 errors.

This can cause a "data on log" situation, where data such as user or system objects is loaded into a log segment. This means that less total log space is available to the database. The data is still accessible, but processing may be slowed, if not totally stopped, by the shortage of log space.

There can also be a "log on data" situation, where portions of the transaction log are loaded into a data segment. This is not as serious as the "data on log" scenario. This is because the portion of the transaction log on the data segment is eventually truncated using the dump transaction command, freeing up the space. All future transaction log allocations will be correctly placed on the log segment.

Suppose the example database were altered as follows:

```
1> alter database example on device_1 = 2
2> go
```

The logical order of the database pages is:

- The first 10MB of pages are data (default and system segments).
- The next 2MB are log.
- The last 2MB are data.

A 2558 error could result if the database were dumped, then dropped and re-created with different segment mappings. For example:

```
1> create database example on device_1 = 12
2> log on device_2 = 2
3> go
```

Although both ways of creating the database (create and alter, or just create) allocate the same amount of data and log space on the same devices, the mapping of data and log space is not the same. The first 12MB are allocated differently: in the first (create and alter) database, 10MB of data are followed by 2MB of log; in the second (create), the first 12MB are data. Dumping the first database and loading it to the second would cause 2MB of log to be mapped onto a data segment, yielding 2558 errors.

Warning! If you load databases in this manner, you can get mapping that you do not expect and that can cause problems. Do not load databases in this manner! If you suspect that such a problem might have occurred, use sp_helpdb to check the mapping of your database.

All space allocations that occur on Database 2 after the load completes are correctly mapped.

sp_placeobject

If you use sp_placeobject to assign an object to a new segment, existing pages will still reside on the old segment.

If you have trace flag 2513 turned on, dbcc checkalloc, dbcc tablealloc, and dbcc indexalloc will display Error 2558 after you use sp_placeobject.

sp_dropsegment

If you have trace flag 2513 turned on, dbcc checkalloc, dbcc tablealloc, and dbcc indexalloc will display Error 2558 if you use sp_dropsegment to remove a segment for which existing objects are still mapped.

Prevention

Consider using the on segment_name option with create table rather than using sp_placeobject to assign an object to a segment.

Do not drop segments when objects are still mapped to them.

Keep an up-to-date copy of the master..sysusages system table as well as scripts to re-create all databases. Use the scripts to ensure that the entries in the master..sysusages table for the database being loaded match the corresponding entries for the database that was dumped. More specifically, the segmap, lstart, and size columns of sysusages must be identical in content and order. Verify this before beginning the load database command.

The entries in sysusages will be correct on the database being loaded into if you execute the following commands with the same parameters and in the same order as they were on the dumped database:

- create database
- alter database
- sp_addsegment
- sp_dropsegment
- sp_extendsegment
- sp_logdevice

Note You cannot depend on the output of the system procedure **sp_helpdb** to compare the contents of the sysusages table, because it does not necessarily display rows from sysusages in the actual logical page order. Instead, do direct selects against the sysusages table.

Action

You only need to resolve 2558 errors if the object taking up space on the wrong segment is a problem for you.

If you received the error because of an sp_dropsegment, you may be able to resolve it using sp_addsegment or sp_extendsegment.

If there are a substantial number of 2558 errors due to an incorrect load database, it might be a good idea to drop and correctly re-create the database, and then reload it.

Use this section if you have 2558 errors and cannot re-create the database from scripts or hard copy.

2558 errors occur on these types of pages:

- Data pages or clustered index pages
- Nonclustered index pages
- text or image data pages
- System tables

- Transaction log pages

The action needed to correct this error depends on the type of page on which the error occurred.

Use the following table to match up the object ID value ("Object" in the error message) and index ID ("indid" from the error message) with the corresponding data page type, and then go to that section for the appropriate action. All action sections follow the table:

Table 1-7: Data types and object/index ID values

Type of Data	Corresponding Object ID and Index ID Values
"Data Pages or Clustered Index Pages"	object ID > 99, index ID = 0 or 1
"Nonclustered Index Pages"	object ID > 99, 1 < index ID < 255
"text or image Data Pages"	object ID > 99, index ID = 255
"System Table Pages"	object ID < 100 (and not equal to 8)
"Transaction Log Pages"	object ID = 8

An object ID value of 100 will not occur.

Data Pages or Clustered Index Pages

(Object ID > 99 and Index ID = 0 or 1)

The error occurred on the data page or on the clustered index of a user table. Resolve the problem with either one of the following procedures:

- Create a clustered index on the table, if you have the available space needed. If a clustered index already exists, drop and re-create it. If not, creating one will clear the error. You can then drop the clustered index.
- Copy the data out of the table, drop and re-create the table, and copy the data back in using bcp or select into.

Nonclustered Index Pages

(Object ID > 99 and 1 < Index ID < 255)

The error occurred on the nonclustered index of a user table. The error can be cleared by dropping and re-creating the index. Creating a clustered index causes all nonclustered indexes to be rebuilt as well. Therefore, if 2558 errors are occurring on several indexes on the same table, you can create a clustered index on that table to clear all these errors.

text or *image* Data Pages

(Object ID > 99 and Index ID = 255)

The error occurred on data which is text or image datatype. Use either procedure below to correct the situation:

1 Select all the information from the old table into a new table.

2 Drop the old table. Any subsequent dbcc checkalloc, dbcc tablealloc, or dbcc indexalloc command will continue to show this error until the old table is dropped.

3 Use sp_rename to rename the new table with the old table name.

Or:

Copy the table out, using bcp or select/into. Then drop and re-create the table, and copy the contents of the table back in, using bcp or select/into.

System Table Pages

(Object ID < 100 and Not 8)

Call Sybase Technical Support.

Transaction Log Pages

(Object ID = 8)

If the error occurs on pages containing transaction logs, do the following:

1 Have all users finish transactions and wait until this procedure is finished to initiate any further transactions.

2 Perform the checkpoint command:

```
1> use database_name
2> go

1> checkpoint
2> go
```

3 Dump the transaction log according to your normal procedures.

4 Notify users that they may resume normal operations.

Additional information

Refer to “Creating and Using Segments” in the *System Administration Guide* for information about segments and devices.

Refer to the *Reference Manual* for information about checkpoint and dump transaction.

Version in which this error is raised

All versions

Error 2559

Severity 16

Message text
```
Data page number %ld is empty but is not the first page.
Status = 0x%x.
```

Explanation

When an object is first created, a single, empty page is allocated to it. Usually there are no other empty pages. Error 2559 occurs when dbcc checktable or dbcc checkdb encounters an empty page while traversing the page chain of an object which is not the first page of the chain.

If you are not encountering any other errors either at run time or from dbcc, it is possible that the empty page will not cause any further problems. Therefore, it is acceptable to wait to correct this problem until nonpeak hours. If other errors are occurring, refer to the information in this manual if applicable or contact Sybase Technical Support.

Action

Correct this situation by forcing the page chain involved to be rebuilt. This can be accomplished in different ways, depending on whether or not the page chain involved is associated with an index or table data. To determine this, refer to Chapter 2, "How to Find an Object Name from a Page Number" to determine the index ID and object ID associated with the page number from the error message.

If the error is on a user table (a system table's object ID is less than 100), go to "User Tables or Clustered Indexes (Index ID is 0 or 1)."

If the error is on a system table and the index ID is *not* 0, the error is on a system table index. Refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index.

If the error is on a system table and the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption or it may be necessary to restore from clean backups.

User Tables or Clustered Indexes (Index ID is 0 or 1)

There are two ways to correct this error, depending on whether you have enough extra disk space to create a clustered index on the table in question.

Solution 1: Drop Index

If a clustered index exists, drop and re-create it. If not, create a dummy clustered index and then drop it. This will remove this problem since creation of a clustered index causes the table to be copied and does not copy the empty pages. Note that clustered index creation does require considerable disk space (refer to "create index" in the *Reference Manual)*.

Solution 2: Use the Bulk Copy Utility

Use bcp to copy the table out of and then back to the Server:

1 bcp the table data out to an operating system file.

2 Empty the table with the truncate table command.

3 (Optional) Drop all indexes on the table.

4 bcp the table data into Adaptive Server.

5 (Optional) Re-create all indexes on the table.

For Nonclustered Indexes (1 < Index ID < 255)

There are two ways to correct this error on a nonclustered index:

Solution 1: Drop Index

1 Execute the following query in the database in question to determine the index name:

```
1> select name from sysindexes
2> where id = object_ID and indid = index_ID
3> go
```

where *object_ID* and *index_ID* are from the output of dbcc page run earlier.

2 Drop and re-create the nonclustered index specified in the query.

Solution 2: Create a Clustered Index

Create a clustered index since this forces all nonclustered indexes to be rebuilt. This method may not be suitable for large tables because of the overhead required for creating a clustered index.

Version in which this error is raised

All versions

Error 2571

Severity

14

Message text

```
Permission denied. Only a user with System Administrator
(SA) authorization can use the DBCC command '%.*s'.
```

Explanation

Only the “sa” account or users with sa_role can issue the following dbcc commands:

- dbcc dbrepair
- dbcc traceon
- dbcc traceoff
- dbcc checkalloc
- dbcc memusage
- dbcc page

Many stored procedures that use dbcc commands require sa_role. Some of these stored procedures are:

- sp_addsegment
- sp_addthreshold
- sp_dbremap
- sp_dropsegment
- sp_dropthreshold
- sp_extendsegment
- sp_logdevice
- sp_modifythreshold
- sp_placeobject

Action

Run the dbcc commands and stored procedures listed above from the "sa" account or make sure you have sa_role granted to you. To grant sa_role to a user, you must have sa_role granted to you. To grant sa_role to a user, type:

```
1> sp_role "grant", sa_role, login_name
2> go
```

where *login_name* is the name of the login account to which sa_role is being granted. When you grant a role to a user, it takes effect the next time the user logs into Adaptive Server. However, the user can immediately enable the role by using the set role command. For example, this command:

```
1> set role "sa_role" on
2> go
```

enables the sa_role role for the user.

Additional information

For information about using the "sa" account, associated roles, and how locking the "sa" account can affect scripts, refer to the *Security Administration Guide*.

Refer to “Roles” and grant in the *Reference Manual* for information about granting roles.

Version in which this error is raised

All versions

Error 2572

Severity

10

Message text

```
Database '%S_DBID' is not in single user mode - may find
spurious allocation problems due to transactions in
progress.
```

Explanation

This warning occurs when you execute dbcc checkalloc on a database that is not in single-user mode. dbcc checkalloc displays this message and then continues to execute.

Executing dbcc checkalloc on a database that is not in single-user mode may cause error messages to be displayed which do not reflect the actual state of your database. These error messages are referred to as spurious. Error messages 2521, 2540, 2546, 7939, 7940, and 7949 may be spurious under these conditions.

Action

No action is required.

If other errors are reported by dbcc checkalloc, refer to this manual for information on how to deal with them.

If your operational environment allows it, put the database in question into single-user mode prior to running dbcc checkalloc.

- For the master database, refer to Chapter 2, “How to Start Adaptive Server in Single-User Mode”.
- For other databases, use sp_dboption as follows:

```
1> use master
2> go
1> sp_dboption database_name, single, true
2> go

1> use database_name
2> go

1> checkpoint
2> go
```

Additional information

Refer to Chapter 2, "How to Fix and Prevent Allocation Errors".

Refer to "dbcc" in the *Reference Manual* for information about dbcc checkalloc.

Version in which this error is raised

All versions

Error 2573

Severity

16

Message text

```
Database '%.*s' is not marked suspect. You cannot drop
it with DBCC.
```

Explanation

This error occurs when you try to use dbcc dbrepair to drop a database that is not marked suspect. Only databases that have been marked suspect can be dropped using this command.

Databases are marked suspect in the following ways:

- Adaptive Server marks a database suspect when some critical errors occur.
- The Sybase System Administrator can mark a database suspect in order to drop it with the dbcc dbrepair command.

Action

Try to drop the database with the drop database command. If this fails, use the procedure described in Chapter 2, "How to Drop a Database When drop database Fails".

Version in which this error is raised

All versions

Error 2574

Severity

16

Message text

```
Index page number %ld is empty.  Status = 0x%x.
```

Explanation

Tables can contain text columns, which are variable-length columns that can hold printable characters, and image columns, which are variable-length columns that can hold binary data. Adaptive Server stores text and image data in a linked list of data pages separate from the rest of the table. Each 2K text or image page stores a maximum of 1800 bytes of data.

Error 2574 occurs when a dbcc check finds problems with a page that is part of a table's text or image page chain. Error 7105 is a related error; it is raised when Adaptive Server attempts data retrieval or another operation on a table with a bad text or image page chain.

Action

Error 2574 is typically due to a small number of bad rows in a table with text or image data, and can be corrected by dropping and re-creating those rows.

Use the following steps to address the problem:

1 Identify the table. Refer to Chapter 2, "How to Find an Object Name from a Page Number" for information.

2 Determine the first page of the text/image chain using dbcc pglinkage, with the starting page being the page number from the error message. For example, assuming dbid=7 and starting page number 531:

```
1> dbcc pglinkage(7, 531, 0, 1, 0, 0)
2> go

Object ID for pages in this chain = 208003772.
Beginning of chain reached.
Page : 531
Page : 532
Page : 530
3 pages scanned. Object ID = 208003772.
Last page in scan = 530.
```

Page 530 is the first page of the text/image chain.

3 Identify the row(s) with the bad data using the following command:

```
1> select column1 from table_name
2> where convert (int, textptr(text_colname))
3> = first_page_of_chain
4> go
```

where *column1* is any column or columns in your table which will uniquely identify the rows.

4 Drop the row(s) in question.

5 Insert the rows back into the table.

6 Run dbcc checktable on the table to verify that the problem has been corrected.

Version in which this error is raised

All versions

Error 2575

Severity

16

Message text

```
The last page %ld in sysindexes for table '%.*s' has
next page # %ld in its page header. The next page #
should be NULL. Please check sysindexes.
```

Explanation

The sysindexes table lists each table and index, and the segment where each table, clustered index, nonclustered index, and chain of text pages is stored. It also lists other information such as the first page and root for the object. For heap tables (those with no clustered index), the root entry points to the last page in the table's data page chain.

Adaptive Server expects to find a value of zero for the next page number on the last page for a table (meaning that there are no more pages in the page chain). Error 2575 is raised when dbcc checktable or dbcc checkdb is run, and the next page number in the table's last page is not zero. This error is caused by an incorrect root page value for the table in sysindexes, or corruption in the table's last page.

Action

If this is the only error message you see when you run dbcc checktable or dbcc checkdb, it is likely that the table's root entry in sysindexes is incorrect. Use one of these options to recover from the problem:

- Select the table data into a new table, drop the old table, and rename the new table to the old table name. Run dbcc checktable to check that the problem has been resolved.
- Bulk copy the affected table out, drop and re-create the table, and bulk copy back in. This is the most efficient solution for a very large table. Run dbcc checktable to confirm that the problem is resolved.

  For more information about how to copy a table in a new table or file, refer to Chapter 2, "How to Rescue Data from a Corrupted Table".
- Call Sybase Technical Support. They *may* be able to determine the correct last page number and patch the table's root entry in sysindexes.

If you see other errors in addition to Error 2575, the problem may be more extensive. Follow the instructions in this manual for handling those errors.

Additional information

Before calling Technical Support, have the following information available:

- Adaptive Server version and SWR version level
- Adaptive Server error log
- Output of dbcc checktable or dbcc checkdb

- Text of all error messages

Version in which this error is raised

All versions

Error 2578

Severity

16

Message text

```
The first page %ld in sysindexes for table '%.*s' has
previous page # %ld in its page header. The previous
page # should be NULL. Please check sysindexes.
```

Explanation

The sysindexes table lists each table and index, and the segment where each table, clustered index, nonclustered index, and chain of text pages is stored. It also lists other information such as the first page and root for the object.

Adaptive Server expects to find a value of zero for the previous page number on the first page for a table (meaning that there are no prior pages in the page chain). Error 2578 is raised when dbcc checktable or dbcc checkdb is run, and the previous page number in the table's first page is not zero. This error is caused by an incorrect first page value for the table in sysindexes, or corruption in the table's first page.

Action

If the object with the error is a system table, call Sybase Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore from clean backups.

If the table is a user table, and this is the only error message you see when you run dbcc checktable or dbcc checkdb, it is likely that the table's first page entry in sysindexes is incorrect. Use one of these options to recover from the problem:

- Select the table data into a new table, drop the old table, and rename the new table to the old table name. Run dbcc checktable to check that the problem has been resolved.
- Bulk copy the affected table out, drop and re-create the table, and bulk copy back in. This is the most efficient solution for a very large table. Run dbcc checktable to confirm that the problem is resolved.

  For more information about how to copy a table in a new table or file, refer to Chapter 2, "How to Rescue Data from a Corrupted Table".

If you see other errors in addition to Error 2578, the problem may be more extensive. Follow the instructions in this manual for handling those errors.

Additional information

Before calling Technical Support, have the following information available:

- Adaptive Server version and SWR version level
- Server error log
- Output of dbcc checktable or dbcc checkdb
- Text of all error messages

Version in which this error is raised

All versions

Error 2582

Severity

16

Message text

```
Data page (%S_PAGE) has been marked as an overflow page;
however, the previous page (%S_PAGE) does not indicate
that there is an overflow page linked to it.
```

Explanation

With a clustered index, Adaptive Server maintains data so that the physical order of rows is the same as their logical, indexed order. The bottom or leaf level of a clustered index contains the table's actual data pages. When an index key has a large number of duplicates, all of the data rows may not fit on the same data page. Due to the structure of a clustered index, it is not possible to simply insert these 'overflow' rows to the beginning of the next data page. Instead, Adaptive Server allocates and inserts a new data page between the original data page and the next data page, marking this an overflow page to store the additional values (and marking the original page as having an overflow page). The original page, and the overflow page linked to it, now both contain at least one row with that key value. The overflow page contains only rows with an overflow index value.

On tables without a clustered index, the second page through the last page of the table are marked as overflow pages; the first page through the next-to-last page are marked as having an overflow page.

When dbcc checktable examines a page that is marked as an overflow page, it checks the previous (original) page for consistency. Error 2582 is raised when the original page is not marked as having an overflow page.

Action

Take the following steps to correct the problem:

1 Use the procedure described in Chapter 2, "How to Find an Object Name from a Page Number" to translate the second page number displayed in the error message into an object ID.

2 If there is a clustered index on the table, run dbcc tablealloc to correct the problem:

```
1> dbcc tablealloc (table_name, full, fix)
2> go
```

Keep in mind that tablealloc can correct the error only when run in the full or optimized mode.

3 If there is no clustered index on the table, use *one* of these options to recover from the problem:

- Select the table data into a new table, drop the old table, and rename the new table to the old table name. Run dbcc checktable to check that the problem has been resolved.
- Bulk copy the affected table out, drop and re-create the table, and bulk copy back in. This is the most efficient solution for a very large table. Run dbcc checktable to confirm that the problem is resolved.

Additional information

For more information about how to copy a table in a new table or file, refer to Chapter 2, "How to Rescue Data from a Corrupted Table".

Version in which this error is raised

All versions

Error 2583

Severity

16

Message text

```
Number of data pages counted by DBCC differs from count
in the data OAM.
```

Explanation

The Object Allocation Map (OAM) tracks information about the allocation and deallocation of storage space for objects (tables and indexes). An OAM entry exists for each table (except syslogs and sysgams) and for the table's indexes.

All pages for an object are logically chained together, each page containing a header that includes the page numbers of the previous and next pages in the chain.

Error 2583 is raised when Adaptive Server detects that the count of data pages in the OAM is different from the actual number of pages in the chain. This problem is detected by dbcc checkdb and dbcc checktable.

Usually this error will not cause any operations to fail at run time. Therefore, it is acceptable to wait to correct the problem until nonpeak hours.

Error 2583 is similar to Error 7949; both are raised due to OAM data inconsistencies. Error 2583 occurs when Adaptive Server finds incorrect used page counts for the object, whereas Error 7949 occurs when the server finds incorrect unused page counts.

Note The instructions below are for fixing 2583 errors once they have occurred. Two easy-to-use strategies exist for detecting this error sooner in the future. Refer to Chapter 2, "Detecting Allocation Errors as Early as Possible" for information about these strategies.

Action

If many of these errors are raised, as might be the case if you run dbcc checkdb and several tables are affected, it is possible to clear them all at once by using the dbcc checkalloc command with the fix option. Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about using dbcc checkalloc.

If only a small number of these errors are raised and you know the table name (that is, if you ran dbcc checktable), start by determining whether the affected table is a user table or a system table:

```
1> select id from sysobjects
2> where name = object_name
3> go
```

User tables have object IDs of 100 or more; system tables have IDs under 100. Corrective action is different in the two cases because the database must be in single-user mode before correcting system tables.

Action for User Tables

If the affected table is a user table, run dbcc tablealloc to correct the error:

```
1> dbcc tablealloc (table_name, full, fix)
2> go
```

Keep in mind that dbcc tablealloc can correct the error only when run in the full or optimized mode.

Refer to "dbcc" in the *Reference Manual* and "Checking Database Consistency" in the *System Administration Guide* for information about dbcc tablealloc.

Action for System Tables

If the affected table is a system table, follow these steps to correct the error:

1 Put the affected database in single-user mode:

- If the database is master, refer to Chapter 2, "How to Start Adaptive Server in Single-User Mode", and then go to step 2.
- If the database is not master, use the sp_dboption stored procedure to put the affected database in single-user mode:

```
1> use master
2> go
1> sp_dboption database_name, single, true
2> go

1> use database_name
2> go

1> checkpoint
2> go
```

2 Run dbcc tablealloc to correct the error:

```
1> dbcc tablealloc (table_name, full, fix)
2> go
```

Keep in mind that for a system table, *dbcc tablealloc* can only correct the error when run in the *full* or *optimized* mode, and with the *fix* option specified.

3 Turn off single-user mode in the database:

- If the database is master, refer to Chapter 2, "Returning Adaptive Server to Multiuser Mode".
- If the database is not master, use the following procedure:

```
1> use master
2> go
1> sp_dboption database_name, single, false
2> go

1> use database_name
2> go

1> checkpoint
2> go
```

Refer to "dbcc" in the *Reference Manual* and "Checking Database Consistency" in the *System Administration Guide* for information about dbcc tablealloc.

Version in which this error is raised

All versions

Error 2591

Severity

16

Message text

```
Could not find index row which has id %d for table
'%.*s'.
```

Explanation

Error 2591 is raised when you run a dbcc command if Adaptive Server is unable to find an index in sysindexes. The 'id %d' parameter in the message text shows the index id of the index that the dbcc command expected to find.

Error 2591 occurs with the following dbcc commands:

- dbcc reindex
- dbcc indexalloc

Error 2591 is raised with the following states:

State	Meaning
1	dbcc reindex is unable to locate a primary index row for a clustered index on a system table in sysindexes.
2	dbcc reindex is unable to locate a primary index row for a non-clustered index on a system table in sysindexes.
3	dbcc indexalloc is unable to locate the specified index in sysindexes.

Action

If the error was raised when using dbcc indexalloc, check the index name you supplied in the command, and retry using the correct index name. To verify the available indexes for the table:

```
1> sp_help <table_name>
2> go
```

If the error was raised when using dbcc reindex, you need to use an alternative method for correcting the index. Note the object name and index id appearing in the message text, and identify the object id as follows:

```
1> select object_id('table_name')
2> go
```

Then use the procedure described in "How to Fix a Corrupted Index on System Tables" in Chapter 2 to correct the index.

Version in which this error is raised

All versions

Error 2596

Severity 16

Message text `%S_PAGE has an incorrect pgfreeoff value of %d. The offset should be %d.`

Explanation

The free byte offset on a page indicates the offset of the first unused byte on that page. Error 2596 occurs when dbcc checktable detects that the free byte offset on a page is incorrect.

This is a serious error as it indicates possible corruption on the affected page. This error can occur in conjunction with other errors, particularly Error 2506.

Action

If Error 2596 occurs on a system table being processed by dbcc checktable, call Sybase Technical Support.

Case 1

If the number of the free byte offset shown in the error message is *lower* than the expected number, some data has been lost. To determine how much data has been lost, find the difference (in bytes) between the actual free byte offset and the expected free byte offset (both numbers are shown in the error message output). The amount of data lost includes 2 bytes overhead per row where there are only fixed-length columns or 4+*n* bytes overhead per row with variable-length columns where *n* is the number of variable-length columns. If you need help, call Technical Support.

Decide whether you should restore the database that contains the affected table from backups. In making that decision, consider the number of rows lost and the number of other corruption errors that have occurred on this table and database.

Case 2

If the free byte offset on the page header is *higher* than the expected value, no data has been lost. However, some extraneous rows may exist.

If you have clean backups of your database, recover from backups. If you do not have clean backups of your database, follow these steps to correct the error:

1 Use the procedure in Chapter 2, "How to Find an Object Name from a Page Number" to identify which table and index correspond to the page number from the error message text.

2 If the object with the error is *not* a system table (a system table's object ID is less than 100), continue with step 3.

If the object with the error is a system table and the index ID is *not* 0, refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index.

If the object with the error is a system table and the index ID is 0, contact Technical Support. They may be able to help you repair the corruption or it may be necessary to restore from clean backups.

3 For user tables, if the index ID is 0 or 255, continue with step 4.

If the index ID is *not* 0 or 255, translate it into an index name:

```
1> use database_name
2> go
```

```
1> select name from sysindexes
2> where id = object_ID and indid = index_ID
3> go
```

To ensure that the information needed to re-create the index is available, run sp_helpindex on the index prior to dropping it.

Drop the index.

Re-create the index. This clears the corruption in most cases.

Run dbcc checktable on the table to verify that the corruption is gone.

4 If the index ID is 255, delete or update the data row to remove the reference to the bad text page.

If the index ID is 0, do one of the following:

- Restore the database from clean backups.
- Refer to Chapter 2, "How to Rescue Data from a Corrupted Table".

Warning! Some data on this page might be lost if you recover your table using bcp or select into (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

Before dumping your database, make sure it is working correctly by running the following commands prior to each dump:

1 dbcc checkdb.

2 dbcc checkalloc or dbcc checkalloc with the fix option. (Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about running these commands in multi-user mode and how to prevent spurious allocation errors from dbcc commands.)

Version in which this error is raised

All versions

Insert Errors

This section contains error messages for Adaptive Server insert commands.

Error 2601

Severity

14

Message text

```
Attempt to insert duplicate key row in object '%.*s'
with unique index '%.*s'%S_EED
```

Explanation

No two rows can have the same index value (including NULL) in a column or columns with a unique index. Adaptive Server checks for duplicate values when the index is created (if data already exists) and checks each time data is added with an insert or update. Error 2601 occurs when you try to put duplicate index values into a column or columns with a unique index.

Action

Using a unique index makes sense only when uniqueness is a characteristic of the data itself. Choose one of the following solutions, depending on whether or not you need a unique index:

- If you need duplicate index values in the indexed column or columns, drop the unique index and create a nonunique index instead.
- If you need a unique index on data which contains duplicate values, you must change some values to remove the duplicates. To do this:
  a Use a select statement to find the row which will be duplicated by your update or insert command.
  b Modify either the data in the table or the data which you want to insert, so that the index values do not match.

Additional information

Refer to “create index” in the *Reference Manual* for information about indexes.

Version in which this error is raised

All versions

Error 2610

Severity

22

Message text

Version 11.5

```
Could not find leaf row in nonclustered index '%.*s'
```

```
page %ld that corresponds to data row from logical data
page %ld, row offset %d during update index attempt
after data page split in dbid %d.
```

Version 11.0.3 and Earlier

```
Could not find leaf row in nonclustered index '%.*s'
that corresponds to data row from logical data page %ld,
row offset %d during update index attempt after data
page split.
```

Note This error may be due to a hardware problem.

Explanation

This error occurs when Adaptive Server fails to update a nonclustered index at the leaf level after a page split (a page split occurs when a page fills up and there is no more room to write on that page). This error message indicates index corruption.

This is a very serious error. You may be able to recover from this error but you may have to restore from known, clean backups.

Action

1 Record the value of the data page and index name (or ID if that is given instead) from the error message. Using this information and the instructions in Chapter 2, "How to Find an Object Name from a Page Number", determine which table corresponds to the data page number from the error message text.

2 If the table determined in step 1 is a system table (object ID is less than 100) and the index ID is *not* 0, refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index. Then go to step 5.

   If the table determined is step 1 is a system table and the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore from clean backups.

3 If the table determined in step 1 is *not* a system table and the message displayed the index ID rather than the index name, find the name of the index involved by executing the following query:

```
1> select name from sysindexes
2> where indid = index_ID
3> and id = object_ID
4> go
```

   where *object_ID* is the ID of the table as determined in step 1 and *index_ID* is the ID of the index specified in the error message.

4 Drop and re-create the index using the table name obtained in step 1 and the index name determined in step 3.

5 To verify that all problems have been resolved on this table, run dbcc checktable and dbcc tablealloc on the affected table.

Additional information

If problems persist, this procedure may not be sufficient to clean up the index corruption, and you should contact Sybase Technical Support. Have the output from the dbcc checkalloc or dbcc tablealloc and dbcc checktable commands available for review.

Version in which this error is raised

All versions

Error 2615

Severity

14

Message text

```
Attempt to insert duplicate row in table '%.*s' with
index '%.*s' in database '%.*s'. Could drop and recreate
index with ignore duprow or allow duprow.
```

Explanation

Tables with clustered indexes do not allow duplicate rows. Error 2615 occurs when you try to insert a duplicate row in a table which already has a clustered index.

Action

If the duplicate values are a result of incorrect data, change the data and rerun the transaction.

Otherwise, decide whether you want to allow or prevent duplicate rows in the table where the clustered index will be created. Use one of the following procedures.

Note ignore_dup_row and allow_dup_row are mutually exclusive options of the create index command. Do not use them together in the same statement, or you will cause Error 1916.

Allow Duplicate Rows

Create the index with the allow_dup_row option:

```
1> create clustered index index_name on table_name
2> with allow_dup_row
3> go
```

Subsequent update and insert commands can create duplicate rows after the allow_dup_row option is set.

Do not use this procedure if you have a unique index on any column or columns in the table. Unique indexes are only used on tables where no duplicate rows exist.

Delete Duplicate Rows

Create the index with the ignore_dup_row option:

```
1> create clustered index index_name on table_name
2> with ignore_dup_row
3> go
```

Existing duplicate rows will be deleted from the table as the index is created. Subsequent attempts to enter a duplicate row with insert or update are ignored and the insert or update is cancelled with an informational message. If the insert or update is part of a larger transaction, the rest of the transaction will continue as usual.

Additional information

Refer to "create index" in the *Reference Manual* for information about ignore_dup_row and allow_dup_row.

Version in which this error is raised

All versions

Error 2619

Severity

20

Message text

Version 11.0.2 and Later

```
Keys did not match overflow page when inserting row at
end of page (%S_PAGE) that has overflow page (%S_PAGE)
that is not disconnected in dbid = %d.
```

Version 11.0.1 and Earlier

```
Keys did not match overflow page when inserting row at
end of page that has overflow page that is not
disconnected.
```

Explanation

With a clustered index, Adaptive Server maintains data so that the physical order of rows is the same as their logical, indexed order. The bottom or *leaf level* of a clustered index contains the table's actual data pages. When an index key has a large number of duplicates, not all of these may fit on the same data page. Adaptive Server then allocates a new data page, marking it an “overflow page” to store the additional values. The original page and the overflow page now both contain at least one row with that key value.

If the data corresponding to the key value on the first page is now updated or deleted so that the key no longer appears on that page, Adaptive Server removes the association between the two pages. The overflow page is now a regular page just like any other data page.

Error 2619 is raised when Adaptive Server attempts to insert a key value at the end of a page, and the key currently in place has an associated overflow page, but the key does not match the key value in the overflow page. It is likely that the page served as an overflow page at some point, but when a row with that key value on the first page was updated or deleted, the overflow page was not marked as a regular (disconnected) page.

When Error 2619 occurs, your connection to Adaptive Server is broken. It is a serious error as it indicates that you have a corrupt clustered index on the table. However, there is a straightforward solution as discussed under “Action” below.

Action

Use the following steps to resolve the error:

1 For Server 11.0.1 and earlier versions, use the server error log to identify the table on which the error occurred. For Versions 11.0.2 and later, refer to Chapter 2, “How to Find an Object Name from a Page Number” to identify the table.

2 Use sp_helpindex to list all indexes associated with the table. Identify the clustered index on the table.

3 Drop and re-create the clustered index.

4 Run dbcc tablealloc on the table to ensure that the problem is resolved. If errors are reported, contact Sybase Technical Support for assistance.

Additional information

When calling Technical support, have the following information available:

- Server release and SWR Rollup level
- Server error log
- Text of all error messages

Version in which this error is raised: All versions

Error 2620

Severity: 21

Message text:

```
The offset of the row number at offset %d does not match
the entry in the offset table of the following page:
%S_PAGE.
```

Explanation

An offset table is stored at the end of every data page and it stores where rows are located on that page. Adaptive Server uses the offset table to quickly access the rows in each page. Error 2620 occurs during an insert operation when Adaptive Server detects inconsistencies in the offset table of the page specified by the error message.

Note This error may be due to a hardware problem.

This is a serious error. You may be able to recover from this error but you may have to restore from known, clean backups.

Action

1 Record the value of the page from the error message. Using this information and the instructions in Chapter 2, "How to Find an Object Name from a Page Number", determine which table and index corresponds to the page number.

2 If the object encountering the error is *not* a system table (a system table's object ID is less than 100), continue with step 3.

If the object is a system table and the index ID is *not* 0, refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index.

If the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption but it may be necessary to restore from clean backups.

3 For user tables, if the index ID is 0 or 255, continue with step 4.

If the index ID is *not* 0 or 255, translate it into an index name:

```
1> use database_name
2> go
1> select name from sysindexes
```

```
2> where id = object_ID and indid = index_ID
3> go
```

To ensure that the information needed to re-create the index is available, run sp_helpindex on the index prior to dropping it.

Drop the index.

Re-create the index. This clears the corruption in most cases.

Run dbcc checktable on the table to verify that the corruption is gone.

4 If the index ID is 0 or 255, do one of the following:

- Restore the database from clean backups.
- Refer to Chapter 2, "How to Rescue Data from a Corrupted Table".

Warning! Some data on this page might be lost if you recover your table using bcp or select into (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

What to do Before you Dump your Database

Before dumping your database, make sure it is working correctly. More specifically, the following commands should be run prior to each dump:

1 dbcc checkdb.

2 dbcc checkalloc or dbcc checkalloc with the fix option. (Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about running these commands in multi-user mode and how to prevent spurious allocation errors from dbcc commands.)

Additional information

If problems persist, this procedure may not be sufficient to clean up the index corruption, and you should contact Sybase Technical Support with the following information:

- Server release and SWR rollup level
- Server error log
- Text of all error messages
- Output from dbcc checkalloc or dbcc tablealloc and dbcc checkdb

Version in which this error is raised

All versions

Error 2626

Severity

22

Message text

```
Illegal attempt to insert duplicate key row in the
clustered index for object '%.*s' in database '%.*s'.
```

Explanation

Error 2626 is raised when you attempt to insert a row into a table with a clustered index, and the key value being inserted matches an existing index key value, which was not expected.

Error 2626 can be caused by:

- a corrupt clustered index.
- an Adaptive Server problem.

Action

Corrective action depends on the type of table, and whether the table is partitioned. If the problem persists after applying the correction, call Sybase Technical Support.

For Partitioned Tables

1 Unpartition the table:

```
1> alter table <table_name> unpartition
2> go
```

where <table_name> is the object named in the message text.

2 Repartition the table:

```
1> alter table <table_name> partition
<number_of_partitions>
2> go
```

For Non-Partitioned Tables

1 If the object encountering the error is a system table, refer to "How to Fix a Corrupted Index on System Tables" in Chapter 2 for instructions on how to repair the index. Then go to step 5.

2 If the object encountering the error is a user table, identify the clustered index on the table:

```
1> sp_helpindex <table_name>
2> go
```

where <table_name> is the object named in the message text.

3 To ensure that the information needed to re-create the index is available, run sp_helpindex on the index prior to dropping it.

4 Drop the index.

Re-create the index, or create a non-clustered index. This eliminates the 2626 error in most cases.

5 Run dbcc checktable on the table to verify that the corruption is gone.

Version in which this error is raised

All versions

Error 2628

Severity

21

Message text

```
In database '%.*s', page %ld is linked forward to page
%ld, but that page is linked backward to page %ld.
Please report this internal error to Sybase Technical
Support.
```

Note This error may be caused by a hardware problem.

Explanation

When updating a table or index, Adaptive Server may determine that there is not enough room on a page to insert additional data. A new page is then allocated for the object. Existing rows may be moved to this page and new rows inserted. Error 2628 is raised when Adaptive Server fetches a page by following the forward pointers, examines the existing page linkage, and finds that the backward pointer does not point to the previous page as expected. Since the linkage problem could be worsened by continuing the page split, the operation is aborted.

When dbcc detects this problem, it raises Error 2503 (Table Corrupt: Page linkage is not consistent; check the following pages: (current page#=%ld; page# pointing to this page=%ld; previous page# indicated in this page=%ld)).

Action

Error 2628 is due to corruption of page pointers for the object. Use the following steps to resolve the error:

1 Identify the object using one of the pages identified in the message. Refer to Chapter 2, "How to Find an Object Name from a Page Number" for information.

2 If the page is a data page (index ID = 0), drop and recreate the clustered index on the table.

3 If this is an index page (index ID is greater than 0), drop and recreate the index.

4 Run dbcc pglinkage on the table to verify that the problem is corrected. Be sure to check both next page and previous page pointers.

Version in which this error is raised

All versions

Create Utilities Errors (continued)

This section contains error messages for Adaptive Server create operations (create table, create database, and so on).

Error 2714

Severity

16

Message text

```
There is already an object named '%.*s' in the database.
```

Explanation

This error occurs when Adaptive Server tries to create an object which already exists.

Error 2714 can occur in the following situations:

- When you are creating procedures, triggers, tables, views, or rules, if you specify an object that already exists.
- If you have changed your sort order to be case-insensitive and then you try to create an object that is spelled the same way as an existing object with a different case (for example, PROC1 and proc1).
- When using temporary tables, the first 13 bytes of the table name, including the pound sign, must be unique for the current session. Long temporary table names are truncated to 13 characters with the pound sign. An attempt to create a second temporary table with the same name in the same session will raise Error 2714. (In contrast, temporary tables created by specifying database_name.owner.table_name may be shared among sessions, not only within one session.)

Action

Make sure the name of the object you are trying to create does not already exist.

If you have stored procedures that create objects, make sure there is a check for the existence of the object in the stored procedure (before it is created). If you wish to re-create an object within a stored procedure, drop the object before you re-create it.

Additional information

Refer to the *Reference Manual* for information about creating and dropping objects.

Version in which this error is raised

All versions

Error 2729

Severity

16

Message text

```
Procedure '%.*s' group number 1 already exists in the
database.  Choose another procedure name or procedure
number.
```

Explanation

Procedures that are meant to be used together in an application can be created with the same name and different version numbers, using an option of the create procedure statement. For example, procedures used with the application ticket can be named ticketproc;1, ticketproc;2, and so on. The numbers are referred to as the group numbers of the procedure.

Error 2729 occurs when:

- You attempt to create a procedure without specifying any group number, and a procedure by that name already exists. Adaptive Server assumes that you want to create a procedure in the same group as the existing procedure, but the command fails since no group number is specified, and the default group number 1 already exists.
- You attempt to re-use an existing group number when creating procedures with the same name.

Action

To resolve the problem, do one of the following:

- Supply a different procedure name.
- Supply the same procedure name, but include a unique procedure group number. For example:

  ```
  1> create procedure myproc;2 ...
  ```

  This will result in two procedures with the same name and group numbers 1 and 2, respectively.

Version in which this error is raised

All versions

Error 2753

Severity

20

Message text

```
Received an error code of '%d' from 'intl_strblist()'.
```

Explanation

intl_strblist is an Adaptive Server function to convert a message from the server's default language to the client's language. Error 2753 is raised when there is an inconsistency in language between the client session and the server's default language, and *intl_strblist* is unable to convert the message to the client format. The error code denotes the type of conversion problem that occurred.

For example, this error may be raised when you invoke an isql session with the -J option to specify a client character set different from us_english and run certain dbcc commands. The error may occur even if Adaptive Server is configured to use the alternate language.

This is not a serious error and does not cause any problems in your database. Depending on the type of error, Adaptive Server may be able to print the message string in us_english.

Action

Use one of the following actions to address the 2753 error:

- Invoke the client session without the -J option.
- Use set language us_english before operations that produce the 2753 error.

Version in which this error is raised

All versions

Error 2762

Severity

16

Message text

```
The '%s' command is not allowed within a multi-statement
transaction in the '%.*s' database.
```

Explanation

SQL commands are grouped into the following categories:

- SQL commands that are not allowed in transactions at all.
- SQL commands, such as Data Definition Language (DDL) commands, that are allowed in transactions only if the required database option (ddl in tran) is set to TRUE.
- SQL commands that are run across databases to create, alter or drop objects in another database, and are allowed in transactions only if the required database option (ddl in tran) is set to TRUE for that database.

Error 2762 typically occurs in the context of data definition commands such as creating, altering, or dropping objects. It occurs when Adaptive Server detects a command that is not allowed in a multi-statement transaction for the specified database. A multi-statement transaction is a set of commands prefaced with the begin transaction command, or when chained mode is on.

Error 226 is a similar error that may be raised along with Error 2762. While Error 2762 typically involves commands that are run across databases, Error 226 is raised when the command affects only the local database.

The following commands are never allowed in multi-statement transactions:

- alter database
- create database
- dbcc reindex, dbcc fix_text
- disk init
- drop database
- dump database, dump transaction
- load database, load transaction
- reconfigure
- select into
- set transaction isolation level
- truncate table
- update statistics

The following DDL commands are not normally allowed in multi-statement transactions but you can use them if you use sp_dboption to set ddl in tran to TRUE first:

- create default, create index, create procedure, create rule, create schema, create table, create trigger, create view
- drop default, drop index, drop procedure, drop rule, drop table, drop trigger, drop view
- grant
- revoke

Action

The literal `%.*s' in the error message is the name of the database specified in the command. Note the database name before choosing any corrective action.

For Commands Involving *tempdb* Database

Error 2762 can be raised when a command creates or drops temporary objects (objects in tempdb) within the context of a multi-statement transaction. For example, this code may generate a 2762 error:

```
1> begin transaction
2> create table #cities
3> (city_name  char(15)  not null)
4> commit transaction
5> go
```

The error is raised when ddl in tran is set to FALSE in tempdb. Use one of the following strategies to correct this error:

- Use a permanent object name, so that tempdb is not affected.
- Execute the command outside a multi-statement transaction.

Warning! Using data definition language commands on tempdb within transactions may cause concurrency problems in tempdb. Always leave ddl in tran set to FALSE in tempdb.

For Commands Not Involving *tempdb* Database

If the DDL command is allowed in a multi-statement transaction when ddl in tran is set to TRUE, set ddl in tran to TRUE before running the transaction. You can check the current setting of ddl in tran with sp_helpdb.

Warning! Data definition language commands hold locks on system tables such as sysobjects and this can affect performance. Avoid using them inside transactions; if you must use them, keep the transactions short.

If the command is never allowed in a multi-statement transaction, execute it outside the multi-statement transaction.

Additional information

Some applications take SQL statements as input and run them for you. If the application uses begin and commit or rollback transaction to surround those statements, Error 2762 may occur. Refer to the user guide for your application to determine if this is the case.

Version in which this error is raised

All versions

Procedure Manager Errors

This section contains error messages for the Adaptive Server Procedure Manager.

Error 2805

Severity

20

Message text

```
Bad pointer 0x%lx encountered while remapping stored
procedure '%.*s'.  Must re-create procedure.
```

Explanation

When a stored procedure is created, Adaptive Server stores an efficient, normalized form of the procedure called a query tree in the sysprocedures table. When you upgrade the server, the query trees for stored procedures and certain other database objects may be remapped to be compatible with the new release.

Error 2805 is seen in the following situations:

- When you upgrade to version 11.0 or higher from a prior release. The error is due to a remap failure during the upgrade.
- If you used dump database and load database to move data from one platform to another, for example Digital OpenVMS VAX to Digital OpenVMS Alpha. Stored procedures executed on the target platform following such a migration may raise Error 2805.

This error may be accompanied by Error 2812.

Action

If Error 2805 occurs following an upgrade, you can resolve the problem by dropping and re-creating the procedure.

Migration to Another Platform

If you see Error 2805 following an attempt to migrate data from one platform to another using dump database and load database, note that dumps and loads are not a supported mechanism for copying data between servers on different platforms. Instead, you will need to re-create and load the database manually using bulk copy.

For more information, refer to the Adaptive Server installation and configuration guide.

Version in which this error is raised

All versions

Error 2806

Severity

20

Message text

```
Stored procedure '%.*s' is corrupted.  Must re-create
procedure.
```

Note This error may be due to hardware problems.

Explanation

Procedure buffers are data structures that manage Adaptive Server stored procedures (and other compiled objects like triggers, rules, defaults, check constraints, and views) in the procedure cache. The server uses one procedure buffer for every copy of a named object stored in the procedure cache.

The Procedure Buffer Manager maintains an array of procedure buffers in memory, and controls and synchronizes access to procedure buffers. It maintains a usage count to determine whether a procedure buffer is in use.

Error 2806 is raised when the Buffer Manager detects one of the following problems:

- While looking for an available procedure buffer for a given stored procedure or a trigger, a buffer with usage count greater than 1 is found. Buffers for a stored procedure or a trigger are not shared and should never have a usage count exceeding 1.
- While reading rows from sysprocedures, the Buffer Manager finds a header but no additional rows for the object.
- While reading rows from sysprocedures, the Buffer Manager finds out-of-sequence rows for the object.

Error 2806 results from corrupted procedure buffers. It may be related to hardware or operating system problems.

Action

1 If this is an isolated error on a user stored procedure or trigger, drop and re-create the procedure to resolve the problem.

2 If the error occurs on one or more system stored procedures (that is, procedures beginning with the "sp_" prefix), you will need to recreate the sybsystemprocs database where system procedures reside. Refer to "Recovering the sybsystemprocs Database" in the *System Administration Guide*.

3 If other errors are reported (2806 errors during installation of a new server are often accompanied by 605 or other errors), follow the directions in this manual for recovering from those errors.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level
- Server error log
- Text of all error messages.

Version in which this error is raised

All versions

Error 2811

Severity

18

Message text

```
Cannot create procedure dbid %d, objid %ld, with a group
number of %d.
```

Explanation

This error occurs when Adaptive Server is unable to insert a row into the *sysprocedures* table. This can happen during creation of a procedure, view, or trigger, or during remapping of a compiled object.

The most common reason for Error 2811 is lack of data space to add the procedure, and therefore it is often raised following an 1105 error.

Action

Determine what errors occurred before the 2811 error by examining:

- The Server error log
- Errors the user saw when running the application
- The application log file

If an 1105 error occurred, refer to Error 1105 for recovery instructions. If errors other than 1105 occurred, follow the recovery instructions in this manual for those errors. If your manual does not include these instructions, call Sybase Technical Support.

Version in which this error is raised

All versions

Error 2812

Severity

16

Message text

```
Stored procedure '%.*s' not found. Specify
```

```
owner.objectname or use sp_help to check whether the
object exists (sp_help may produce lots of output).
```

Explanation

This error occurs when an attempt is made to execute a stored procedure that does not exist. If the procedure actually does exist, that is, it appears when sp_help is run with no parameters, Error 2812 can occur if you do not fully qualify the procedure name.

Adaptive Server's naming convention for database objects is as follows:

```
database.owner.object_name.column_name
```

The default value for database is the current database and the default value for owner is the current user. Remember that the owner is part of the object, so two different users can have two different procedures with the same object_name in the same database, for example user1.proc1 and user2.proc1.

If the procedure is not owned by the user attempting to execute it, and it is not owned by the database owner ("dbo"), then all references to it must include the owner name. For example, suppose user1 creates a procedure called proc1. Any users other than user1 must prefix the procedure name with the owner name when executing it, that is, exec user1.proc1.

The system stored procedures are an exception to the naming convention. These procedures reside in the sybsystemprocs database, are owned by the Sybase system administrator, and have names starting with "sp_". You can execute system stored procedures from any database and they will reference the system catalog of the current database. The search path for all "sp_" procedures is:

1 The current database

2 The *sybsystemprocs* database

3 The master database

If Error 2812 occurs while running the installmaster script, it means that the **sp_configure** procedure was not found. You can ignore this.

If Error 2812 occurs while you are trying to dump or load a database, you may be trying to dump to an Adaptive Server rather than to a Backup Server.

Action

If you do not know who owns the procedure, use sp_help to display the owner. Run this procedure without any parameters to display objects owned by other users.

Alternatively, you can use the following query to determine who owns the stored procedure you are attempting to execute:

```
1> select name,owner=user_name(uid)
```

```
2> from sysobjects
3> where name = "procedure_name"
4> go
```

If the procedure does not appear in the output of this query, it is either in a different database or does not exist.

If you do not own the procedure, qualify the procedure name with the owner name:

```
1> execute owner_name.procedure_name
2> go
```

For procedures used by many users of a database, it is usually easiest if the database owner creates the procedure. This will allow any user to find the procedure without specifying an owner name.

If the procedure is not in the database where it is executed, fully qualify the procedure name with the database name:

```
1> exec database_name.user_name.procedure_name
2> go
```

The owner name is not needed if you or the database owner own the procedure:

```
1> exec database_name..procedure_name
2> go
```

Execute permission must be provided so other users can execute this procedure, but no permissions are required to see the text of the procedure.

System Stored Procedures

If this error occurs on system stored procedures, it might be resolved by running the installmaster script. This installs all of the system procedures and initializes various other Adaptive Server structures.

Backup Server

If Error 2812 occurs while you are trying to do a dump or load, you may be trying to dump to an Adaptive Server rather than to a Backup Server. Check your Backup Server assignment with the following command:

```
1> sp_helpserver SYB_BACKUP
2> go
```

The network_name column is the interfaces file that contains the entry for this Server.

Check the status of this Server:

```
1> exec SYB_BACKUP...sp_ps
```

```
2> go
```

Open Server applications have the stored procedure sp_ps while Adaptive Servers do not. Therefore, if this command fails because sp_ps is not found, most likely the Server defined for SYB_BACKUP is an Adaptive Server and not a Backup Server. If that is the case:

1 Find out how the Backup Server should be defined.

2 Drop the SYB_BACKUP entry using sp_dropserver.

3 Add the correct entry using sp_addserver.

Additional information

Refer to the *Transact-SQL User's Guide* for information about naming conventions.

Refer to the *Reference Manual* for information about sp_dropserver and sp_addserver.

Version in which this error is raised

All versions

Error 2824

Severity

```
19
```

Message text

```
Process %d cannot hold PROC_BUF 0x%1x named '%.*s'
because it holds another PROC_BUF 01%1x named '%.*s'. A
process can only hold one view, rule, or default at a
time.
```

Explanation

Error 2824 can be caused by the following situations:

A process can hold many procedures and triggers but it can only hold one view, rule, or default at a time. If a process already holds a view, rule, or default, and Adaptive Server tries to add another one for the process, Error 2824 is raised.

Error 2824 can be caused in the following situations:

- A view referenced by stored procedures or triggers has been dropped and re-created. (Stored procedures and triggers are not reresolved when a view they reference is dropped and re-created.)
- Procedure cache corruption has occurred.

Action

Use sp_help view_name to determine whether a view that stored procedures or triggers reference has been dropped and re-created. If the creation time of the view is later than the creation time of the procedure or trigger that refers to it, drop and re-create the stored procedures and triggers that reference that view.

If the 2824 error persists, there is probably procedure cache corruption. Stop and restart Adaptive Server to clear the corruption.

Version in which this error is raised

All versions

Error 2835

Severity

16

Message text

```
The procedure tree is old.  Use the sp_remap procedure
to remap all the procedures in this database.
```

Explanation

When you execute a stored procedure, Adaptive Server checks to determine whether the procedure is already in cache. If the procedure is not in cache, Adaptive Server reads the procedure into cache from sysprocedures and recompiles it. Before Adaptive Server recompiles the procedure, it checks the Server release number of the procedure to make sure it is up to date for the current release. Error 2835 occurs when the procedure is not up to date for the current release.

This error can occur because the query remapping phase of an upgrade failed. If that is the case, all the query trees might need to be remapped.

Action

Use sp_remap to remap each procedure, trigger, rule, default, and view in the current database:

```
1> use database_name
2> go

1> sp_remap object_name
2> go
```

where *database_name* is the name of the database where the object resides and *object_name* is the name of the object to remap.

If you have many objects to remap, refer to Chapter 2, "How to Remap All Objects in a Database".

Additional information

Refer to the *Reference Manual* for information about sp_remap.

Version in which this error is raised

All versions

dump and *load* Errors

This section contains error messages for Adaptive Server dump and load commands.

Error 3101

Severity

16

Message text

```
Database in use. A user with System Administrator (SA)
role must have exclusive use of database to run load.
```

Explanation

You cannot use load database while a database is in use by any user (even yourself). Error 3101 occurs when you try to load a database dump while other users are still accessing the database.

Note This error does not occur on master.

Action

Wait until all the users have finished using the database, then run the load database command again.

Make sure that you are not in the database being loaded while you issue the load database command. Sybase recommends running all load statements from within the master database, although any database other than the one being loaded can be used.

Version in which this error is raised

All versions

Error 3105

Severity

16

Message text

```
Data on dump will not fit into current database. Need
%d Mbyte database.
```

Explanation

This error occurs when an attempt is made to load a database dump into a database smaller than the dumped database. You cannot use dump database or load database to shrink a database, regardless of the amount of actual data in the database.

Action

If your goal is to load the database from the dump, you must increase the size of the database being loaded into to at least the configured size of the dumped database. You can do this with either the alter database command or by dropping and re-creating a larger database.

If your goal is to shrink the total size of a database, you cannot use dump and load. You must copy all the data out of the database using bcp, drop the database and re-create a smaller database, then re-create the database structure with scripts, and reload the data with bcp.

Additional information

The database being loaded into must have identical data and log mapping, and identical segment mappings, as the dumped database. Refer to Error 2558 and the *System Administration Guide* for more information.

Version in which this error is raised

All versions

Error 3120

Severity

16

Message text

```
The database you are attempting to LOAD was DUMPed under
a different sort order ID (%d) or character set ID (%d)
than the ones running on this server (SortOrd = %d,
CharSet = %d). If the sort orders differ, at least one
of them is non-binary.
```

Explanation

A database cannot be loaded into a Server using load database unless the sort order of the Server that dumped the database matches the sort order of the Server being loaded into. Error 3120 occurs when you try to load a database on an Adaptive Server with a different sort order from that of the dumped database.

Action

You have the following options:

- Move the data via bcp rather than load.
- Change the sort order of the Adaptive Server. Refer to “Configuring Character Sets, Sort Orders, and Languages” in the *System Administration Guide* for instructions.
- If the server into which you are attempting to load is a new server, you can:
  - Reinstall the server with the correct sort order
  - Re-create the databases with the load option

- Retry the load

Version in which this error is raised: All versions

Error 3201

Severity: 16

Message text: No dump device has been specified.

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server is unable to access a dump device during a database dump . This error is fatal and stops the dump from completing.

Error 3201 is most likely caused by one of the following:

- The device you specified for the dump is offline or otherwise unavailable to Adaptive Server.
- Permissions for the dump device are not set correctly for user “sybase” or the user performing the dump (read and write privileges are required).
- A previous dump aborted and Adaptive Server believes the dump device is still in use.

Action

1 Check to make sure the device you specified is defined for your Adaptive Server:

```
1> select * from master..sysdevices
2> where status = 16 or status = 24
3> go
low  high   status   cntrltype   name        phyname     mirrorname
---  ------ -------  ----------  ----------  ----------  ---------
0    20000  16       3           tapedump1   /dev/rmt4   NULL
0    20000  16       2           tapedump2   /dev/rst0   NULL
```

If the device is not there, you can use sp_addumpdevice to add it.

2 At the operating system level, check the permissions for the dump device for user “sybase” or the user performing the dump (read and write privileges are required).

3 If you have still not solved the problem, check your operating system error log for operating system errors. If no errors are logged in your operating system error log and there are no other errors in your Adaptive Server or Backup Server error log, contact Sybase Technical Support for assistance.

Additional information

Refer to "Specifying the Database and Dump Device" in the *System Administration Guide* for information about dump devices.

If you call Technical Support, have the following information on hand:

- Server version and SWR Rollup level
- Backup Server version
- Adaptive Server and Backup Server error logs
- Text of all error messages
- `select * from sysdevices` output
- Operating system error log

Version in which this error is raised

All versions

Error 3203

Severity

`16`

Message text

`Backup Server has detected a SQL Server error.`

Note This error may be caused by a hardware problem.

Explanation

Error 3203 is displayed when a problem occurs with a remote procedure call to the Backup Server. In this case, the error is on the Adaptive Server side.

The error message output includes:

- *vsn* – the virtual socket number.
- *return* – the return value: 0 means successful; -2 means failure.
- *status* – the ending status, displayed in decimal value. The most common value is 524288 which usually means an I/O error.

Some causes are:

- Write protect is turned on (it must be off when loading a tape because the device is opened read/write).
- No dump exists on the media being accessed (for example, a blank tape or the wrong device specified).
- The media that contains the dump is not readable.

Action

- Check the Adaptive Server and Backup Server error logs to determine the cause of the error being sent from the Adaptive Server.
- Test the connection between the Backup Server and the Adaptive Server by logging into the Adaptive Server through isql and typing:

```
1> execute backupserver...sp_ps
2> go
```

  where *backupserver* is the name of your Backup Server. This executes sp_ps on the Backup Server.

Additional information

If you need to call Technical Support, have the following information on hand:

- Server version and SWR Rollup level
- Backup Server version
- Adaptive Server and Backup Server error logs
- Text of all error messages
- Operating system error log

Version in which this error is raised

All versions

Error 3211

Severity

20

Message text

```
DUMP failed in database %.*s:  error occurred executing
RPC.
```

Explanation

When a dump database command is processed, the Backup Server scans for allocated pages to determine what pages to dump. It begins this scan as soon as it receives an RPC (remote procedure call). Error 3211 is raised when an Adaptive Server RPC to the Backup Server returns an error.

Error 3211 occurs with the following states:

State	Meaning
1	An error occurred in the RPC stream.
2	The RPC itself returned an error.

Action

- Check the Adaptive Server and Backup Server error logs to determine the cause of the error being sent from Adaptive Server.
- Test the connection between the Backup Server and the Adaptive Server by logging into the Adaptive Server through isql and typing:

```
1> execute backupserver...sp_ps
2> go
```

  where *backupserver* is the name of your Backup Server. This executes sp_ps on the Backup Server.

Additional information

If you need to call Technical Support, have the following information on hand:

- Server version and SWR Rollup level
- Backup Server version
- Adaptive Server and Backup Server error logs
- Text of all error messages
- Operating system error log

Version in which this error is raised

All versions

Error 3212

Severity

20

Message text

```
DUMP failed in database %.*s:  unexpected status
executing RPC: %d.
```

Explanation

Dumps and loads are performed by an Open Server program known as Backup Server. Adaptive Server communicates with Backup Server using Remote Procedure Calls (RPC's). When a dump database command is processed, for example, Adaptive Server uses RPC's to send parameters to Backup Server. Error 3212 is raised when an Adaptive Server RPC to the Backup Server returns an unexpected status which Adaptive Server cannot process.

Action

- Check the Adaptive Server and Backup Server error logs to determine the cause of the error.

- Test the connection between the Backup Server and the Adaptive Server by logging into the Adaptive Server through isql and typing:

  ```
  1> execute backupserver...sp_ps
  2> go
  ```

  where *backupserver* is the name of your Backup Server. This executes sp_ps on the Backup Server.

Additional information

If you need to call Technical Support, have the following information on hand:

- Server version and SWR Rollup level
- Backup Server version
- Adaptive Server and Backup Server error logs
- Text of all error messages
- Operating system error log

Version in which this error is raised

All versions

Error 3216

Severity

10

Message text

```
Backup Server session id is:  %ld.  Use this value when
executing the 'sp_volchanged' system stored procedure
after fulfilling any volume change request from the
Backup Server.
```

Explanation

Dumps and loads are performed by an Open Server program known as Backup Server. During the backup and restore process, it may be necessary to change tape volumes. If Backup Server detects the need to change the currently mounted volume, it requests a volume change by sending messages to either the client or its operator console. After mounting another volume, the operator notifies Backup Server by executing the sp_volchanged system procedure on Adaptive Server.

3216 is an informational message that Backup Server issues when starting a dump or load operation. The message notifies the operator of the current session ID; the operator needs this information if it becomes necessary to execute sp_volchanged during the dump/load operation.

Action	This is an informational message. No action is necessary.
Version in which this error is raised	All versions

Error 3225

Severity	21
Message text	`I/O failed while attempting to clear pages starting at page %ld of disk number %ld.`

Note This error may be caused by a hardware problem.

Explanation	During a load database, create database, or alter database command, Adaptive Server initializes the non-allocated database pages. Error 3225 occurs if an I/O error occurs during this process. This error can be caused by an operating system problem such as insufficient disk space (for a file system) or a disk going offline.
Action	Check your operating system error logs to determine whether you have run out of disk space or are having other problems that could explain the I/O errors. Refer to “Checking the Operating System Error Log” in Encyclopedia of Tasks for assistance.
Version in which this error is raised	All versions

Error 3230

Severity	24
Message text	`For logical disk device '%.*s', neither primary device '%.*s' nor secondary device '%.*s' is active. %s cannot continue.`
Explanation	During dumps and loads, Adaptive Server must translate the logical device names into physical device names and pass these to Backup Server, which performs the dump/load. For each device, Adaptive Server first checks whether the primary physical device is in the active state; if not, it checks whether the secondary (mirror) device is active.

Error 3230 is raised when neither the primary nor secondary devices are active. This is a fatal error which breaks your connection to Adaptive Server.

The error may be caused by:

- the device being offline
- an Adaptive Server problem, which results when you create a device using a previously used *vdevno*.

Action

Examine the Adaptive Server error log and the operating system log to see if the device(s) used in the dump/load are offline or otherwise unavailable to the server. See Chapter 2, "Checking the Operating System Error Log" for assistance.

If the device(s) are active, shut down and restart Adaptive Server, and retry the dump/load operation.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level
- Text of all error messages.

Version in which this error is raised

All versions

Error 3233

Severity

```
17
```

Message text

```
%s for database '%S_DBID' failed:  an alarm function
could not be installed.
```

Explanation

Dumps and loads are performed by an Open Server program named Backup Server. In preparation for the dump, Adaptive Server must install an alarm function, which is a timing mechanism required during the dump. Using the alarm function, Backup Server can go into a `sleep' state at an appropriate point and awaken periodically to check the progress of the dump.

Error 3233 is raised when you perform a dump, but there are not enough alarms available to allow Adaptive Server to install an alarm function for the dump, causing the dump to abort.

Action

The available alarms are controlled by the configuration parameter number of alarms. To resolve the problem, change the value of number of alarms to a number larger than the current value.

Refer to No More Alarms Available Error for directions on increasing the number of alarms.

Since the waitfor command utilizes alarms, frequent use of the command can also lead to 3233 errors. Check use of this command in statement blocks, stored procedures and transactions.

Version in which this error is raised

All versions

Error 3240

Severity

18

Message text

```
Unable to obtain the SQL Server's current working
irectory.  The SQL Server must have search permission
on its current working directory and the name must be
shorter than %d bytes.
```

Explanation

Dumps and loads are performed by an Open Server program known as Backup Server. When you issue a dump database or load database command, Adaptive Server communicates certain information about its working environment to Backup Server. Error 3240 is raised when Adaptive Server is unable to determine its current working directory for Backup Server.

Action

This error commonly occurs when Adaptive Server is started from a directory other than the sybase home directory; the user ID of the user starting Adaptive Server may not have the appropriate permissions for the directory from which it was started.

To resolve the error, check that the *SYBASE* variable is set to the correct sybase home directory and is no longer than 255 bytes.

Make corrections if necessary, and restart Adaptive Server.

Additional information

Scripts that start up Adaptive Server must always set the directory to the sybase home directory before starting Adaptive Server.

Version in which this error is raised

All versions

Commit and Abort Errors

This section contains error messages for Adaptive Server commit and abort transaction operations.

Error 3301

Severity 21

Message text `Invalid log record found in syslogs (logop %d).`

Note This error may be caused by a hardware problem.

Explanation

The syslogs table is the Adaptive Server transaction log, containing a record of changes made in the database in the order in which they occur. Adaptive Server scans the records in the transaction log during database recovery and certain other activities. Error 3301 occurs if Adaptive Server encounters an invalid record while scanning the transaction log. *logop* refers to the type of record (log operation) encountered. For example, a *logop* of 18 is logged by a save transaction command.

Error 3301 is raised with the following states:

State	Meaning
1	An unexpected log record was encountered during recovery.
2	An unexpected log record was encountered while rescanning portions of the log during recovery.
3	Recovery encountered a savepoint record it did not expect to find. The record marks the beginning of a portion of a transaction that is to be rolled back (and should not be seen during recovery).
4	An obsolete record was encountered during recovery.
5	While scanning a block of records for a transaction, recovery encountered a savepoint record pointing back to a previous savepoint record which was missed during the scan.
6	Recovery encountered a savepoint record it did not expect to find. The record marks the end of a portion of a transaction that is to be rolled back (and should not be seen during recovery).
7	Nested abort records were found in the log during a rollback.
8	An unexpected log record was found during a rollback.

State	Meaning
9	An unexpected log record was encountered while translating log records associated with an obsolete log format (for a previous Adaptive Server version) to the format for the current Server version.
10	An out-of-range logop number was encountered while translating log records associated with an obsolete log format (for a previous Adaptive Server version) to the format for the current Server version.

This is a serious error as it indicates a problem with the transaction log. Corrective action depends on the state with which the error is raised.

Action

A 3301 error during recovery marks the database suspect. The database is identified in the error log. Restore the database from a good backup as follows:

1 Drop the database. If the drop fails, follow the instructions in "How to Drop a Database When drop database Fails".

2 Create a database for load. Make sure the database you create has sizes as least as large as those in sysusages for the original database (and that all other sysusages values match the original values). For more information, refer to "create database" in the *Reference Manual*.

3 Load the database from backup. (Refer to "load database" in the *Reference Manual*.)

4 Use the online database command to make the database available for use.

A 3301 error with these states indicates a run-time error when Adaptive Server attempted to roll back a single transaction. Contact Sybase Technical Support for assistance.

A 3301 error with these states typically occurs when upgrading via a load. To avoid this error a standard procedure such as upgrading via sybinit, or bulk copying data out of the old server and into the new server, is recommended.

Additional information

Refer to "Developing a Backup and Recovery Plan" in the *System Administration Guide* for complete information about how to safely create, dump, load, and re-create databases.

Version in which this error is raised

All versions

Error 3307

Severity 21

Message text

```
Process %d was expected to hold logical lock on page %ld
in database %S_DBID for object %S_OBJID, index number
%d.
```

Explanation

This error occurs when Adaptive Server, while committing or aborting a transaction, attempts to release a lock on the page displayed by the error message and the page is not locked.

This error may occur during recovery as well as during run time.

Common causes of this error are:

- A software failure occurred.
- When Adaptive Server was performing a rollback, it expected a logical lock on the page displayed in the error message but the lock could not be found.
- When Adaptive Server attempted to deallocate a page, it expected a logical lock on that page but one did not exist.
- Adaptive Server tried to undo a page deallocation.
- Adaptive Server tried to roll back a user transaction which has executed a stored procedure that required reresolution. Reresolution is required for stored procedures that reference objects which have been dropped and re-created between executions.

Action

Look for other errors in the Adaptive Server error log as well as in your operating system error log to find out the specific source of the problem and clear those errors first, as they might be the actual cause of the error. Restarting Adaptive Server will release any locks that active transactions might still hold. However, this will not clear the root cause of this error.

If there are no other errors in the Adaptive Server error log or in your operating system error log, contact Sybase Technical Support for assistance.

Additional information

For more information about transactions, refer to transaction in the *Transact-SQL User's Guide*.

Before calling Technical Support, have the following information available:

- Server version and SWR Rollup level
- Server error log
- Hardware error log
- Output of sp_lock and sp_who before restarting Adaptive Server

- Text of all error messages

Version in which this error is raised

All versions

Recovery Errors

This section contains error messages for Adaptive Server database recovery. Database recovery occurs during Adaptive Server start-up, load database, and load transaction.

Error 3401

Severity

21

Message text

```
Rec_logbounds: getnext SCAN_RID of last checkpoint
failed on Rid from sysdatabases. %S_RID.
```

Explanation

When you restart Adaptive Server, recovery scans the log to find the last checkpoint record to determine the active portion of the log. Error 3401 is raised when recovery is unable to find a checkpoint record.

Error 3401 occurs with the following states:

State	Meaning
1	An attempt to start a log scan failed.
2	A log scan was started, but an attempt to find a checkpoint record in the log failed.

Action

Good Backup Available

If you have a good backup, restore the database as follows:

1 Drop the database. If the drop fails, follow the instructions in Chapter 2, "How to Drop a Database When drop database Fails".

2 Create a database for load. Make sure the database you create has sizes as least as large as the values in sysusages for the original database (and that all other sysusages values match the original values). For more information about creating a database for load, refer to create database in the *Adaptive Server Reference Manual*.

3 Load the database from backup. (Refer to load database in the *Adaptive Server Reference Manual*.)

4 Use the online database command to make the database available for use.

No Good Backup Available

If you do not have a good backup, contact Sybase Technical Support. They may be able to help you recover from this error.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level
- Text of all error messages.

Version in which this error is raised

All versions

Error 3403

Severity

22

Message text

```
During recovery initialization, page %ld was
encountered.  This page belongs to object %ld, not the
log.
```

Note This error may be caused by a hardware problem.

Explanation

During recovery, Adaptive Server scans the transaction log until the last page is found. During this scan, if a page is found that is allocated to syslogs but the object ID on the page header is not the same as that of syslogs, Error 3403 occurs.

Situations that may cause this error include the following:

- A bad page allocation or write occurred due to an Adaptive Server problem.
- Adaptive Server went down after the allocation page was updated but before the transaction log page was written. The database was later rebuilt without the transaction log pages being cleared and entries from the old log still exist.

This is a serious error and you will have to rebuild the affected database either using bcp or from clean backups.

Action

If you have a good backup, restore the database from your backup:

1 Drop the database. If the drop fails, follow the instructions in Chapter 2, "How to Drop a Database When drop database Fails".

2 Create a database for load. Make sure the database you create has sizes as least as large as those in sysusages for the original database (and that all other sysusages values match the original values). For more information about creating a database for load, refer to "create database" in the *Adaptive Server Reference Manual*.

3 Load the database from backup. (Refer to "load database" in the *Reference Manual*.)

4 Use the online database command to make the database available for use.

No Good Backup

If you do not have a good backup, copy the data from the corrupted table into a new (dummy) table or into a file. Then rename your old, corrupted table and copy the data back into a new table using the original name. For information about doing this, refer to Chapter 2, "How to Rescue Data from a Corrupted Table".

To prevent some occurrences of this error in the future, checkpoint each database that is being used before shutting down Adaptive Server.

Additional information

Refer to "Developing a Backup and Recovery Plan" in the *System Administration Guide* for complete information about how to safely create, dump, load, and re-create databases.

Version in which this error is raised

All versions

Error 3404

Severity

21

Message text

```
Rec_complete: Could not open controlling database (id
%d) of controlling database in multi-db transaction.
```

Explanation

The Adaptive Server transaction log records all changes in the database. When database activity occurs in the context of a transaction, BEGINXACT and ENDXACT records are written in the transaction log to mark the transaction boundary. If the transaction spans databases, it is known as a multi-database transaction; the database where the transaction was started (in the case of explicit transactions, the database where the first begin transaction was issued) is referred to as the controlling database.

When Adaptive Server is brought back up after a shutdown, the recovery process uses the transaction log to bring all the databases to a consistent state. Error 3404 is raised when Adaptive Server attempts to recover a multi-database transaction, but is unable to open the controlling database for the transaction.

Specifically, Error 3404 results when Adaptive Server is attempting to recover a transaction that started in a user database but modified data in master database. When this happens, Adaptive Server is unable to access the transaction log for the user database since the device on which the log resides is not yet open. Consequently master cannot be recovered. This error is frequently accompanied by Errors 913 and 3414.

Action

Use one of these options to correct this problem:

- Restore the database specified in Error 3404 from a clean backup.
- Contact Sybase Technical Support. They may be able to help you re-create the database in question.

To avoid a recurrence of this error, do not start transactions in user databases that modify tables in master.

Additional information

Before calling Technical Support, have the following information available:

- Server release and SWR Rollup level
- Server error log
- Text of all error messages

Version in which this error is raised

All versions

Error 3414

Severity

```
10
```

Message text

```
Database '%.*s' (dbid %d): Recovery failed. Check the
SQL Server errorlog for further information as to the
cause.
```

Explanation

This error occurs during startup when Adaptive Server could not complete the recovery of the database listed in the error message.

You cannot use the database until whatever caused the error has been corrected because Adaptive Server marks the database suspect.

Error 926 is a related error which is raised when attempting to access a database that could not recover. Refer to Error 926.

Action

To determine why recovery failed, examine the Adaptive Server error log for any errors prior to the 3414 error. It is important to identify the errors before the first occurrence of the 3414 error because subsequent attempts to start Adaptive Server may not give the detailed error information you need to diagnose the problem.

If you do not have sufficient information to recover from the previous errors, you can recover from known, clean backups or contact Sybase Technical Support for assistance.

Version in which this error is raised

All versions

Error 3418

Severity

21

Message text

```
Not enough deses to open '%.*s'.
```

Explanation

This error occurs during startup when Adaptive Server could not complete recovery of the system database listed in the error message.

Error 3418 is called for the following system databases which existed in sysdatabases prior to starting recovery:

- model
- sybsecurity
- sybsystemprocs

Possible reasons for Error 3418 include:

- During database recovery, Adaptive Server failed to find the corresponding row in sysdatabases (State 1 or State 3 is returned with the 3418 error).
- Adaptive Server was unable to open a database for recovery (State 2 is returned with the 3418 error).

When Error 3418 occurs, Adaptive Server shuts down automatically. Since the affected databases are required for Adaptive Server to start successfully, manual intervention is required to start Adaptive Server when this error occurs.

Action

To determine why recovery failed, examine the Adaptive Server error log.

Contact Sybase Technical Support for assistance with the manual intervention required to start Adaptive Server and recover from this error.

Version in which this error is raised

All versions

Error 3425

Severity

21

Message text

```
Transaction (%ld, %d) not found in transaction table.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs during recovery when an end (commit or rollback) transaction log record was found that does not have a corresponding begin transaction record. Therefore, the transaction could not be rolled back or committed and recovery or load could not complete for that database.

You cannot use the affected database until whatever caused the error has been corrected because Adaptive Server marks the database suspect.

Action

If the error occurred during start-up (rather than during load database or load transaction), determine which database had the error by looking at your Adaptive Server error log.

If you have a clean backup, restore your database using that backup.

If you do not have a clean backup, call Sybase Technical Support.

Additional information

If you need assistance from Sybase Technical Support, have the following information available when you call:

- Server release and SWR Rollup level
- Server error log
- Text of all error messages

Version in which this error is raised

All versions

Error 3429

Severity

21

Message text

```
Error recovering database '%.*s' - could not connect to
commit service to check completion status of xact:
%S_RID.
```

Explanation

The two-phase commit service allows an application to coordinate updates within a Server or among two or more Adaptive Servers. The commit service uses one Adaptive Server, the *commit server*, as a central recordkeeper that helps the application determine whether to commit or roll back transactions in case of failure.

During database recovery, each outstanding transaction is handled in one of the following ways:

- If the transaction committed, it is redone.
- If the transaction aborted or failed to complete, it is rolled back.

If the transaction reached the prepare-to-commit stage of the two-phase commit, Adaptive Server must find out whether the controlling transaction committed. If the controlling transaction is at the same site, Adaptive Server can directly examine the log. But if the transaction was handled by the commit service, Adaptive Server must query the commit service to find out whether the controlling transaction has been committed.

Adaptive Server queries the commit service using the probe utility. As a standard client, probe requires the following:

- A valid interface must exist for the commit server.
- The commit service must be running and not in single-user mode.
- The probe version must be the same as the Adaptive Server version.
- In version 11.0.x, the probe binary must exist in the *$SYBASE/bin* directory. (Not required in versions 11.5 and higher.)

Error 3429 occurs when Adaptive Server uses probe to query the commit service to find out whether the controlling transaction committed, and probe cannot connect to the commit service because the connection is refused or times out.

After Error 3429 occurs, Error 3414 is raised and the database status is set to “suspect.”

Action

1 Solve the problem that is preventing probe from connecting to the commit service. Make sure that the commit service Adaptive Server is running. If it is hung or otherwise inaccessible, shut down and restart the commit service Adaptive Server.

2 For the affected database, execute one of the procedures supplied in Chapter 2, "How to Reset a Database's "suspect" Status".

3 Shut down and restart Adaptive Server to recover the database.

Additional information

Refer to the *Open Client DB-Library/C Reference Manual* for information about the two-phase commit service.

Version in which this error is raised

All versions

Error 3434

Severity

20

Message text

```
Cannot change sortorder. Server shutting down.  Restart
to continue with sortorder unchanged.
```

Explanation

When you restart a server after changing the server's sort order, the recovery process must rebuild system tables in both master database and user databases, and rebuild indexes that are affected by the sort order change. Error 3434 occurs when recovery is unable to recreate a system index or table following the sort order change.

The error can be raised:

1 When you change from a case-sensitive sort order to a case-insensitive sort order, and conflicts arise between key values that are differentiated only by case. For example, key values 'Joe' and 'joe' will cause a conflict when rebuilding the index with a case-insensitive sort order.

   Error 3436 ("Cannot rebuild index %d for the '%.*s' table in the '%.*s' database.") precedes Error 3434 in this situation.

2 When one or more databases for load exist in the server. You cannot change sort order when the server contains a database that was created for load.

Error 3434 rolls back the transaction in progress, reverts to the original sort order, and shuts down the server.

Action

Check the error log to obtain more information about the error. Actions depend on the likely cause of the error:

1. If the error is related to key value conflicts, determine the table and index from the 3436 error message, and check for '*duplicate key*' messages in the error log. Correct the reported conflict and retry the sort order change.
2. If the server contains a database created for loading a dump (for load database), take the following steps to correct the problem:
   a. Drop the for-load database(s).
   b. Manually change the sort order on the server. See "How to Manually Change Sort Order or Default Character Set" in Chapter 2 for details.
   c. Restart the server.

After changing the sort order, refer to "If You Changed the Sort Order or Default Character Set" in the *System Administration Guide* and do the steps described there. It is very important that you do these steps to guarantee the integrity of your data.

Additional information

1. You can achieve similar results without changing the sort order to case insensitive by using the upper or lower functions in your queries. For example, if tableX has col1 with values "sybase", "Sybase", and "SYBASE", you could run the following isql statement:

   ```
   select cols from tableX where lower(col1) = "sybase"
   ```

   This treats all of the values as lowercase for the query.
2. In many cases, you cannot reload your data from a database dump after reconfiguring the sort order. Refer to "Database Dumps and Configuration Changes" in the *System Administration Guide* for details.

Version in which this error is raised

All versions

Error 3445

Severity

```
10
```

Message text

```
SQL Server could not bring database '%.*s' online.
```

Explanation

This error occurs when Adaptive Server cannot bring a database online for one of the following reasons:

- An attempt to update the database log version fails. Adaptive Server tries to update the log version if it is earlier than the current Adaptive Server log version.
- Adaptive Server failed to clear the offline status bit after an upgrade.

Action

Look at the error messages that precede the 3445 message to determine why the 3445 error occurred, and resolve those problems.

If you cannot solve the problem:

1 Determine the current log version for the database. For example, to determine the log version of a database called test_db, use the following commands:

```
1> load database test_db from 'test.dump' with headeronly
2> go
Backup Server session id is:  6.  Use this value when executing
the `sp_volchanged' system stored procedure after fulfilling any
volume change request from the Backup Server.
Backup Server: 6.28.1.1: Dumpfile name `test_db952820A2F8' section
number 0001 mounted on disk file
`/remote/solaris/rel1100/install/test.dump'
This is a database dump of database ID 6, name `test_db', from
Oct  9 1995 11:35AM.  SQL Server version:
SQL Server/11.0/B/Sun_svr4/OS5.2/1/OPT/Fri Aug 1805:10:26 PDT 1995.
Backup Server version:
Backup Server/11.0/B/Sun_svr4/OS5.4/1/OPT/Thu Aug 17 21:54:21 PDT 1995.
Database contains 1536 pages; checkpoint RID=(Rid pageid = 0x405;
row num = 0xd); next objectID=3031; sort order ID=50, status=0;
charset ID=1.
Database log version=2; database upgrade version=1.
```

The *log version* and *upgrade version* have values as follows:

Version	Value Is 0	Value Is 1	Value Is 2
Database log version	10.0.x and earlier	10.1	11.0 or later
Database upgrade version	10.x or earlier	11.0 or later	Not used

2 Call Sybase Technical Support. Have the following information ready:

- Server version and SWR Rollup level
- Server error log
- Text of all error messages
- Output from step 1.

Version in which this error is raised	All versions

Error 3446

Severity	16
Message text	`You do not have privilege to bring database '%.*s' online.`
Explanation	The online database command marks a database available for public use after a normal load sequence and, if needed, upgrades a loaded user database and transaction log dumps to the current version of Adaptive Server. Only a System Administrator, Database Owner, or Operator can execute online database. Error 3446 occurs when you do not have any of these privileges and you attempt to execute an online database command.
Action	Make sure you have the needed privileges or ask a user who does have the needed privileges to run the online database command.
Additional information	Refer to the *Reference Manual* for information about online database. Refer to the *Security Administration Guide* for information about using the "sa" account and roles in Adaptive Server. Refer to "grant" in the *Adaptive Server Reference Manual* for information about granting roles.
Version in which this error is raised	All versions

Error 3447

Severity	10
Message text	`Database `%.*s' appears to be in the process of being loaded; SQL Server will not bring it online automatically. Use the ONLINE DATABASE command to bring this database online.`

Explanation	During a load database, Adaptive Server takes the database being loaded offline. When the load is complete, you bring the database back online manually with the online database command. If your backup strategy involves loading a sequence of transaction logs after a load database, the fact that the database will be unavailable for use until you enter the online database command will allow you to complete your loads without interference from other processes changing that database. Error 3447 occurs during the recovery phase of Adaptive Server start-up, when the recovery process attempts to bring all the recovered databases online automatically and Adaptive Server sees that the database named in the error message is being loaded. This error does not occur for master, as Adaptive Server always brings master online.
Action	Finish the load database sequence and then use online database to make the database available for use.
Additional information	Refer to the *Reference Manual* for information about online database.
Version in which this error is raised	All versions

Error 3449

Severity	10
Message text	`Database '%.*s': beginning upgrade step: %S_MSG %s [ID %d].`
Explanation	This informational message tells you that Adaptive Server is beginning an upgrade step. The text of the message includes the following information about the object Adaptive Server is upgrading: • Database name (*%.*s*) • Which upgrade step is being performed (*%S_MSG*) • The name of the object being upgraded (*%s*) • The task ID of the Server internal upgrade function (*ID %d*) The message also tells you the index name, if applicable.
Action	No action is required.

Version in which this error is raised: All versions

Error 3452

Severity: 20

Message text:

```
Database '%.*s': upgrade item %d depends on item %d,
which could not be installed. Please refer to previous
messages for the cause of the failure, correct the
problem and try again.
```

Explanation

Databases can be upgraded during an online database command or when you are using the sybinit upgrade utility. Many upgrade steps depend on other upgrade steps having been done previously.

Error 3452 occurs when the first upgrade step shown in the message depended on the second reported step, and that second step has failed. The database upgrade does not complete when this error occurs.

You will not be able to bring the database named in the error message online until you resolve the problem that led to the 3452 error.

Action

1 Find the previous message describing the failure of the dependent (second) step.

2 Fix the problem that caused the failure. If appropriate, follow the directions for a failed upgrade in the Adaptive Server installation and configuration guide.

3 If you were running online database when Error 3452 occurred, issue the online database command to restart the database upgrade.

Version in which this error is raised: All versions

Error 3454

Severity: 10

Message text:

```
Database '%.*s': SQL Server could not completely upgrade
this database; upgrade item %d could not be installed.
```

Explanation

Databases can be upgraded during an online database command or when you are using the sybinit upgrade utility. Error 3454 is an informational message telling you that Adaptive Server could not completely upgrade the database listed in the message. The 3454 error is preceded by other errors that show why the upgrade failed.

You cannot bring the database named in the error message online until you resolve the problem that led to the 3454 error.

In the following example, these messages appeared during start-up, when Adaptive Server attempted to bring a database online but failed to upgrade it:

```
00:95/11/21 18:34:36.51 server Database 'test' appears to be at an older
revision than the present installation; SQL Server will assess it,
and upgrade it as required.
00:95/11/21 18:34:36.59 server Error: 1105, Severity: 17, State: 1
00:95/11/21 18:34:36.61 server Can't allocate space for object
'sysattributes' in database 'test' because the 'system' segment is full.
If you ran out of space in syslogs, dump the transaction log.
Otherwise, use ALTER DATABASE or sp_extendsegment to increase
the size of the segment.
00:95/11/21 18:34:36.63 server Error: 3460, Severity: 20, State: 1
00:95/11/21 18:34:36.65 server  Database 'test': upgrade could not
record the installation of upgrade item '80'. Please refer to
previous error messages to determine the problem.
Fix the problem, then try again.
00:95/11/21 18:34:36.68 server Error: 3451, Severity: 20, State: 1
00:95/11/21 18:34:36.70 server Database 'test': upgrade has failed for
this database.  Please refer to previous messages for the cause of
the failure, correct the problem and try again.
00:95/11/21 18:34:36.73 server Error: 3454, Severity: 20, State: 1
00:95/11/21 18:34:36.75 server Database 'test': SQL Server could not
completely upgrade this database; upgrade item 80 could not
be installed.
00:95/11/21 18:34:36.76 server SQL Server could not bring database
'test' online.
```

Action

1. Look at the error messages preceding the 3454 error message to determine what happened.
2. Refer to the recovery steps suggested for those errors.
3. Correct the problem.
4. If the error occurred when you were running online database, issue the online database command for that database after the problem is fixed. Adaptive Server will automatically try to finish the upgrade.

If the error occurred during an Adaptive Server upgrade, follow the recovery directions in the Adaptive Server installation and configuration guide.

Version in which this error is raised

All versions

Error 3470

Severity

10

Message text

```
SQL Server could not completely upgrade database '%.*s',
but the database was online when upgrade began, so it
will be left online.
```

Explanation

When Adaptive Server tries to bring an already online database to the online state and it finds that some elements of the database have not been upgraded, it tries to complete the database upgrade. Error 3470 occurs when that attempt to upgrade the database fails.

Action

1 Determine the current upgrade version for the database. For example, to determine the upgrade version of a database called test_db, use the following commands:

```
1> load database test_db from 'test.dump' with headeronly
2> go

Backup Server session id is:  6.  Use this value when executing
the `sp_volchanged' system stored procedure after fulfilling any
volume change request from the Backup Server.
Backup Server: 6.28.1.1: Dumpfile name `test_db952820A2F8'
section number 0001 mounted on disk file
`/remote/solaris/rel1100/install/test.dump'
This is a database dump of database ID 6, name `test_db', from
Oct  9 1995 11:35AM.  SQL Server version:
SQL Server/11.0/B/Sun_svr4/OS5.2/1/OPT/Fri Aug 1805:10:26 PDT 1995.
Backup Server version:
Backup Server/11.0/B/Sun_svr4/OS5.4/1/OPT/Thu Aug 17 21:54:21 PDT 1995.
Database contains 1536 pages; checkpoint RID=(Rid pageid = 0x405;
row num = 0xd); next object ID=3031; sort order ID=50, status=0;
charset ID=1.
Database log version=2; database upgrade version=1.
```

The *log version* and *upgrade version* have values as follows:

Version	Value is 0	Value is 1	Value is 2
database log version	10.0.x and earlier	10.1	11.0 or later
database upgrade version	10.x or earlier	11.0 or later	Not used

2 Call Sybase Technical Support. Have the following information ready:

- Server version and SWR Rollup level
- Server error log
- Text of all error messages
- Output from step 1
- Output from sp_configure "upgrade version".

Version in which this error is raised

All versions

Error and Exception Handling Errors

This section contains error messages for Adaptive Server error and exception handling.

Error 3604

Severity

10

Message text

```
Duplicate key was ignored.
```

Explanation

This error occurs when you try to insert duplicate key values in a column or columns which have a unique index that was created with the option **ignore_dup_key**.

Adaptive Server ignores the commands in a transaction which try to insert a duplicate index. The rest of the commands in the transaction will complete normally.

Action

No action is necessary.

If you want the row inserted, use one of the following procedures:

- Eliminate the unique index.
- Remove duplicate values. If the duplicate values are a result of incorrect data, change the data and rerun the transaction.

Additional information

Refer to “create index” in the *Reference Manual* for information about indexes.

Version in which this error is raised

All versions

Error 3621

Severity

10

Message text

```
Command has been aborted.
```

Explanation

This error occurs when Adaptive Server is unable to complete a command. Error 3621 may be the last message displayed by your application even though another message was sent by Adaptive Server when the error actually occurred. The rest of the commands in the transaction will complete normally.

Some situations that may cause Error 3621 include:

- Division by 0 in your application
- Entry of a duplicate row during use of your application
- Other illegal functions in your application

Action

Look for an error before the 3621 error occurs. Resolve the problem that is causing the error before the 3621 error.

If client applications receive this error, ask your Adaptive Server system administrator to check the Adaptive Server error log.

Version in which this error is raised

All versions

Error 3624

Severity

10

Message text

```
Truncation error occurred.
```

Explanation

The precision of a datatype denotes the maximum number of digits allowed in columns of that datatype. For the exact numeric types dec, decimal and numeric, the allowed range of data depends on the column's precision as well as its scale, that is, the maximum number of digits that are allowed to the right of the decimal point.

Error 3624 occurs when an operation inserts data into a target column or variable but the precision or scale of the target are too small for the data. The operation fails and the command is aborted.

Action

Examine the data involved in the arithmetic operation to ensure that the target column or variable is sized appropriately.

Additional information

The set arithabort and set arithignore options specify behavior following loss of precision or (for exact numeric types) loss of scale. For more information see:

- "Arithmetic Errors" in the *Transact-SQL User's Guide.*
- The *Reference Manual.*

Version in which this error is raised

All versions

Error 3626

Severity

17

Message text

```
The transaction was aborted because it used too much
stack space. Either use sp_configure to increase the
stack size, or break the query into smaller pieces.
spid: %d, suid: %d, hostname: %.*s, application name:
%.*s
```

Explanation

At start-up, Adaptive Server allocates one stack for each configured user connection. These stacks are in contiguous areas of memory, with a guard area at the end of each stack. Adaptive Server periodically checks to determine whether the stack point for a user connection has entered the stack guard area adjacent to that connection's stack. Error 3626 occurs when this has happened.

Because Error 3626 occurs before any stack corruption, Adaptive Server does not shut down. Instead, it aborts the user process and rolls back any transactions that existed for that process.

An Adaptive Server process usually encounters this error for one of two reasons:

- The stack size is too small. The information placed on the stack is valid but its total size exceeds the specified stack size. This is often caused by a single, complex query which requires more than the average amount of stack space. Long where clauses, long select lists, and deeply nested stored procedures can contribute to this type of stack overflow and Error 3626.
- An incorrectly handled error results in recursive calls which eventually fill the stack, regardless of its size.

If this error occurs frequently, there is increased risk a process may overflow its stack area and corrupt another process' stack. This situation will cause Adaptive Server to shutdown with the error:

```
kernel:  *** Stack guardword corrupted.
```

See Stack Guardword Error for more information.

Action

If this error occurs with a transaction which contains a where or where-in clause, you may need to rewrite the query. See Error 404 for additional actions which may be required.

If this error occurs because a complex query has run out of stack area, you can resolve it in one of two ways:

- Break up the query into smaller pieces.
- Use the following query to increase the stack size:

```
1> sp_configure "stack size", new_value
2> go
```

Shut down and restart Adaptive Server to make these changes take effect.

When changing the stack size on Adaptive Server, consider the following facts:

- The stack size should always be an even increment of 2K (2,048 bytes). If a specified stack size is not divisible by 2K, sp_configure will round it up to the next 2K increment.
- Increase the stack size in small increments to avoid making it too large. If the stack size is too large, you will lose cache space and you may be unable to start Adaptive Server.
- Adaptive Server requires more memory for each configured user connection and device when the stack size increases. Therefore, increasing the stack size without increasing the total memory configuration parameter will result in less total cache space for Adaptive Server. This can affect performance.

Additional information

Refer to "sp_configure" in the *Reference Manual* and "Setting Configuration Parameters" in the *System Administration Guide* for information about sp_configure and stack size and stack guard size configuration parameters.

Refer to the *Performance and Tuning Guide* and "Configuring Memory" in the *System Administration Guide* for information about configuring Adaptive Server memory.

Version in which this error is raised

All versions

drop Errors

This section contains error messages for Adaptive Server drop commands.

Error 3701

Severity

11

Message text

```
Cannot drop the %S_MSG '%.*s', because it doesn't exist
in the system catalogs.
```

Explanation

This error occurs when you try to drop an object that is not found in at least one system table where Adaptive Server expected to find it.

Error 3701 can occur due to the following circumstances:

- The object you are trying to drop does not exist
- Inconsistent system catalog tables
- An Adaptive Server problem has occurred

Action

Make sure you entered the object name correctly in your drop command.

If you are entering the object name correctly and the drop fails with Error 3701, try to re-create the object. If your create command displays the message:

```
Msg 2714, Level 16, State 1:
Line 1:
There is already an object named 'object_name' in the
database.
```

then your system tables are incorrect with regard to this object. If this occurs, run dbcc checkcatalog and dbcc checkdb. Then call Sybase Technical Support. They will probably be able to help you delete the object that is causing the error. However, because other objects may reference that object, deleting it cleanly may be difficult. If this is the case, you may need to recover from backups.

Additional information

Before calling Technical Support, have the following information available:

- Server release and SWR Rollup level
- Text of all error messages
- Output of dbcc checkdb and dbcc checkcatalog

Version in which this error is raised

All versions

Error 3702

Severity

16

Message text

```
Cannot drop the %S_MSG '%.*s', because it is currently
in use.
```

Explanation

This error occurs when Adaptive Server tries to drop a database, user table, procedure, rule, default, trigger, or view that is currently in use.

Error 3702 can occur in the following circumstances:

- During a drop command such as drop database, drop table, drop procedure, or drop rule
- During a dbcc dbrepair (database_name, dropdb) command

Error 3702 can occur for the following reasons:

- The object you are trying to drop is in use
- The process using the object was abnormally terminated
- Data corruption has occurred
- An Adaptive Server problem has occurred

Action

Use sp_who to determine whether the database or object name in your drop command is currently in use. If it is in use, contact the user shown in the sp_who output.

If the process using the object was abnormally terminated (for example, with an operating system kill command), the object may not have been completely removed from the Adaptive Server system tables. The process would remain invisible while keeping the object open and in use. This prevents the drop command from succeeding. In this case, shutting down and restarting Adaptive Server should clear the 3702 error. After you restart Adaptive Server, run dbcc checkcatalog to determine whether corruption has occurred.

If Error 3702 persists, call Sybase Technical Support.

Additional information

Before calling Technical Support, have the following information available:

- Server release and SWR Rollup level
- Text of all error messages
- Output of dbcc checkdb, dbcc checkcatalog, and dbcc checkalloc for the database that contains the object named in the error message

Version in which this error is raised

All versions

Error 3703

Severity 11

Message text

```
Cannot drop the %S_MSG with object-id %ld in database
%d, because it doesn't exist in the system catalogs.
```

Explanation

This error occurs when Adaptive Server fails to drop a database, user table, procedure, rule, default, trigger, or view because the object being dropped does not exist in the appropriate system table.

Error 3703 can occur during a drop command such as drop database, drop table, drop procedure, or drop rule.

Error 3703 occurs with the following states.

State	Meaning
1	During a drop trigger command, if Adaptive Server fails to find the trigger or the target table of the trigger, Error 3703 occurs with State 1.
2	During a drop index command, if Adaptive Server fails to find the object ID of the table for which the index is being dropped in sysobjects, Error 3703 occurs with State 2.
4	During any drop command, if Adaptive Server fails to find a sysobjects entry for the object being dropped, Error 3703 occurs with State 4.
5	The sysconstraints table has one row for each referential and check constraint associated with a table or column. You can drop constraints using the alter table command. If, during such an alter table command, Adaptive Server fails to find an entry for the constraint in sysconstraints, Error 3703 occurs with State 5.
6	The sysreferences table has one row for each referential integrity constraint declared on a table or column. If the object being dropped is a referential integrity constraint, Adaptive Server searches sysreferences for the referential integrity constraint ID. If the ID is not found, Error 3703 occurs with State 6.
7	When Adaptive Server fails to drop a database, user table, procedure, rule, default, trigger, or view that is not some other object's subsidiary part (for example, dropping a trigger because its associated table is being dropped), and the object ID is not found in sysobjects, Error 3703 occurs with State 7.
8	If the database being dropped has referential dependencies, Adaptive Server tries to drop the references from the other databases. If the ID of the reference from the other database is not found in sysreferences, Error 3703 occurs with State 8.

Action

Recovering from Error 3703 requires manually modifying one or more system tables. Call Sybase Technical Support for assistance in determining which tables to modify and how to do it.

Additional information

Before calling Technical Support, have the following information available:

- Server release and SWR Rollup level

- Text of all error messages

Version in which this error is raised

All versions

Error 3704

Severity

16

Message text

```
Cannot drop the %S_MSG '%.*s' because you are not the
owner or a user with System Administrator (SA) role.
```

Explanation

Error 3704 occurs when Adaptive Server is unable to carry out a command, either implicit or explicit, to drop a database object. The error can involve temporary or permanent database objects.

Error 3704 with Permanent Database Objects

Error 3704 is raised when you attempt to drop a permanent table or other object, but you are not authorized to do so because you are not the owner of the object or a user with System Administrator (SA) role.

Error 3704 is raised in the following states:

State	Meaning
1	The object being dropped is a user table, view, stored procedure or trigger.
2	The object being dropped is an index.
3	The object being dropped is a user database.

Error 3704 with Temporary Tables

Error 3704 can be raised if Adaptive Server is unable to automatically drop temporary tables created within a stored procedure.

Error 3704 with temporary tables is due to an Adaptive Server problem.

Action

If Error 3704 is raised for operations involving permanent objects, this means you have encountered a permission problem when trying to drop the object. Check with your System Administrator.

If Error 3704 is raised for operations involving temporary tables, contact Sybase Technical Support for assistance.

Additional information

If Errors 3704 is raised repeatedly and accompanied by Error 216 ("Attempt to automatically drop temporary table failed."), Adaptive Server may become unusable. If this happens, shut down Adaptive Server using shutdown with nowait and restart. This will clear unnecessary tables from tempdb.

If you need to call Technical Support, have the following information on hand:

- Server version and SWR Rollup level
- Server error log
- Text of all error messages
- Text of stored procedures involved
- Operating system error log

Version in which this error is raised

All versions

Error 3731

Severity

16

Message text

```
Cannot drop table '%.*s' because it is partitioned.
```

Explanation

By default, Adaptive Server stores a heap table's data in one doubly linked chain of database pages. Adaptive Server inserts all new rows into the last page of the chain. A transaction holds an exclusive lock on the last page while inserting new rows, which can block other, concurrent transactions from being inserted into the table.

The partition clause of the alter table command allows you to partition user tables that do not have a clustered index. Partitioning creates additional page chains on the table, each with its own last page. This reduces page contention for concurrent inserts. It can also reduce I/O contention if the table is distributed over multiple physical devices.

Error 3731 is raised when you try to drop a table that is partitioned.

Action

If you want to drop the table, use the unpartition clause of the alter table command to concatenate all partitions:

```
1> alter table table_name unpartition
2> go
```

Then drop the table.

Additional information

Refer to the *Reference Manual* for information about alter table.

Version in which this error is raised	All versions

Transaction Errors

This section contains error messages for Adaptive Server transaction processing.

Error 3904

Severity

21

Message text

```
Can't unsplit logical page %ld in object '%S_OBJID' in
database '%S_DBID' - both pages together contain more
data than will fit on one page.
```

Explanation

Data and index pages are split when Adaptive Server determines that a row cannot be inserted due to insufficient room on the existing page. During transaction rollback (either by a user request or during database recovery), Adaptive Server tries to undo the changes done by the transaction. Error 3904 occurs when Adaptive Server fails to unsplit pages which were split during the transaction.

Error 3904 is caused by an Adaptive Server problem.

When this error occurs, data integrity may have been compromised.

Action

1 Use the procedure in Chapter 2, "How to Find an Object Name from a Page Number" to identify which table and index correspond to the page number in the error message text.

2 If the object encountering the error is *not* a system table (a system table's object ID is less than 100), continue with step 3.

If the object with the error is a system table and the index ID is *not* 0, refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index.

If the object with the error is a system table and the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore from clean backups.

3 For a user table, if the index ID is *not* 0, run the sp_helpindex procedure on the index to get information for rebuilding it.

Drop the index.

Re-create the index. This clears the corruption in most cases.

Run dbcc checktable on the table to verify data integrity. If errors occur, call Sybase Technical Support.

4 For a user table, if the index ID is 0, run dbcc checktable on the table to verify data integrity. If errors occur, restore the database from clean backups.

Additional information

Refer to drop index and create index in the *Reference Manual* for information about dropping and re-creating indexes.

Refer to "Developing a Backup and Recovery Plan" in the *System Administration Guide* for information about how to safely create, dump, load, and re-create databases

Version in which this error is raised

All versions

Error 3905

Severity

21

Message text

```
Can't unsplit logical page %ld in object '%S_OBJID' in
database '%S_DBID' - row number %d is used on both
pages.
```

Explanation

When Adaptive Server determines that a row cannot be inserted due to insufficient room on the existing page, data and index pages are split to accomodate the insertion. If the transaction is rolled back (either by a user request or during database recovery), Adaptive Server tries to undo the changes done by the transaction. The undo steps include restoring the offset table which contains the location of rows on the page.

Error 3905 occurs when Adaptive Server is unable to restore the offset table on a data page, thus failing to unsplit pages which were split during the transaction.

Error 3905 is due to an Adaptive Server problem and is sometimes seen when buffer pools larger than 2K are in use. When this error occurs, data integrity may have been compromised.

Action

Use the steps described below to recreate the affected table, then call Sybase Technical Support for help in resolving the server problem:

1 If the object in question is a system table (a system table name starts with 'sys' and the object ID is less than 100), call Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore from clean backups.

   If the object is a user table, continue with Step 2.

2 Use sp_help(*<table_name>*) to determine the table's lock scheme: allpages-locked (APL) table or data-only-locked (DOL) table.

3 Identify the row(s) corresponding to the affected page using:

```
dbcc page(<dbid>, <page_num>, 1)
```

   where *dbid* and *page_num* are the database ID and page number in the error message respectively.

4 Use one of the following methods to recover from the problem:

   - Select the table data into a new table, drop the old table, and rename the new table to the old table name. Omit the affected rows from Step 3 when selecting data from the old table.

     For a DOL table, if there is an index on the table, force this index when copying data.

   - Bulk copy the affected table out, drop and re-create the table, and bulk copy back in. This is the most efficient solution for a very large table. Omit the affected rows from Step 3 when copying in.

5 Run dbcc checktable to check the new table.

Additional information

Avoid using named caches or large I/O buffer pools until the problem is resolved.

Have the following information ready when you call Sybase Technical Support:

- Adaptive Server version and SWR rollup level.
- Server error log.

Version in which this error is raised

All versions

Error 3906

Severity

16

Message text

```
Attempt to BEGIN TRANsaction in database '%.*s' failed
because database is READ ONLY.
```

Explanation

Error 3906 is raised when a transaction, explicit or implicit, is attempted while the database is in read-only mode.

Error 3906 is raised with the following states:

State	Meaning
1	Attempted to start a transaction when the database is in read-only mode. Applies to all possible transaction types.
2	Attempted to start a subordinate transaction (child transaction started by a parent transaction in a parallel setting) when the database is in read-only mode.

Error 3906 can be raised in three distinct cases, described below.

Case 1: User Transaction Attempted in Read-Only Database

A user may have attempted to start a transaction while the database is in read-only mode. This is not allowed.

Case 2: During ALTER DATABASE (Version 11.0.3.2 and Earlier)

On Adaptive Server Version 11.0.3.2 and earlier, this error may be raised:

- When you use alter database to extend log space on a read-only database. The database is altered, and log space added, but the server is unable to start a transaction to update the last-chance threshold (LCT) for the database.
- On a warm standby server which is periodically re-loaded with dumps from a production server. If the production database grows in size, loading dumps to the warm standby may require extending the log. The database is successfully altered, and log space is added, but the server is unable to start a transaction to update the last chance threshold for the database.

Note On later versions, you will see other messages when extending the log on a warm standby. See Additional Information below for details.

Case 3: Re-resolution in a Read-Only Database (All Versions)

The error may occur when a view or procedure must be resolved in a read-only database. Re-resolution must update the sysprocedures table, but since the database is read-only, Error 3906 is raised.

Action

Corrective action depends on the scenario in which Error 3906 was raised, as described in the above cases.

Case 1: User Transaction in Read-Only Database (All Versions)

Re-try the transaction when the database is no longer read-only.

Case 2: During *alter database* (Version 11.0.3.2 and Earlier)

If you altered a read-only database and the log segment was extended:

1 Determine if the LCT needs to be modified. If this database is normally operated in read-only mode, no updates occur on the database and so there is no need to update the LCT.

However, if this is a production database which was temporarily in read-only mode during ALTER DATABASE, you should re-establish the LCT after the database is made available for updates. Continue with Step 2.

2 Use the following command to determine the new last-chance threshold:

```
1> select lct_admin("reserve", <log size in pages>)
2> go
```

For example, on a 20MB transaction log (containing 10240 2KB pages):

```
1> select lct_admin("reserve", 10240)
2> go
 -----------
616
```

The resulting value is the amount of space required to dump the transaction log.

3 Modify the LCT using:

```
1> select lct_admin('lastchance', <new threshold>)
2> go
```

Refer to *Creating a Free-Space Threshold for the Log Segment* in the *System Administration Guide* for guidelines on setting thresholds.

Case 3: Object Re-resolution in a Read-Only Database (All Versions)

If Error 3906 was raised when you attempted to select from a view or execute a stored procedure, you can correct the problem by making the database temporarily available for writes and re-compiling objects. Keep in mind, however, that if the database is used as a warm standby, this activity will create additional records in the log, which will invalidate the database as a warm standby.

To re-resolve compiled objects in the database:

1 Put the database in read-write mode.

```
1> sp_dboption <database_name>, "read only", false
2> go
```

2 Execute all compiled objects such as stored procedures and views once.

3 Return the database to read-only mode.

```
1> sp_dboption <database_name>, "read only", true
2> go
```

Additional information

Different versions of Adaptive Server differ in their handling of alter database (to effect log extends) in a warm standby database. As described above, version 11.0.3.2 and earlier raise Error 3906. On early releases of 11.0.3.3, alter database may fail, raising Error 5033 (“Attempt to alter database ’%S_DBID’ failed because it is in read-only mode.”).

Later versions raise Error 5034 ("Cannot modify the last-chance threshold because the database ’%S_DBID’ is in read-only mode.") followed immediately by Error 5035 ("Remember to execute SELECT LCT_ADMIN(’LASTCHANCE’, %d), when the database is made writable."). Follow the steps in Case 2 above.

Starting with Version 11.9.x, which introduced the standby access mode to allow access to the database during load sequences, you can run alter database while in standby access. However since transactions cannot be started in this mode, you must re-establish the LCT after the database is fully online following the completion of the load sequence. Follow the steps in Case 2 above.

Version in which this error is raised

All versions

Error 3908

Severity

16

Message text

```
Attempt to BEGIN TRANsaction in database ’%.*s’ failed
because database is in BYPASS RECOVERY mode.
```

Explanation

Adaptive Server can start various types of transactions such as local (in response to a local command), internal (for the server’s own use), external (started by an external client like a TP Monitor) and subordinate (child transactions started by a parent transaction in parallel).

Bypass recovery starts Adaptive Server without recovering one or more databases. A database in this mode can be accessed even though it is not recovered. You use this mode to allow access to the database for problem analysis, to copy out data, and so on.

Error 3908 is raised when a transaction is attempted while the database is in bypass recovery. Although certain commands are permissible in this mode, no transactions - regardless of type - can be started when the database is in bypass recovery.

Error 3908 is raised with the following states:

State	Meaning
1	Attempted to start a transaction when the database is not available for update. Applies to all possible transaction types.
2	Attempted to start a subordinate transaction when the database is not available for update.

Action

Do not attempt to start a transaction while the database is in bypass recovery. Check if any automated jobs may have submitted a batch job that attempted such a transaction. First restore the database to the original working status it had (prior to bypass status) before attempting any transactions.

Use the following steps to reset the database status:

Note Although a status of zero is usually appropriate, it is not the only possible status value you can use. Reset the database to the actual, normal status it had prior to entering bypass recovery status.

1 Reset the database status to 0:

```
1> sp_configure "allow updates", 1
2> go

1> begin tran
2> go

1> use master
2> go

1> update sysdatabases set status=0
2> where name=<database_name>
3> go
```

Check that the above update command affected only one row (if more than one row was affected, issue a rollback transaction.) Then commit the transaction and shut down Adaptive Server as follows:

```
1> commit tran
2> go

1> shutdown with nowait
2> go
```

2 Restart Adaptive Server. Run dbcc checkdb and dbcc checkalloc on the affected database to make sure there are no errors.

Version in which this error is raised: All versions

Error 3917

Severity: 16

Message text:

```
An attempt was made to startup a multi-database update
in database '%S_DBID' after starting up the master
transaction in tempdb. This could render the database
'%S_DBID' unrecoverable and this update is therefore not
allowed.
```

Explanation: Error 3917 is raised when you initiate a transaction in tempdb and the transaction attempts to update one or more tables in another database. For example:

```
1> use tempdb
2> go
1> begin tran
2> go
1> update geo..nation
2> set n_comment="This is a grade V earthquake zone"
3> where n_nationkey = 4
4> go
Msg 3917, Level 16, State 1:
Line 1:
An attempt was made to startup a multi-database update in database
'geo' after starting up the master transaction in tempdb.
This could render the database 'geo' unrecoverable and this
update is therefore not allowed.
```

You cannot start update transactions in tempdb because tempdb is truncated upon recovery. If the transaction is left in a prepare state and it becomes necessary to restart the server, this will render the user database (the one being updated) unrecoverable.

3917 errors may also be seen when the tempdb transaction log gets filled up, and are accompanied by Error 1105 (Can't allocate space for object '%.*s' in database '%.*s' because '%.*s' segment is full/has no free extents. If you ran out of space in syslogs, dump the transaction log. Otherwise, use alter database or sp_extendsegment to increase size of the segment.).

Action

Check the ASE error log to help determine the circumstances in which the error is raised.

1 If Error 3917 was raised when you attempted a transaction containing an update or insert statement, it means that the transaction was incorrectly initiated from tempdb. Take one of the following steps, as applicable:

   - Add 'use <*database_name*>' in the SQL script to change the database from tempdb to the database being updated, and retry the transaction; or
   - Change the default database for the user executing the transaction to the database where the insert is taking place.

2 If Error 1105 in tempdb accompanies the 3917 error, this is probably due to an Adaptive Server problem which prevents the checkpoint process from executing because the transaction log is filled up. Shut down and restart the server. Call Sybase Technical Support to upgrade to a version in which the problem is fixed.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Adaptive Server version and SWR rollup level
- Server error log

Version in which this error is raised

All versions

dataserver Errors

This section contains general error messages for the Adaptive Server dataserver utility program.

Error 4001

Severity 11

Message text

```
Cannot open default database '%.*s'.
```

Explanation

This error is raised when Adaptive Server is unable to open the user's default database at login.

After raising the error, the server checks if the user is allowed to log in to master. If the master database is forbidden, the login request is denied.

Common reasons why the default database cannot be opened include:

- The database may not be available for use. For example, it may be in recovery.
- The database may be up, but not available for general use. For example, it may be undergoing failover as part of the High Availability (HA) subsystem (Version 12.0).
- Permission checks failed. For example, if the database was created by loading from another server, but the login information in source and target servers does not match, login may fail.

Action

Retry the login in the event that the database was in recovery or temporarily unavailable.

Since the 4001 error is usually a symptom of some other problem, check the server error log for any other errors that may have been raised prior to 4001. Check for the following errors and retry:

- 921 errors (Database '%.*s' has not been recovered yet - please wait and try again.) indicate that the database is in recovery, or is marked suspect due to an earlier problem. Refer to the Error 921 writeup in this manual for instructions.
- 2206 errors (Database %.*s with dbid %d is already shut down.) indicate that the database is not available for general use. These errors are raised in Version 12.0 and higher. Check for other HA subsystem messages to determine the cause.

- 10351 errors (Server user id %d is not a valid user in database '%.*s') indicate a permissions problem. If the database was loaded from another server, the problem may be due to an incorrect suid/uid combination in sysusers. Check syslogins entries in source and target server, respectively. Check the *Security Administration Guide* and "mismatched suids" in the *System Administration Guide* for more information.

Version in which this error is raised: All versions

Error 4002

Severity: 14

Message text: `Login failed.`

Explanation: This error occurs when a login request is rejected by Adaptive Server.

Some reasons for Error 4002 are:

- User name does not exist in syslogins (names are case-sensitive).
- An incorrect password has been entered (passwords are case-sensitive).
- The account for the user trying to log in is locked.
- The user is trying to log in to an Adaptive Server that:
  - does not exist, or
  - the user does not have access to, or
  - is not running.
- Adaptive Server is in single-user mode and only "sa" can login.
- Adaptive Server shutdown is in progress and only "sa" can login.
- Remote Adaptive Server names are translated from the login record to a site ID. If the Adaptive Server name is not found in sysservers, then the login fails.
- A remote login is being attempted and sysremotelogins is missing.
- A remote login is being attempted and the user cannot be found in sysremotelogins.
- For "sa," your Replication Server password is different from your Adaptive Server password.

Action

Determine whether any of the reasons listed above are causing the 4002 error. If you have checked each of these and you still cannot login to Adaptive Server, call Sybase Technical Support.

Following are instructions for checking on the above reasons:

- Login name does not exist in syslogins (names are case-sensitive).
  - When Adaptive Server is first installed, the only login is “sa” and it has no password.
  - If using isql, make sure the user is specifying the -U option and that the correct name and case are being specified.
  - Check syslogins to make sure the user is spelling the login name correctly:

    ```
    1> select * from syslogins
    2> go
    ```

- An incorrect password has been entered (passwords are case-sensitive).
  - If using isql, make sure the user is specifying the -P option and the correct password and case are being specified.
  - If a user cannot remember his or her password:

    Change the password using sp_password (only the System Security Officer can change someone else's password):

    ```
    1> sp_password caller_passwd,
    2> new_passwd, login_name
    3> go
    ```

    where *caller_passwd* is the password of the login account that is currently executing sp_password and *new_passwd* is the new password for *login_name* (the name of the user whose password you are changing).
  - If you cannot remember your “sa” password and you cannot get into Adaptive Server with any other login:

    Run dataserver -p to generate a new password for a System Security Officer's account. You start Adaptive Server with -p, immediately log into Adaptive Server with the new random password, and execute sp_password to reset your password to a more secure one. Refer to the Adaptive Server utility programs manual for details.
- The account for the user trying to log in is locked. To determine whether this is the case, type:

```
1> sp_displaylogin login_name
2> go
```

where *login_name* is the login name of the user. Check the value for the *status* column:

- If the value is 1, the password is less than 6 characters or NULL.
- If the value is 2, the account is locked.
- If the value is 3, the account is expired.

To unlock a user account:

```
1> sp_locklogin login_name "unlock"
2> go
```

(Only System Administrators and System Security Officers can use sp_locklogin to lock and unlock logins.)

- The user is trying to log into a Server that does not exist or that the user does not have access to. Check the user's environment file to make sure the value of *DSQUERY* is defined to be the Server the user is trying to log in to.
- When Adaptive Server is in single-user mode, only “sa” can login.
  - Refer to Chapter 2, “Returning Adaptive Server to Multiuser Mode”.
  - Check your runserver file to make sure the -m option is *not* specified.
- When Adaptive Server shutdown is in progress, only “sa” can login. Tell the user to wait until Adaptive Server has come back up.
- Remote Server names are translated from the login record to a site ID. If the Server name is not found in sysservers, then the login fails. Use sp_helpserver to list the Servers in sysservers.
- If a remote login is being attempted and sysremotelogins is missing, then the login fails. To determine whether this is the case, type:

  ```
  1> sp_helpdb sysremotelogins
  2> go
  ```

- If a remote login is being attempted and the user cannot be found in sysremotelogins, then the login fails. To determine whether this is the case, type:

  ```
  1> sp_helpremotelogin
  2> go
  ```

- For “sa,” your Replication Server password is different from your Adaptive Server password. Make sure the “sa” passwords are the same.

Additional information

For information about logins and the procedures mentioned in this writeup, refer to the *Security Administration Guide.*

Version in which this error is raised

All versions

Error 4020

Severity

20

Message text

```
While reading characters from the client, the receiving
buffer has been detected to be too small.
```

Explanation

All client/server communication occurs over a network via packets. Clients initiate a connection to the server. The connection sends client requests and server responses.

Error 4020 is raised when the server detects an overflow in the destination buffer while reading one or more characters from the client. The error breaks the connection to the server.

Possible reasons for Error 4020 include:

- application errors
- corrupt packets being transmitted
- incorrect network configuration

Action

Collect all relevant information about the circumstances in which the error occurred: character set differences on client and server, changed or new applications, recent changes in the network configuration, and so on.

Check for coding errors in the application which may be causing incorrectly sized packets to be sent.

Check for possible packet corruption within the network.

Check the network configuration. For example, if there is 7-bit communication at one end, and 8-bit at the other end, check that the network is properly configured to handle this setup.

Check that Adaptive Server is correctly configured to work with the underlying network packet size or buffer size on your system. Relevant server parameters include:

```
default network packet size
maximum network packet size
```

`total memory`

Version in which this error is raised: All versions

dump and *load* Errors (continued)

This section contains error messages for Adaptive Server dump and load commands.

Error 4204

Severity

17

Message text

```
Unable to continue logged version of DUMP TRAN. No space
left in database; use NO_LOG option or run ALTER
DATABASE. Else, an old transaction may be active; see
Syslogshold table.
```

Explanation

This error occurs when Adaptive Server runs out of space in the transaction log during a dump transaction command.

Action

You have the following options for recovering from Error 4204:

- If any old transactions are active, determine whether they can be terminated. If such transactions do exist and they can be terminated, then you can dump the transaction log, which may free up space in the database.
- Perform a dump transaction with the no_log option. dump transaction with no_log removes the inactive part of the log without making a backup copy and without recording the procedure in the transaction log (performs a checkpoint).
- Extend the transaction log segment.

Each of these options is described below.

Check for Old Transactions

In the master database, there is a table called syslogshold. This table contains information about each database's oldest active transaction (if any) and Replication Server truncation point (if any) for the transaction log. This table is built dynamically when you query it.

Check syslogshold for old transactions for the database for which the error occurred:

```
1> use master
2> go
1> select * from syslogshold
2> where dbid = database_ID
```

```
3> go
```

Determine whether the oldest active transaction can be terminated (perhaps it was left active unintentionally). If you can terminate this transaction and then execute a dump transaction with no_log command, space may be returned to the database. Continue this process until there are no other old transactions that can be terminated. If, after terminating some old transactions, you still get Error 4204, try one of the other options described in this writeup.

Recovery Using *dump transaction with no_log*

1 Use the following method to determine approximately how many pages the transaction log occupies:

```
1> use database_name
2> go

1> select data_pgs (8, doampg)
2> from sysindexes where id=8
3> go
```

where "8" is the ID for syslogs. The result is the number of data pages (2K pages on most operating systems) that the transaction log occupies.

Note The query results may be inaccurate by as many as 16 pages, but using the query is much faster than counting the rows in the transaction log.

2 Dump the transaction log using the no_log option.

3 Repeat step 1. If the transaction log now occupies significantly fewer pages, continue with the next step now.

If, however, there are still a large number of pages in the syslogs table, an outstanding transaction is probably preventing the log from being cleared. If this is the case, do one of the following:

- Follow the instructions in Chapter 2, "How to Detect and Clear Long-Running Transactions".
- Restart Adaptive Server and repeat step 2. When Adaptive Server starts and this database is recovered, the outstanding transaction is rolled back, allowing the log to be cleared by a subsequent dump transaction.

4 Dump the database now, because the no_log option prevents subsequent changes recorded in the log from being used to recover from a media failure. You must do a dump database now to ensure the recoverability of subsequent database modifications using transaction logs.

Note This database dump is not required if your backup and recovery procedures do not include saving the transaction logs for media failure recovery.

If you are concerned that your transaction log is not large enough for your Adaptive Server, refer to information about estimating the transaction log size within "Creating and Managing User Databases" in the *System Administration Guide* before deciding to increase its size. If you choose to increase the size, refer to the *Adaptive Server Reference Manual* for more information about the alter database command used to increase log size. Remember that once they have been increased, log and database sizes cannot be decreased.

Extending the Transaction Log Segment

Warning! If you use this option to extend the transaction log segment, you will not be able to get that space back later.

1 Refer to "Estimating the Transaction Log Size" within "Creating and Managing User Databases" in the *System Administration Guide* before deciding whether to increase the transaction log size.

2 To get information about the log segment, enter a command like the following and replace "test" with the name of the database in the error message:

```
1> use test
2> go
1> sp_helpsegment logsegment
2> go
segment  name                           status
-------  -----------------------------  ------
      2  logsegment                          0
device                          size                    free_pages
------------------------------ ---------------------- ----------
dev1                            1.0MB                          512
table_name                  index_name                     indid
--------------------------- ------------------------------ ------
syslogs                     syslogs                             0
```

3 If you have more space in the current log device fragment, increase the amount of space allocated for the transaction log with commands like the following:

```
1> use master
2> go
1> alter database database_name
2> log on device_name = additional_space
3> go
```

where:

- *database_name* is the name of the database for which the log device fragment is to be extended.
- *device_name* is the name of the database device on which to locate the database extension.
- *additional_space* is the number of megabytes to extend the log device fragment.

4 If you do *not* have any more space in the current log segment, you can use sp_extendsegment to extend the range of the log segment to another database device (the database must already exist on that device fragment before you run these commands):

```
1> use database_name
2> go
```

```
1> sp_extendsegment logsegment, database_name,
2> device_name
3> go
```

where:

- *database_name* is the name of the database where the segment is to be extended.
- *device_name* is the name of the additional database device for the log segment.

Warning! If you use sp_extendsegment to extend the range of the log segment to another database device and the other device contains data, the segment on which that data exists will become a log segment. This situation can cause you to run out of space in the database earlier than you expected and may cause other problems, including 2558 errors.

Version in which this error is raised	All versions

Error 4205

Severity	16
Message text	`Syslogs does not exist in its own segment in database '%S_DBID' with segmap '%ld' with logical start page number of '%ld'. You cannot use DUMP TRANSACTION in this case, use DUMP DATABASE instead.`
Explanation	The dump transaction command copies the transaction log, which resides in the syslogs table, recording any database changes made since the last database dump or transaction log dump. Error 4205 is raised when you attempt dump transaction, but the log and data share the same device. Transaction log dumps are only possible when data and log reside on separate devices. Since transaction logs are used for up-to-date recovery in the event of a disk failure, they are useful only if kept separate from the data device.
Action	Use dump database. This will make a copy of the entire database, including both data and the transaction log.
Additional information	Unless you are creating small, noncritical databases, always place the log on a separate device. This allows you to dump transaction logs, ensures full recovery from hard disk crashes and has many other benefits including performance improvement. For details, refer to Placing the Transaction Log on a Separate Device in the *System Administration Guide.* When data and log share device space, it is possible to move the transaction log to its own device by using the system stored procedure sp_logdevice. For related information, see the writeup for Error 4222 in this manual.
Version in which this error is raised	All versions

Error 4207

Severity	`16 State 2`
Message text	`Dump transaction is not allowed because a non-logged operation was performed on the database. Dump your database or use dump transaction with truncate_only`

```
until you can dump your database.
```

Explanation

This error occurs when you have performed a minimally logged operation such as bcp in or select into on the database, followed by a dump transaction to a dump device. The dump transaction is not allowed because your database contains data that cannot be recovered from the log. The log could not be used to recover the database properly. For instance, if a user does a select into followed by an update, data for the select into is not recorded in the log, so the update cannot be recovered from the log.

Note A minimally logged operation is one in which changes to the data itself are not logged to the transaction log. All changes to the allocation structures resulting from the operation are logged to the transaction log.

Note that simply enabling the select into / bulkcopy option does not, in itself, cause this condition. The option must be set and a minimally logged operation executed before this error message is reported during a dump transaction to a dump device.

Dumping the database creates a new save point on which subsequent transaction logs are based. Therefore, dumping the database ensures that the minimally logged changes are recoverable and dumping the transaction log to a device is once again allowed.

Action

Perform a dump database. Disabling the select into/bulkcopy option using **sp_dboption** will not allow you to dump the log. You must dump the database to clear the 4207 error condition. You may want to disable the select into/bulkcopy option to prevent any further nonlogged operations in your database.

You may need to use dump transaction with truncate_only until you can perform a full database dump. If you must do this, you will lose the ability to recover up-to-the-minute changes in data in case of a media failure.

Version in which this error is raised

All versions

Error 4216

Severity

```
21
```

Message text

```
%s failed in database %.*s:  unexpected end of log
```

```
encountered on page %ld while scanning for page %ld.
```

Note This error may be caused by a hardware problem.

Explanation

The transaction log consists of a series of log pages chained together in a doubly linked list; each log page has a page header with pointers to the next and previous log pages. The previous page pointer for the first page is zero, as is the next page pointer for the last page, indicating the beginning and end of the log chain respectively.

Error 4216 is raised when Adaptive Server finds what it considers to be a premature end of the transaction log chain during execution of a command. This error is raised during dump transaction or dump database and is often seen in conjunction with Error 4219.

Error 4216 can be caused by a bad page chain or allocation (syslogs corruption).

Action

1 Determine the status of syslogs:

```
1> use database_name
2> go

1> dbcc checktable (syslogs)
2> go

1> dbcc tablealloc (syslogs)
2> go
```

where *database_name* is the name of the database appearing in the error message.

dbcc reports any problems found in syslogs, and the number of data pages being used by the log. If your log is on a separate device, dbcc checktable also tells you how much space is used and how much is free.

2 If no problems were reported, the transaction log is clean. Call Sybase Technical Support for assistance.

3 If problems were reported, try to fix them using the sections in this manual that describe those errors. If errors are still reported, call Sybase Technical Support.

Additional information

For information about estimating and managing the transaction log size, refer to “Creating and Managing User Databases” in the *System Administration Guide*.

Before calling Technical Support, have the following information available:

- Server version and SWR Rollup level
- Server error log
- Output of dbcc checktable and dbcc tablealloc
- Text of all error messages

Version in which this error is raised

11.0 and higher

Error 4221

Severity

16

Message text

```
DUMP TRANsaction to a dump device is not allowed where
a truncate-only transaction dump has been performed
after the last DUMP DATABASE.  Use DUMP DATABASE
instead.
```

Explanation

The dump database command makes a copy of the entire database, including both the data and the transaction log. dump database does not truncate the log.

The dump transaction command copies the transaction log, recording any database changes made since the last database or transaction log dump. After copying the log, it truncates the inactive portion. The dump transaction with truncate_only command, on the other hand, removes transactions from the log without making a backup copy.

Error 4221 occurs in the following situations:

- If you attempt a dump transaction command, and the most recent dump consisted of a dump transaction with truncate_only. This is not allowed because the dump transaction would produce a transaction dump that could not be applied, since the prior dump transaction with truncate_only removed transactions upon which the current transactions may depend. To ensure recoverability, you must dump the database with dump database each time it is necessary for you to run a dump transaction with truncate_only command.
- If you use the log on clause to create database to place the transaction log on a separate database device, and you attempt a dump transaction prior to dumping the database.

Action

Perform a dump database first to back up the database and log. Follow this up with the appropriate dump transaction command.

Make sure the trunc log on chkpt database option is *off*. When this option is on you cannot dump the transaction log.

Version in which this error is raised

All versions

Error 4222

Severity

16

Message text

```
DUMP TRANSACTION for database '%S_DBID' failed: log page
(%ld) is on a data-only segment. This is probably due
to an incomplete move of the log to a device separate
from the data. Complete the log migration or contact the
System Administrator.
```

Explanation

The dump transaction command copies the transaction log, recording any database changes made since the last database dump or transaction log dump. After copying, it truncates the inactive portion of the log. Transaction log dumps are only possible when data and log reside on separate devices.

If data and log share device space, it is possible to move the transaction log to its own device by using the system stored procedure sp_logdevice.

Error 4222 is raised when you execute dump transaction, but the first page of the transaction log is found to be on a data-only segment of the database. Since the log is not isolated from the data, the transaction dump fails.

The likely cause of this error is that the steps to move the log to a separate device were initiated, but were not complete at the time you issued dump transaction. To understand this, note that sp_logdevice moves future allocation for the transaction log to the new device. However the log remains on the original device until the extent that is currently in use has been filled and the transaction log has been dumped. Error 4222 indicates that the log in its current state cannot be migrated to its own device.

Action

Take the following steps to resolve the error:

1 Run dbcc checktable (syslogs) to determine which page is currently in use.

2 Execute enough transactions to fill the extent currently in use. The following commands will fill at least one extent on the log (8 pages):

```
1> create table dummy1 (c1 char(255), c2 char(255))
2> go
1> insert dummy1 values ("a", "a")
2> go 16
```

```
1> drop table dummy1
2> go
```

3 Make sure that there are no currently active transactions on the database device, then run dump transaction with truncate_only.

> **Note** If an active transaction causes a begin tranto be written onto the log pages in question, a transaction dump will not be able to truncate those pages. This is why you should avoid having activity on the device while you migrate the log off the data segment.

4 Run the system procedure sp_helplog to ensure that the complete log is on the new device.

Additional information

Refer to the *System Administration Guide* for details on moving the transaction log to another device.

Version in which this error is raised

All versions

Error 4305

Severity

16

Message text

```
Specified file '%.*s' is out of sequence. Current time
stamp is %S_DATE while dump was from %S_DATE.
```

Explanation

During a load transaction, Adaptive Server checks the timestamp on the dumped database and each dumped transaction log to determine whether the sequence is correct. If the load sequence of the transaction logs is incorrect, Error 4305 occurs and the load fails.

If you are missing a transaction log dump, you can only restore up to the transaction log dump before the missing dump.

In addition, you cannot load multiple transaction log dumps that have been created with the no_truncate option. The no_truncate option assumes that nothing further will be done with the transaction log after that point: it is only useful for up-to-the-minute recovery of a database after failure of the data device. If you have used the no_truncate option and your database is still okay, do a dump database as soon as possible.

This error is serious as the load processing stops and the load transaction fails.

Action	Load the log backups in the same sequence in which they were made. Check your transaction log dumps to determine whether one was loaded out of order. If you find the correct dump, load it and continue with the rest of the dumps, in order.
Additional information	For more information, refer to “load transaction” in the *Reference Manual*.
Version in which this error is raised	All versions

Error 4306

Severity	16
Message text	`There was activity on database since last load, unable to load. Must restart load sequence with the load database to continue.`
Explanation	If user transactions modify the transaction log between transaction dump loads to a database, Error 4306 occurs and the load fails. In 11.0.x and later versions of the server, Error 4306 would only occur if someone ran the online database command for the database, someone did work in the database, and then you tried to do a load transaction for the database.
Action	Begin again with a load database or proceed to use the database without recovering all transactions.
Additional information	For more information, refer to “load transaction” in the *Reference Manual*.
Version in which this error is raised	All versions

Error 4322

Severity	16
Message text	`SQL Server cannot load a dump of the master database that was dumped from an older version of SQL Server. The current master database upgrade version is %ld, the dump master database upgrade version is %ld.`
Explanation	Error 4322 is raised when you try to load a dump of the master database that was created on an earlier version of Adaptive Server than your currently installed version.

Action

If you know you do not have a dump of master from your current version level because you did not dump master after you upgraded, and your master database is corrupted, call Sybase Technical Support.

Otherwise, check the current upgrade version of the master dump you are trying to load. For example:

```
1> load database master from 'master.dump' with headeronly
2> go
Backup Server session id is:  6.  Use this value when executing the
`sp_volchanged' system stored procedure after fulfilling any volume
change request from the Backup Server.
Backup Server: 6.28.1.1: Dumpfile name `master_952820A2F8' section
number 0001 mounted on disk file
`/remote/solaris/rel1100/install/master.dump'
This is a database dump of database ID 6, name `master', from
Oct  9 1995 11:35AM.  SQL Server version:
SQL Server/11.0/B/Sun_svr4/OS5.2/1/OPT/Fri Aug 1805:10:26 PDT 1995.
Backup Server version:
Backup Server/11.0/B/Sun_svr4/OS5.4/1/OPT/Thu Aug 17 21:54:21 PDT 1995.
Database contains 1536 pages; checkpoint RID=(Rid pageid = 0x405;
row num = 0xd); next object ID=3031; sort order ID=50, status=0;
charset ID=1.
Database log version=2; database upgrade version=1.
```

The *log version* and *upgrade version* have values as follows:

Version	Value Is 0	Value Is 1	Value Is 2
Database log version	10.0.x or earlier	10.1	11.0 or later
Database upgrade version	10.x or earlier	11.0 or later	Not used

Additional information

If you need to call Sybase Technical Support, have the following information ready:

- Server version and SWR Rollup level
- Server error log
- Text of all error messages
- Output from load database...with headeronly command if you have what you thought was a current dump
- Output from sp_configure "upgrade version"

Version in which this error is raised

All versions

truncate table Errors

This section describes errors for the truncate table command.

Error 4716

Severity

```
16
```

Message text

```
Cannot truncate table '%.*s' because it is partitioned.
```

Explanation

By default, Adaptive Server stores a heap table's data in one doubly linked chain of database pages. Adaptive Server inserts all new rows into the last page of the chain. A transaction holds an exclusive lock on the last page while inserting new rows, which can block other, concurrent transactions from being inserted into the table.

The partition clause of the alter table command allows you to partition user tables that do not have a clustered index. Partitioning creates additional page chains on the table, each with its own last page. This reduces page contention for concurrent inserts, and can also reduce I/O contention if the table is distributed over multiple physical devices.

Error 4716 is raised when you try to truncate a table that is partitioned.

Action

If you want to truncate the table, use the unpartition clause of the alter table command to concatenate all partitions:

```
1> alter table table_name unpartition
2> go
```

Then truncate the table.

Additional information

Refer to the *Reference Manual* for information about the alter table command.

Version in which this error is raised

All versions

Bulk Copy Utility Errors

This section describes errors for the Bulk Copy Utility (bcp).

Error 4801

Severity

20

Message text

```
Bulk_main: opentable on BULK INSERT table failed.
Dbid=%d name='%.*s'.
```

Explanation

Error 4801 is raised when you attempt to bulk copy data into a table, but bcp is unable to locate the target table.

Possible causes include:

- a bcp syntax error
- more than one session attempting to bcp into the same table
- the target table was deleted or renamed by a concurrent user

Action

1. Check the syntax of your bcp in command and the target table named in the command.
2. Check that the target table exists, and that no other user is attempting to bulk copy to the table.
3. If you are making a substantial change, such as loading a large number of tables, you may need to increase the number of open objects configuration parameter. See the *System Administration Guide* for details.

Version in which this error is raised

All versions

Error 4806

Severity

16

Message text

```
You cannot run the non-logged version of bulk copy in
this database. Please check with the DBO.
```

Explanation

This error occurs when the select/into bulkcopy option is set to false and you use "fast" bulk copy into a table that has no indexes or triggers.

Note By default, the select into/bulkcopy option in newly created databases is set to the same as that in model.

Action

You have the following options for recovering from this error.

Note Not all options are available in all versions.

Turn On *select into/bulkcopy (all versions)*

Use the sp_dboption stored procedure to set the select into/ bulkcopy option to true:

```
1> sp_dboption database_name,
2> "select into/bulkcopy", true
3> go
```

```
1> use database_name
2> go
```

```
1> checkpoint
2> go
```

Once a minimally logged operation such as "fast" bulk copy runs in the database, you are not allowed to dump the transaction log to a device, because unlogged changes are not recoverable. Instead, you should do a dump database as soon as possible to restore recoverability and allow transaction dumps to devices again.

Create an Index or Trigger for the Table (Version 11.5.x and earlier)

Put an index or trigger on the table you are trying to copy into. This causes inserts to be logged. You can create an insert trigger that does not perform any actions.

When you copy into a table that has indexes or triggers, a slower version of bcp is automatically used. The slow version, which does log data inserts in the transaction log, can cause the transaction log to become very large. You may need to use dump transaction with truncate_only until you can perform a full database dump. If you must do this, you will lose the ability to recover up-to-the-minute changes in data in case of a media failure. You may also consider using smaller batch sizes.

Create an Index for the Table (Version 11.9.2 and later)

Put an index on the table you are trying to copy into. This causes inserts to be logged.

When you copy into a table that has indexes, a slower version of bcp is automatically used. The slow version, which does log data inserts in the transaction log, can cause the transaction log to become very large. You may need to use dump transaction with truncate_only until you can perform a full database dump. If you must do this, you will lose the ability to recover up-to-the-minute changes in data in case of a media failure. You may also consider using smaller bcp batch sizes.

Version in which this error is raised

All versions

alter table Errors

This section contains error messages for the alter table command.

Error 4950

Severity: 16

Message text:

```
Cannot partition table '%.*s' because it is a system
table.
```

Explanation:

By default, Adaptive Server stores a heap table's data in one doubly linked chain of database pages. Adaptive Server inserts all new rows into the last page of the chain. A transaction holds an exclusive lock on the last page while inserting new rows, which can block other, concurrent transactions from being inserted into the table.

Partitioning creates additional page chains on the table, each with its own last page. This reduces page contention for concurrent inserts, and can also reduce I/O contention if the table is distributed over multiple physical devices.

The partition clause of the alter table command allows you to partition user tables that do not have a clustered index. Partitioning is not allowed for system tables.

Error 4950 occurs when you try to partition a system table.

Action: No action is required.

Additional information: Refer to the *Reference Manual* for information about the alter table command.

Version in which this error is raised: All versions

Error 4951

Severity: 16

Message text:

Version 11.5

```
Cannot create %d partitions on the table '%.*s' with
clustered index because the maximum number of partitions
allowed on a table with clustered index is %d.
```

Version 11.0.x and Earlier

```
Cannot partition table '%.*s' because it has a clustered
index.
```

Explanation

By default, Adaptive Server stores a heap table's data in one doubly linked chain of database pages. Adaptive Server inserts all new rows into the last page of the chain. A transaction holds an exclusive lock on the last page while inserting new rows, which can block other, concurrent transactions from being inserted into the table.

Partitioning creates additional page chains on the table, each with its own last page. This reduces page contention for concurrent inserts, and can also reduce I/O contention if the table is distributed over multiple physical devices.

Version 11.5 and Later

The partition clause of the alter table command allows you to partition tables with or without a clustered index. However, after partitioning a clustered table, the entire page chain is placed in the first partition.

Error 4951 occurs when you attempt to partition a table with a clustered index, and exceed the maximum allowed partitions.

Version 11.0.x and Earlier

The partition clause of the alter table command allows you to partition user tables that do not have a clustered index. A clustered index requires that all leaf pages be part of one page chain. Since partitioning splits the table into separate page chains, one chain per partition, partitioning is not allowed for tables with a clustered index.

Error 4951 occurs when you try to partition a user table with a clustered index.

Action

Version 11.5

Reenter your partition command using a value for number of partitions that does not exceed the specified maximum.

Version 11.0.x and Earlier

If you want to partition the table, drop the clustered index and then reenter your partition command:

1 Identify the clustered index:

```
1> use database_name
2> go
```

```
1> sp_helpindex table_name
2> go
```

2 Drop the clustered index:

```
1> drop index table_name.index_name
2> go
```

3 Partition the table.

Additional information Refer to the *Reference Manual* for information about the drop index and alter table commands.

Version in which this error is raised All versions

Error 4953

Severity 16

Message text

```
Cannot partition table '%.*s' because it is already
partitioned.
```

Explanation The partition clause of the alter table command allows you to partition user tables that do not have a clustered index. You cannot partition a table that is already partitioned.

Error 4953 occurs when you try to partition a table that is already partitioned.

Action If you want to change the number of partitions for the table, follow these steps:

1 Concatenate all existing page chains:

```
1> use database_name
2> go
```

```
1> alter table table_name unpartition
2> go
```

2 Repartition the table:

```
1> alter table table_name partition
2> number_of_partitions
3> go
```

Additional information Refer to the *Reference Manual* for information about the alter table command.

Version in which this error is raised All versions

Error 4954

Severity

16

Message text

```
Cannot unpartition table '%.*s' because it is not
partitioned.
```

Explanation

The partition clause of the alter table command allows you to partition user tables that do not have a clustered index. The unpartition clause of the alter table command allows you to concatenate all existing page chains for these tables. You cannot unpartition a table that is not partitioned.

Error 4954 occurs when you try to unpartition a table that is not partitioned.

Action

No action is required.

If you want to determine whether a table is partitioned, use the following commands:

```
1> use database_name
2> go
```

```
1> sp_help table_name
2> go
```

Additional information

Refer to the *Reference Manual* for information about the alter table command.

Version in which this error is raised

All versions

Error 4956

Severity

16

Message text

```
Cannot alter table if level 0 scans are active.
```

Explanation

Isolation level 0 allows transactions to read uncommitted data (“dirty reads”). You can specify isolation level 0 for your queries using the at isolation syntax or you can specify level 0 for your session as part of the transaction isolation level option of the set command. Adaptive Server's default transaction isolation level is 1.

Since alter table allows you to modify a table's schema, it would not be safe to allow the use of alter table when level 0 scans are active for that table.
Error 4956 is raised if you enter an alter table command when the table in your command is currently being used for reading uncommitted data.

Action

Try your alter table command again later, after the process that is using level 0 scans on the table is no longer active.

Additional information

Refer to the *Adaptive Server Enterprise Transact-SQL User's Guide* for information about selecting an isolation mode for either a query or for your session.

Version in which this error is raised

All versions

Error 4964

Severity

17

Message text

```
Lock scheme conversion failed due to insufficient SQL
Server memory. Please retry later when there is less
load/users on the SQL server, or ask your System
Administrator to reconfigure SQL Server with more
memory.
```

Explanation

Adaptive Server 11.9 provides two new locking schemes:

- Datapages locking, which locks only the data pages
- Datarows locking, which locks only the data rows.

Since neither scheme locks index pages, they are referred to together as the data-only locking scheme.

The pre-11.9 locking scheme, known as allpages locking, continues to be available with 11.9. This scheme locks the data and index pages affected by queries, and is the default locking scheme.

When you use the alter table command to change a table's locking scheme from allpages locking to data-only locking, Adaptive Server creates memory structures to build a new table with the desired locking scheme and transfers data from the existing table.

Error 4964 indicates that there is not enough memory available in Adaptive Server's data cache to effect the lock scheme change. The problem can occur at various stages of table conversion:

- when initializing a new Object Allocation Map (OAM) page
- when initializing a new index structure
- when building the new table.

Action

Check the activity on your server and re-try the lock scheme conversion when there is less load on the system.

If the problem persists, check the memory allocation on your server. You may need to increase the memory allocated to the data cache.

Use sp_configure to view the current values of memory-related parameters on your system:

```
1> sp_configure "Memory Use"
2> go
```

Refer to "Configuring Memory" in the *System Administration Guide* for details on monitoring and adjusting memory allocation parameters.

Version in which this error is raised

11.9 and later

Error 4981

Severity

16

Message text

```
You cannot drop a column from table '%.*s', as this
table does not use the allpages lock scheme.
```

Explanation

As of Version 11.9, Adaptive Server provides two new locking schemes:

- Datapages locking, which locks only the data pages
- Datarows locking, which locks only the data rows.

They are referred to together as the data-only locking scheme.

The pre-11.9 locking scheme, known as allpages locking, continues to be supported. This scheme locks the data and index pages affected by queries, and is the default locking scheme.

Error 4981 is raised in version 11.9.x when you attempt to drop a column in a data-only locked table.

Warning! In version 11.9.x and earlier, the command to directly drop a column is undocumented, and is not supported by Sybase. Do not attempt to drop a table column using an unsupported command, regardless of the locking scheme. Failure to heed this warning may lead to recovery problems and database corruption with 605 and 821 errors.

Action

In version 11.9.x and earlier, use one of the following to remove a column from the table (these methods apply to both data-only locked and allpages locked tables):

1 Select the whole table into a new table, excepting the column you wish to drop. Drop the old table and rename the new table.

2 Bulk copy the affected table out, excepting the column you wish to drop; then drop and re-create the table, and bulk copy back in. This is the most efficient solution for a very large table.

> **Note** You will need a mechanism – such as a format file – to exclude a column when bulk copying out the table.

Additional information

For more information about how to copy a table in a new table, refer to Chapter 2, "How to Rescue Data from a Corrupted Table".

Version in which this error is raised

11.9.x

alter database Errors

This section contains error messages for the alter database command.

Error 5006

Severity

16

Message text

```
Could not find enough space on disks to extend database
%.*s.
```

Explanation

This error occurs when an alter database command refers to an Adaptive Server device that is completely full.

If the device is not full but has less space than the alter database command requests, the database will be extended to fill up the available space on the device and no error message will be displayed. Instead, you will see the usual status message from alter database, but it will report fewer pages than you requested.

This error message does not usually appear in the error log. However, if the exhausted device also uses up the available physical disk drive space as well as the space allocated for Adaptive Server, kernel messages declaring the file system full will occur in the error log. If this happens, fix the file system problem as well as the lack of space for Adaptive Server.

Action

To recover from Error 5006, do one of the following steps:

- Use a device name in the alter database command which is large enough. Use sp_helpdevice and sp_helpdb to find out which devices have enough room.
- Add a new device to Adaptive Server with the disk init command. Then refer to this device with the alter database command.
- Although there is no particular performance benefit to keeping a database entirely on a single device, if you need to do this at your site, use the following procedure:

  a Back up all the databases on the device by dumping them to a valid dump device with the dump database command.

  b Drop the device with sp_dropdevice.

  c Initialize a larger device with disk init.

d Create the old databases on the new device using standard procedure. It is important to create the databases exactly like the old ones or errors will occur. Refer to Error 2558 for more information.

e Reload each database with the load database command.

f Use the online database command for each database to make the databases available for use.

Refer to "Developing a Backup and Recovery Plan" in the *System Administration Guide* for information about how to safely create, dump, load and re-create databases.

Version in which this error is raised

All versions

Error 5013

Severity

16

Message text

```
Cannot extend the MASTER database onto any device other
than 'master'. The ALTER DATABASE was aborted.
```

Explanation

This error occurs when you try to extend the master database onto a device other than the master device.

Note It is recommended that you keep user objects out of the master database. If you keep user databases off the master device, you allow space in case the master database needs to grow. In addition, if you ever need to rebuild the master device, it will be easier if it does not contain user databases.

Action

Adaptive Server users can move any "home-grown" system procedures that start with "sp_" to sybsystemprocs (by dropping them from the master database and creating them in sybsystemprocs).

Extend the master database only if absolutely necessary! If you are sure you must increase the master database size and have no room on the current master device, use the following procedure to remove user databases from the master device.

Move User Databases

1 Dump the user databases with the dump database command.

2 Rename the dumped databases on the master device with sp_renamedb.

3 Re-create the databases with their original names on another device with create database. Be sure they are created exactly like the old databases, to avoid 2558 and other errors. Refer to Error 2558 for more information.

4 Load the dumps with load database.

5 Use the online database command for each database to make the databases available for use.

6 Check the databases in their new location to make sure the load was successful (that is, perform a simple query with isql), and if everything loaded successfully, drop the old databases from the master device.

You can now try to increase the size of the master database on the master device with the alter database command.

Increase Master Device Size

If the master device contains only the master database and the master device is too small, then use the following procedure:

Warning! Altering the master device is extremely risky! Avoid it if at all possible. Be familiar with the recovery methods in Chapter 1, "System Database Recovery," in case you lose your master database or master device.

---

1 Back up the master database with the dump database command.

2 Save the contents of key system tables such as sysdatabases, sysdevices, sysusages, and syslogins.

3 Use the buildmaster utility to build a new master device with enough extra space so that you will never need to increase the master device again. When buildmaster completes, a new master database will exist on the new master device.

4 Expand the size of the new master database with the alter database command, if necessary, so that it matches the size of the dumped master database.

5 Execute the following command in isql:

```
1> select name, high from master..sysdevices
2> where name = "master"
3> go
```

and note the "high" value for the master device.

6 Start Adaptive Server in single-user mode. Refer to Chapter 2, "How to Start Adaptive Server in Single-User Mode" for instructions.

7 Load the dump of the master database. Test the database to make sure it is functioning correctly before deleting any extra backups of master.

8 Allow updates to the system catalog:

```
1> sp_configure, "allow updates", 1
2> go
```

9 Reset the "high" value in master..sysdevices:

```
1> begin transaction
2> go
1> update master..sysdevices
2> set high = <value_of_high_from_step_5>
3> where name = "master"
4> go
```

10 If the previous update affected only one row, commit the transaction.

11 Restart Adaptive Server.

12 Turn off allow updates:

```
1> sp_configure, "allow updates", 0
2> go
```

Version in which this error is raised: All versions

Error 5018

Severity: 17

Message text:

```
Caution:  You have set up this database to include space
on disk %d for both data and the transaction log.  This
can make recovery impossible if that disk fails.
```

Explanation: Error 5018 is raised when you execute alter database to add storage space, and the disk map for the database is found to contain both data and log segments.

The message can also be raised in another situation. When you load a database from dump, Adaptive Server compares the segments on the dump with the device allocations in the target database. If they match, the server maintains the allocations. If, however, there is a mismatch, the load will remap segments to the free device allocations. The remapping may result in mixed data and log segments, raising the 5018 error.

Why is recovery difficult when data and log are mixed? To understand this, consider first what happens when you keep data and log separate, and make a complete database backup (say) every 7 days. If the device containing the data crashes, you can restore the database from the latest backup, then apply the committed transactions from the log device, restoring the database to practically the last moments before failure.

In contrast, if data and log are mixed and the device fails, you can only restore the database up to the time of the last full backup, since there is no way to apply the transaction log.

Action

If Error was Raised when Creating a Database

Unless you are creating small, noncritical databases, always place the log on a separate device. This allows you to dump transaction logs, ensures full recovery from hard disk crashes and has many other benefits including performance improvement. For details, refer to “Placing the Transaction Log on a Separate Device” in the *System Administration Guide.*

If Error was Raised when Loading a Database

It is likely that there is a mismatch in the segment mappings between the dump’s source and target databases. To correct this problem:

1 In the source server (where the dump is to be executed), run the following query:

```
1> use master
2> go
1> select * from sysusages where dbid=db_id("<database_name>")
2> go
```

Each line of output corresponds to a segment, with a *segmap* of 4 denoting a log segment, and any other segmap denoting a data segment. Save the output.

2 Dump the database.

3 Create the target database for load, specifying the segments in the same order and with the same sizes as defined by the query output in Step 1. For simplicity, make sure that all non-log segments have a segmap value of 3; if you had user-defined segments, the load will automatically remap these segments.

   Check that the segment mappings are the same, by running the following query for the new database:

```
1> select * from sysusages where dbid=db_id("<database_name>")
2> go
```

4 Load the database.

Version in which this error is raised — All versions

Error 5034

Severity — 16

Message text

```
ALTER DATABASE was successful. Cannot modify the last-chance threshold because the database '%.*s' is in read-only mode.
```

Explanation

Each database that stores its transaction log on a separate segment has a last-chance threshold (LCT). This is an estimate of the number of free log pages that would be required to back up the transaction log.

Error 5034 is raised when you use alter database to extend log space on a read-only database. The database is altered and log space added, but the LCT cannot be modified because this requires an update to systhresholds, which is not possible in a read-only database.

Error 5034 is followed immediately by Error 5035 ("Remember to execute SELECT LCT_ADMIN('LASTCHANCE', %d), when the database is made writable.").

Note Version 11.0.3.2 and earlier raise Error 3906 ("Attempt to BEGIN TRANsaction in database '%.*s' failed because database is READ ONLY.") when alter database is used to extend log space. The database is altered and log space added, but the LCT cannot be modified. See the writeup of Error 3906 for details.

Action

Re-establish the LCT after the database is available for updates using the following steps:

1 Use the following command to determine the new last-chance threshold:

```
1> use <database name>
2> go
1> select lct_admin("reserve", <log size in pages>)
2> go
```

For example, on a 20MB transaction log (containing 10240 2KB pages):

```
1> select lct_admin("reserve", 10240)
2> go
-----------
616
```

The resulting value is the amount of space required to dump the transaction log.

2 Modify the LCT using

```
1> select lct_admin('lastchance', <new threshold>)
2> go
```

Additional information

Refer to Creating a Free-Space Threshold for the Log Segment in the *System Administration Guide* for additional information on setting thresholds.

For databases normally operated in read-only mode, no updates occur on the database and so there is no need to update the LCT.

Version in which this error is raised

11.0.3.3 and later

Disk Errors

This section contains error messages pertaining to disk activation and system administrator disk commands.

Error 5115

Severity

16

Message text

```
I/O error during disk initialization.  PHYSNAME '%.*s'
may be incorrect or %ld (VSTART + SIZE) may exceed the
size of the device.  Please consult the SQL Server error
log for more details.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when a disk init command fails. It can occur for the following reasons:

- Incorrect physical device name
- Incorrect size
- Virtual device number not unique
- Permission problem
- UNIX kernel not configured for asynchronous I/O.

Action

1 Check the Adaptive Server error log for information that might help determine the cause of the error.

2 If necessary, check the items mentioned in the sections below.

3 Make any needed changes and then run the command again.

Incorrect Device Name

Make sure you have the correct physical device name. Refer to the Adaptive Server installation and configuration guide or the *System Administration Guide* for instructions.

Incorrect Size

The size parameter of the disk init command must be specified in units of 2K blocks (size of most data pages). There are 512 2K blocks in 1MB. Use the formula:

```
#_of_Megabytes * #_of_data_pages_per_Megabyte =
#_of_2K_blocks
```

For example, to initialize a 10MB Sybase device, specify 5120 for the size:

```
10MB * 512 = 5120
```

To determine the number of 2K blocks to use for a raw partition on a UNIX system, convert the size of the partition, in sectors, to 2K blocks. The size of a sector varies on different UNIX systems. For example, it is 512 bytes on Sun, so use sectors/4 for the size parameter. It is 1024 bytes on HP, so use sectors/2. Also, see the information below about virtual device numbers. You must restart Adaptive Server or use a different virtual device number when you reissue the disk init command.

Refer to the Adaptive Server installation and configuration guide for more information about Sybase and your operating system's partitions.

Note Not all operating systems have partitions as described above.

Virtual Device Number

The vdevno parameter of the disk init command must be unique. Also, if a disk init command fails, Adaptive Server saves the vdevno used, even though the disk init command failed. To make the vdevno available again, restart Adaptive Server. Use sp_helpdevice to determine which virtual device numbers have been used.

The available number of virtual devices is equal to the number of configured devices minus one. For example, if the number of devices configuration parameter is set to 10, then there are 9 available devices, 1 through 9.

Permission Problem

If Adaptive Server was started by the “sybase” account, it must have permission to write to the raw partition or to create a file in the directory specified by disk init. Raw partitions used by Sybase should always be owned by the “sybase” account. Correct the permissions problem and execute the command again, using a different vdevno. Refer to “Virtual Device Number” for more information.

UNIX Kernel

If the UNIX kernel (on SunOS 4.x/Sun Solaris 1.x systems only) is not configured for asynchronous I/O, the 5115 error occurs only if your master device is on a UNIX file, and you are trying to do a disk init on a raw partition. Have your UNIX administrator rebuild the kernel to allow for asynchronous I/O. Refer to Error 823 for more information.

Version in which this error is raised

All versions

Error 5123

Severity

16

Message text

```
DISK INIT encountered an error while attempting to
open/create the physical file. Please consult the SQL
Server error log (in the SQL Server boot directory) for
more details.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when a disk init command fails while trying to initialize an operating system file or a raw partition for use as a database device. It can occur for the following reasons:

- Incorrect physical device name
- Permission problem
- Incorrect size specified (for raw devices, when you request initialization of a device larger than the size of the raw partition)
- Not enough room on the device.

Action

Check each section below, make any needed changes and then run the command again.

Incorrect Device Name

Make sure you have the correct physical device name. Refer to the Adaptive Server installation and configuration guide or the *System Administration Guide* for instructions.

Permission Problem

If Adaptive Server was started by the “sybase” account, it must have permission to write to the raw partition or to create a file in the directory specified by disk init. Raw partitions used by Sybase should always be owned by the “sybase” account. Correct the permissions problem and execute the command again, using a different vdevno. Refer to “Virtual Device Number” for more information.

Note Adaptive Server should *not* be started by “root.” If it is started by root, you will get Error 5123 if the device or file is owned by “sybase.”

Incorrect Size

When you request initialization of a raw device, make sure the size you specify is not larger than the size of the raw partition.

The size parameter of the disk init command must be specified in units of 2K blocks (size of most data pages). There are 512 2K blocks in a megabyte. Use the formula:

```
#_of_Megabytes * #_of_data_pages_per_Megabyte =
#_of_2K_blocks
```

For example, to initialize a 10MB Sybase device, specify 5120 for the size:

```
10MB * 512 = 5120
```

To determine the number of 2K blocks to use for a raw partition on a UNIX system, convert the size of the partition, in sectors, to 2K blocks. The size of a sector varies on different UNIX systems. For example, it is 512 bytes on Sun, so use sectors/4 for the size parameter. It is 1024 bytes on HP, so use sectors/2. Also, see the information below about virtual device numbers. You must restart Adaptive Server or use a different virtual device number when you reissue the disk init command.

Refer to the Adaptive Server installation and configuration guide for more information about Sybase and your operating system's partitions. For example, on UNIX, there is a size limitation of 2GB for a device file.

Note Not all operating systems have partitions as described above.

Additional information

Refer to “Initializing Database Devices” in the *System Administration Guide* for information about disk init.

Version in which this error is raised: All versions

Error 5142

Severity: 16

Message text:

```
Mirroring for device '%.*s' is not currently enabled.
```

Explanation: Error 5142 is raised when you execute a disk unmirror command in retain mode, but the device in question is not currently mirrored.

Disk mirroring is automatically deactivated when a read or write to a mirrored device fails. You may see Error 5142, possibly accompanied by additional messages at the console or in the error log:

- if there are transient device failures
- if the mirror was set up improperly, for example if an incorrect device name was specified.

Action: Take the following steps:

1. Look at the server error log. Check for additional errors preceding the 5142 error to determine the probable cause. For example, a udunmirror kernel error indicates a physical disk problem leading to I/O failure; for corrective action, refer to the writeup for udunmirror Errors in this manual.
2. If the device was recently mirrored or remirrored, check the procedure used to set up the mirror. Since a database device and its mirror constitute one logical device, you should not initialize the mirror device with disk init. Also check the syntax used for the disk mirror command, ensuring that the path to the mirror is properly specified. Finally, check disk configuration and make sure the disk is properly labeled.

Additional information: To retain use of asynchronous I/O, always mirror devices that are capable of asynchronous I/O to other devices capable of asynchronous I/O. In most cases, this means mirroring raw devices to raw devices and operating system files to operating system files.

Refer to “Initializing Mirrors” in the *System Administration Guide* for details.

Version in which this error is raised: All versions

ASTC Errors

This section contains writeups for common Adaptive Server Transaction Coordinator (ASTC) errors.

Error 5602

Severity

17

Message text

```
Could not find an available distributed transaction
participant descriptor. Retry later, or please contact
a user with System Administrator (SA) role to
reconfigure your system for more descriptors.
```

Explanation

A distributed transaction (DTX) participant is an internal memory structure that the Adaptive Server Transaction Coordination (ASTC) service uses to manage a remote transaction branch. As transactions propogate to remote servers, the ASTC service must obtain new DTX participants (also known as DTX descriptors) to manage those branches.

Error 5602 is raised when there are not enough participants in the DTX descriptor pool to start or continue a remote distributed transaction. In-progress distributed transactions may abort due to a shortage of DTX participants.

Action

Update the number of dtx participants configuration parameter. To strike a balance between too many participants (wasting memory) and too few (failing transactions), use sp_monitorconfig during a peak period:

```
1> sp_monitorconfig "number of dtx participants"
2> go
```

If the *#Free* value is zero or very low, consider increasing the value to increase the number of remote transaction branches that Adaptive Server can handle.

Version in which this error is raised

12.0 and later

Open Client Errors

This section contains Adaptive Server error messages for Open Client applications.

Error 5701

Severity

10

Message text

```
Changed database context to '%.*s'.
```

Explanation

Adaptive Server sends this informational message to a DB-Library or CT-Library client:

- When the client logs in (the message displays the default database)
- When the client changes the current database

The message is displayed unless the client application suppresses the message via its message handler.

Action

If you wish to suppress this message, add the following line to your client's Adaptive Server message handler:

```
if (msgno==5701) return (0);
```

This message cannot be suppressed on the Server side; it must be handled on the client side.

Version in which this error is raised

All versions

Error 5702

Severity

10

Message text

```
The SQL Server is terminating this process.
```

Explanation

This informational message appears in conjunction with other errors such as 813, 8211, or stack traces. Error 5702 indicates that Adaptive Server is terminating the client process since the accompanying error was fatal.

Action

Check for and resolve the errors raised in conjunction with the 5702 error. Check the Adaptive Server error log if the other errors were not displayed on your screen.

Version in which this error is raised	All versions

Error 5703

Severity	10
Message text	`Changed language setting to '%.*s'.`
Explanation	Adaptive Server sends this informational message to a DB-Library or CT-Library client: • When the client logs in (the message displays the default language) • When the client changes the current language
Action	If you wish to suppress this message, add the following line to your client's Adaptive Server message handler: `if (msgno==5703) return (0);` This message cannot be suppressed on the Server side; it must be handled on the client side.
Version in which this error is raised	All versions

Error 5704

Severity	10
Message text	`Changed client character set setting to '%.*s'.`
Explanation	Adaptive Server sends this informational message to a DB-Library or CT-Library client: • When the client logs in (the message displays the default character set) • When the client changes the current character set
Action	If you wish to suppress this message, add the following line to your client's Adaptive Server message handler: `if (msgno==5704) return (0);` This message cannot be suppressed on the Server side; it must be handled on the client side.

Version in which this error is raised	All versions

Configuration Errors

This section contains error messages for Adaptive Server configuration.

Error 5808

Severity

10

Message text

```
WARNING: Dynamic loading of caches and pools through
loading a new file are not supported. However, the
loadfile '%s' will be inspected for consistency. Refer
to 'sp_cacheconfig' and 'sp_poolconfig' to create or
alter pools and caches.
```

Explanation

Configuration options are either dynamic or static. Dynamic parameters go into effect as soon as you execute sp_configure. Static parameters require Adaptive Server to reallocate memory; thus, they take effect only after Adaptive Server has been restarted.

Creating, dropping, and changing the size of data caches requires a restart of Adaptive Server for the configuration to take effect. Changes to data caches that take effect without a restart include changing the type, creating, dropping, and resizing memory pools with sp_poolconfig, changing the wash percentage of the pools, and binding and unbinding objects.

Error 5808, an informational message, is displayed while Adaptive Server is loading configuration information from a configuration file. Although dynamic loading of caches by using sp_configure to load a new configuration file is not supported, Adaptive Server checks cache configuration information as it is read from the file to make sure valid values are specified.

Action

No action is required.

To activate cache and pool entries already made in the configuration file, shut down and restart Adaptive Server.

In the future, you can use sp_cacheconfig and sp_poolconfig to configure caches and pools and avoid the 5808 message. You will still have to shut down and restart Adaptive Server to make cache changes take effect.

Additional information

Refer to the *Reference Manual* for information about sp_cacheconfig and sp_poolconfig.

Version in which this error is raised

All versions

Error 5824

Severity 16

Message text

```
Cannot reconfigure server to use sort order ID %d,
because the row for its underlying character set (ID %d)
does not exist in syscharsets.
```

Explanation

The syscharsets system table contains one row for each character set and sort order defined for use by Adaptive Server. One of the sort orders is marked in master..sysconfigures as the default sort order, which is the only one actually in use.

This error is raised when you are reconfiguring Adaptive Server to use a different character set and/or sort order, and one of the following is true:

- A corrupt index exists on syscharsets.
- The specified sort order or character set does not exist in syscharsets.

The error occurs with the following states:

State	Meaning
1	Adaptive Server could not find the specified character set in syscharsets
2	Adaptive Server could not find the specified sort order in syscharsets

Action

1 Determine the character sets that exist in syscharsets (note that a character set is defined by id, whereas a sort order is defined by a combination of the sort order ID and the character set ID csid):

```
1> select name, id from master..syscharsets
2> where csid = 0
3> go
```

This query can return more than one row if alternate character sets are installed. In this case, check the sysconfigures table or use sp_helpsort to determine the default character set.

If Error 5824 is raised with State 1, and the specified character set does not exist, you will need to install it. Refer to “Configuring Character Sets, Sort Orders, and Message Language” in the *System Administration Guide*.

2 Using the character set from step 1, determine whether the specified sort order exists in syscharsets:

```
1> select name, id, csid, description
2> from master..syscharsets
3> where csid = [charset_id]
4> go
```

where [charset_id] is the id from step 1.

3 If the character set and/or sort order data exist in syscharsets, check the table integrity:

```
1> dbcc checktable (syscharsets)
2> go
```

If errors are returned, recover from them by using the sections in this manual that describe those errors.

4 If errors persist, contact Sybase Technical Support. Depending on the type of problem you found, you may be able to take corrective action, such as manually changing the character sets or rebuilding the index on syscharsets.

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR Rollup level
- Server error log
- Output of dbcc checktable and the contents of syscharsets
- Text of all error messages
- sybinit log, if the sybinit utility was used to attempt changing the character set and/or sort order

Version in which this error is raised

All versions

Error 5846

Severity

16

Message text

```
Illegal value '%ld' specified for configuration option
'%s'. The legal values are between '%ld' and '%ld'.
```

Explanation

Error 5846 occurs when you specify a value for a configuration parameter that is lower than the lowest value allowed or higher than the highest value allowed.

Note Adaptive Server checks to make sure you have not set configuration parameters to values outside the allowed ranges. However, if you set a parameter such as number of open objects to a value that is too high for the amount of memory you have configured and that value is within the allowed range, Error 5846 will not be raised and you may not be able to start Adaptive Server. Use sp_configure with the verify or read option to point out this sort of problem after you make configuration parameter changes and before you restart Adaptive Server.

Action

Refer to "Setting Configuration Parameters" in the *System Administration Guide* for legal value ranges for configuration parameters. Reset the value of the parameter named in the 5846 error message.

Additional information

Refer to the *Reference Manual* for information about sp_configure.

Version in which this error is raised

All versions

Error 5847

Severity

16

Message text

```
Configuration Error: Configuration file '%s' does not
exist or you do not have the required read/write
permission for the file.
```

Explanation

You can set or change the value of Adaptive Server configuration parameters in one of two ways:

- By executing the system procedure sp_configure with the appropriate parameters and values.
- By directly editing your configuration file and then invoking sp_configure with the configuration file option.

Error 5847 occurs when Adaptive Server cannot find the configuration file you specified in an sp_configure command or the user who started Adaptive Server does not have the required operating system-level read or write permission for the file.

Action

Check the name of the configuration file you specified in the sp_configure command. If the name was wrong, enter the command again with a valid file name.

If the file name was not the problem, check the permissions on the configuration file you specified in the sp_configure command. If the user who started Adaptive Server does not have the required operating system-level read or write permission for the file, ask your System Administrator to change the protection on the file. Or have a different user, who has the appropriate access to the configuration file, start Adaptive Server.

Additional information

Refer to "Setting Configuration Parameters" in the *System Administration Guide* for information about configuration parameters and files.

Refer to the *Adaptive Server Reference Manual* for information about sp_configure.

Version in which this error is raised

All versions

Error 5848

Severity

16

Message text

```
Warning: Unknown parameter '%s' found on line '%ld' of
configuration file '%s'. This parameter is ignored. Edit
the configuration file to remove or correct this
parameter.
```

Explanation

You can set or change the value of Adaptive Server configuration parameters in one of two ways:

- By executing the system procedure sp_configure with the appropriate parameters and values.
- By directly editing your configuration file and then invoking sp_configure with the configuration file option.

Error 5848 occurs during startup or run time when Adaptive Server detects a line in the configuration file that contains an unknown parameter. Adaptive Server ignores the unknown parameter.

Action

Edit your configuration file. Using the parameter name and line number specified in the error message, remove or correct the incorrect entry.

Additional information

Refer to "Setting Configuration Parameters" in the *System Administration Guide* for information about configuration parameters and files.

Version in which this error is raised	All versions

Error 5849

Severity	16
Message text	`Verification failed for parameter '%s'.`
Explanation	Adaptive Server runs a verification check on configuration parameters in the following circumstances: • When you use sp_configure with the verify option to perform validation checking on the values in the configuration file you specify. This can help you avoid configuring Adaptive Server with invalid configuration values. • When you use sp_configure with the read option to perform validation checking on the values in the configuration file you specify and to read those values that pass validation into Adaptive Server. If any parameters are missing from the configuration file you specified, current values for those parameters are used. Error 5849 occurs in one of those situations when Adaptive Server detects a verification error for the parameter named in the error message.
Action	Use sp_configure or edit your configuration file to correct the incorrect value for the parameter named in the error message. Run the verification check again to confirm that the problem has been corrected.
Additional information	Refer to “Setting Configuration Parameters” in the *System Administration Guide* for information about configuration parameters and files.
Version in which this error is raised	All versions

Error 5850

Severity	16
Message text	`Notification failed for parameter '%s'.`
Explanation	You can use sp_configure with the verify or read option to perform validation checking on the parameters in a configuration file.

Configuration parameters are either dynamic or static. Dynamic parameters go into effect as soon as you execute sp_configure. Static parameters require Adaptive Server to reallocate memory, and thus they take effect only after Adaptive Server has been restarted.

During validation checking, if the value of a dynamic parameter is different in the configuration file than the current run value, Adaptive Server notifies itself that it should perform validation checking on the new value and notifies the user by writing an entry in the error log. Error 5850 occurs when that notification fails. It is caused by an Adaptive Server problem.

Action

Try the sp_configure with verify or read command again. If Error 5850 occurs again, call Sybase Technical Support.

Additional information

Refer to "Setting Configuration Parameters" in the *System Administration Guide* for information about configuration parameters and files.

Before calling Technical Support, have the following information available:

- Server version and SWR Rollup level
- Server error log
- Text of all error messages
- A copy of the configuration file you are using

Version in which this error is raised

All versions

Error 5851

Severity

16

Message text

```
Configuration Error: Configuration file '%s' has an
unknown format on line %d.
```

Explanation

You can set or change the value of Adaptive Server configuration parameters in one of two ways:

- By executing the system procedure sp_configure with the appropriate parameters and values.
- By directly editing your configuration file and then invoking sp_configure with the configuration file option.

- By starting Adaptive Server with the -c option to specify a configuration file.

> **Note** For some parameters in the configuration file, the keyword "default" must be in capital letters (for example, for wash size).

The syntax for parameters in the configuration file is:

```
parameter_name={value | DEFAULT}
```

where:

- *parameter_name* is the name of the parameter you want to specify.
- *value* is the numeric value to which you want to set *parameter_name*.
- DEFAULT specifies that you want to use the default value for *parameter_name*. If you specify DEFAULT, do not include *value*.

Error 5851 occurs during start-up or when you run sp_configure with the verify or read option when Adaptive Server detects a syntax error in a line in the configuration file. Adaptive Server ignores the line with the syntax error.

Action

Edit your configuration file. Using the configuration file name and line number specified in the error message, remove or correct the incorrect entry.

Additional information

Refer to "Setting Configuration Parameters" in the *System Administration Guide* for information about configuration parameters and files.

Version in which this error is raised

All versions

Error 5852

Severity

```
16
```

Message text

```
Changing the value of '%s' is not allowed since it is a
static option.
```

Explanation

You can set or change the value of Adaptive Server configuration parameters in one of two ways:

- By executing the system procedure sp_configure with the appropriate parameters and values.
- By directly editing your configuration file and then invoking sp_configure with the configuration file option.

- By starting Adaptive Server with the -c option to specify a configuration file.

Configuration options are either dynamic or static. Dynamic parameters go into effect as soon as you execute sp_configure. Static parameters require Adaptive Server to reallocate memory; thus, they take effect only after Adaptive Server has been restarted.

Static parameters include:

Note Not all parameters listed here may be available in your Adaptive Server version.

- additional network memory
- address lock spinlock ratio
- allow nested triggers
- allow resource limits
- allow sql server async i/o
- audit queue size
- cpu grace time
- default character set id
- default database size
- default fill factor percent
- default language id
- default network packet size
- default sortorder id
- disable character set conversions
- disk i/o structures
- enable cis
- enable HA
- enable java
- enable xact coordination
- engine adjust interval

- esp execution stack size
- event buffers per engine
- identity burning set factor
- lock shared memory
- max async i/os per engine
- max async i/os per server
- max cis remote connections
- max cis remote servers
- max network packet size
- max number network listeners
- max online engines
- max roles enabled per user
- max SQL text monitored
- maximum dump conditions
- memory alignment boundary
- memory per worker process
- min online engines
- number of alarms
- number of aux scan descriptors
- number of devices
- number of dtx participants
- number of extent i/o buffers
- number of languages in cache
- number of large i/o buffers
- number of locks
- number of mailboxes
- number of messages
- number of pre-allocated extents

- number of worker processes
- number of open databases
- number of open indexes
- number of open objects
- number of pre-allocated extents
- number of remote connections
- number of remote logins
- number of remote sites
- number of user connections
- open index spinlock ratio
- open index hash spinlock ratio
- open object spinlock ratio
- page lock spinlock ratio
- partition groups
- partition spinlock ratio
- permission cache entries
- print recovery information
- procedure cache percent
- remote server pre-read packets
- shared memory starting address
- size of global fixed heap
- size of process object fixed heap
- size of shared class heap
- SQL Perfmon Integration *(Windows NT only)*
- sql server clock tick length
- stack guard size
- stack size
- strict dtm enforcement

- table lock spinlock ratio
- tape retention in days
- tcp no delay
- time slice
- total data cache size
- total memory
- txn to pss ratio
- user log cache size
- user log cache spinlock ratio
- user security services *(Windows NT only)*

Error 5852 occurs when you run sp_configure with the read or verify option and Adaptive Server determines that you have changed the value of the static option mentioned in the error message and have not yet shut down and restarted Adaptive Server.

Action

No action is required unless you want the value of the static parameter to be used now. If you do, shut down and restart Adaptive Server.

Additional information

Refer to "Setting Configuration Parameters" in the *System Administration Guide* for information about configuration file parameters and options.

Version in which this error is raised

All versions

Error 5853

Severity

16

Message text

```
'%s' has been modified from '%ld' to '%ld' by the
verification routine.
```

Explanation

Adaptive Server runs a verification check on configuration parameters in the following circumstances:

- When you use sp_configure with the verify option to perform validation checking on the values in the configuration file you specify. This can help you avoid configuring Adaptive Server with invalid configuration values.

- When you use sp_configure with the read option to perform validation checking on the values in the configuration file you specify and to read those values that pass validation into Adaptive Server. If any parameters are missing from the configuration file you specified, current values for those parameters are used.

Error 5853 is an informational message that is displayed in one of those situations when Adaptive Server changes the value for a configuration parameter because of rounding of the value. Rounding can occur when Adaptive Server needs to align to page boundaries.

Action

No action is required.

Additional information

Refer to "Setting Configuration Parameters" in the *System Administration Guide* for information about configuration parameters and files.

Version in which this error is raised

All versions

Error 5854

Severity

18

Message text

```
Configuration Error: Cannot write out file '%s' due to
system error '%s'.
```

Explanation

You can set or change the value of Adaptive Server configuration parameters in one of two ways:

- By executing the system procedure sp_configure with the appropriate parameters and values.
- By directly editing your configuration file and then invoking sp_configure with the configuration file option.
- By starting Adaptive Server with the -c option to specify a configuration file.

Adaptive Server:

- Writes information to the configuration file when you change a parameter using sp_configure or other stored procedures such as sp_cacheconfig.
- Writes information to the configuration file when you run sp_configure with the verify or read option and Adaptive Server needs to change the value for a configuration parameter because of rounding of the value.

- Creates a new configuration file when you use sp_configure with the write and restore options.

Error 5854 occurs when Adaptive Server is unable to write information to the configuration file in the error message because of the system error in the error message. The system error is an error returned by the operating system such as permission denied or no such directory or file.

Action

Refer to your operating system documentation for information about resolving the system error.

Resolve the problem that caused the system error and try your action again.

Additional information

Refer to "Setting Configuration Parameters" in the *System Administration Guide* for information about configuration commands and files.

Version in which this error is raised

All versions

Error 5857

Severity

0

Message text

```
Changing the value of '%s' is not allowed since its
value is calculated or is read only.
```

Explanation

Adaptive Server calculates the value for a number of configuration parameters. Some other configuration parameters are read only.

Calculated parameters include:

- allow syb_sendmsg function
- executable codesize + overhead
- syb_sendmsg port number
- total data cache size

Read-only parameters include:

- o/s asynch i/o enabled
- o/s file descriptors
- master device name
- master device mirror name
- configuration file

Error 5857 occurs when you try to change the value of the parameter named in the error message, since that parameter is either calculated by Adaptive Server or it is read-only.

Action No action is necessary.

Additional information Refer to "Setting Configuration Parameters" in the *System Administration Guide* for information about configuration parameters.

Version in which this error is raised All versions

Error 5859

Severity 16

Message text `The configuration number '%d' is invalid.`

Explanation Each Adaptive Server configuration parameter has a unique identification number assigned to it. For user-settable parameters, this number is stored in the config column of the sysconfigures table.

Adaptive Server uses these configuration numbers to identify the option to be displayed, loaded, modified, or verified when you run sp_configure. Error 5859 occurs when the configuration number
Adaptive Server tries to use to identify the option is invalid.

This error can be caused by corruption of sysconfigures or by an Adaptive Server problem.

Action Call Sybase Technical Support.

Have the following information ready:

- Server version and SWR Rollup level
- Server error log
- Text of all error messages
- select * from sysconfigures output

Version in which this error is raised All versions

Error 5861

Severity

16

Message text

```
The 'total memory', '%ld', is not enough to change the
parameter '%s' to '%ld'.
```

Explanation

The total memory configuration parameter sets the size of memory, in 2K units, that Adaptive Server allocates from the operating system.

Error 5861 occurs when you try to change the value of the parameter named in the error message to the value shown in the error message because the value of the total memory parameter is not high enough to support the change.

If the error occurs when you are running sp_configure, the value you asked to be changed is not changed. If the error occurs during start-up (after you modified the configuration file), Adaptive Server does not start.

Action

If the error occurred when you were running sp_configure, specify a lower value for the parameter you are changing or raise the value of the total memory parameter. Then enter your original command again.

If the error occurred while you were starting Adaptive Server, edit the configuration file to make the changes, and then start Adaptive Server.

Note There is a worksheet in the *System Administration Guide* to help you determine appropriate values for memory-related configuration parameters.

Additional information

Refer to “Setting Configuration Parameters” in the *System Administration Guide* for information about configuration parameters.

Refer to the *Performance and Tuning Guide* and “Configuring Memory” in the *System Administration Guide* for information about how Adaptive Server uses memory.

Version in which this error is raised

All versions

Error 5863

Severity

18

Message text

```
Invalid group encountered.
```

Explanation

You can use sp_configure with a group name to display values for the configuration parameters belonging to that group. As of Version 11.9.2, groups are:

- Backup/Recovery
- Cache Manager
- Component Integration Services
- Disk I/O
- Error Log
- Extended Stored Procedures
- General Information
- Languages
- Lock Manager
- Memory Use
- Meta-Data Caches
- Network Communication
- O/S Resources
- Parallel Query
- Physical Memory
- Physical Resources
- Processors
- Rep Agent Thread Administration
- SQL Server Administration
- Security Related
- User Environment

When Adaptive Server retrieves information for a configuration parameter group, it attempts to verify the group name. Error 5863 occurs when Adaptive Server is unable to verify the group name.

This error is caused by an Adaptive Server problem.

Action

Call Sybase Technical Support.

Have the following information ready:

- Server version and SWR Rollup level
- Server error log
- Text of all error messages

Version in which this error is raised: All versions

Error 5865

Severity: 18

Message text: `error in os_get_cpu_count: %d`

Explanation: Before Adaptive Server checks the configuration parameters min online engines and current number online engines to make sure they contain valid values, it checks the number of online CPUs on the machine. Error 5865 occurs when Adaptive Server queries the operating system for the number of online CPUs and the operating system returns an error.

This error is caused by an operating system problem.

Action: If the error occurred during Adaptive Server start-up, restart Adaptive Server.

Retry the command you were running when the error occurred.

If you are still getting the 5865 error, call Sybase Technical Support. Have the following information ready:

- Server version and SWR Rollup level
- Text of all error messages

Version in which this error is raised: All versions

Error 5866

Severity: 10

Message text:
```
Configuration file '%s' has been written and the
previous version has been renamed to '%s'.
```

Explanation: Adaptive Server:

- Writes information to the configuration file when you change a parameter using sp_configure.
- Writes information to the configuration file when you run sp_configure with the verify or read option and Adaptive Server needs to change the value for a configuration parameter because of rounding of the value.
- Creates a new configuration file when you use sp_configure with the write and restore options.

Each time you modify a configuration parameter with sp_configure, a new configuration file is created, using the naming convention *file_name.001*, *file_name.002*, *file_name.003* and so on, up to *file_name.999*.

Error 5866 is an informational message that is written to the Adaptive Server error log when Adaptive Server writes information to the configuration file in the message and renames the previous version of the file.

Action

No action is required.

Additional information

Refer to "Setting Configuration Parameters" in the *System Administration Guide* for information about configuration commands and files.

Version in which this error is raised

All versions

Error 5867

Severity

16

Message text

```
The wash size for the %dk buffer pool in cache %s has
been incorrectly configured. It must be a minimum of %d
buffers and a maximum of %d percent of the number of
buffers in the pool.
```

Explanation

Wash size is the point in the cache at which Adaptive Server writes dirty pages to disk. You can use sp_poolconfig to place an entry in the configuration file to specify a wash size for a memory pool.

The default value for wash size depends on your server version. For example, in 11.5 the default wash size for a pool size less than 300MB is 20 percent of the buffers in the pool. (A buffer is a block of pages equal to the I/O size of the pool. All pages in a buffer are read from disk, written to disk, or flushed from the cache simultaneously.)

The minimum value for wash size is 10 buffers and the maximum value is 80 percent of the size of the pool.

Error 5867 occurs when you specify a value for wash size that is smaller than the minimum allowed size or larger than the maximum allowed size.

This error occurs with the following states:

State	Meaning
1	Error 5867 occurs with State 1 when you run sp_configure with the verify or read option after specifying a value for wash size that is smaller than the minimum allowed size or larger than the maximum allowed size.
2	Error 5867 occurs with State 2 when you use sp_poolconfig or update the configuration file directly to specify a value for wash size and that value is smaller than the minimum allowed size or larger than the maximum allowed size.

Action

Determine a valid value for wash size for the buffer pool and cache specified in the error message and run your command again.

Additional information

Refer to the *Reference Manual* for information about sp_poolconfig and wash sizes.

Refer to the *Performance and Tuning Guide* and "Configuring Data Caches" in the *System Administration Guide* for information about the wash area and wash configuration.

Version in which this error is raised

All versions

Error 5868

Severity

16

Message text

```
Invalid type configured for cache %s. Cache type may be
one of mixed OR log only but NOT both.
```

Explanation

You can create, configure, reconfigure, and drop data caches by using the Adaptive Server procedure sp_cacheconfig or by modifying your configuration file.

There are two types allowed for user-defined caches:

- logonly – only transaction log (syslogs) information is stored here.
- mixed – both transaction log (syslogs) information and user data is stored here. This is the default.

Error 5868 occurs when Adaptive Server determines that you have tried to mix different types of cache. For example, Error 5868 would occur if you try to define a cache as log only as well as mixed at the same time.

Action

Use sp_cacheconfig or modify your configuration file to define a valid type for the cache named in the error message.

Additional information

Refer to the *Reference Manual* for information about sp_cacheconfig and cache types.

Version in which this error is raised

All versions

Error 5893

Severity

16

Message text

```
The sum, (%s) + (%s %s) + (%s) + (%s), must be no greater
than '%ld'.
```

Explanation

Error 5893 is raised when the number of file descriptors available to Adaptive Server is less than the number needed by the process.

When you start up Adaptive Server or run sp_configure to change certain configuration parameters (such as number of user connections), the following criterion must be met:

```
number of remote sites + max number network listeners
+ number of user connections
+ (number of devices * max online engines * 2)
<= o/s file descriptors - descriptors for environment files
```

Note On platforms that use light weight engines (NT systems), use 1 instead of max online engines for the calculation.

descriptors for environment files is not a configuration parameter. It accounts for file descriptors used by the server error log and other operational files. It is equal to 10 if NETMIGRATE feature is on and 8 otherwise.

Action

Increase the available file descriptors or reduce the value of one or more configuration parameters in the above expression.

	Many operating systems allow you to configure the amount of file descriptors available per process. See the Adaptive Server installation and configuration guide for your platform for details on how to increase the available file descriptors.
Additional information	The o/s file descriptors parameter indicates the maximum per-process number of file descriptors configured for your operating system. This parameter is read-only and cannot be configured through Adaptive Server.
Version in which this error is raised	All versions

Process Kill Errors

This section contains error messages pertaining to killing Adaptive Server processes.

Error 6103

Severity

17

Message text

```
Unable to do cleanup for the killed process; received
Msg %d.
```

Explanation

Upon termination of any Adaptive Server process, the server attempts to clean up by releasing resources used by the terminated process. Error 6103 occurs when the server is unable to complete this cleanup. The “Msg %d” appearing in the 6103 message text identifies the error that was received during cleanup.

Error 6103 usually occurs after some other error or when an Adaptive Server process is aborted. This error can be serious as it can cause data corruption or an abnormal Adaptive Server shutdown.

Error 6103 is usually caused by the errors that occurred before it or by an Adaptive Server problem.

Action

Review the error log to identify errors that occurred before the 6103 error. Use the information in this manual to recover from those errors. If errors are not covered in this manual, call Technical Support for assistance.

To recover from Error 6103 after reviewing other errors in the error log, shut down and restart Adaptive Server.

If you cannot perform an Adaptive Server shutdown, kill the Server process at the operating system level.

Run dbcc checkdb for the database that was being accessed by the killed process to determine whether data corruption has occurred. If other errors are displayed when dbcc checkdb is run, refer to the writeups in this manual for those errors to determine how to recover from them.

If the 6103 error occurs again or resources (such as locks) are not released, call Sybase Technical Support.

Version in which this error is raised

All versions

Error 6107

Severity

14

Message text

```
Only User processes can be KILLed or SYB_TERMINATEd.
```

Explanation

A process is a task that is being carried out by Adaptive Server. Processes can be initiated by a user giving a command, or by Adaptive Server itself. You can see information about processes by running the system procedure sp_who.

Error 6107 is raised following an improper attempt to kill an Adaptive Server process; for example, it may be raised if you try to kill a background process. The error is raised with the following states:

State	Meaning
1	Attempted to kill a process that is not a user process (pre-12.0). Attempted to kill a process that is not a user process or an agent process for quiesce database (12.0).
2	Attempted to kill an agent process for quiesce database, but the spid requesting the kill is not the parent process of the agent (12.0).

Action

Use sp_who to check the currently running processes and identify the process in question. Refer to “Killing Processes” in the *System Administration Guide* for a description of process status values and the effect of the kill command. If the error was raised when you attempted to clear a "killable" user process, take the following steps:

1 Perform a checkpoint in each database:

```
1> use database_name
2> go

1> checkpoint
2> go
```

2 Shut down the server using the nowait option:

```
1> shutdown with nowait
2> go
```

3 Restart the server.

Additional information

If you need to terminate a batch program or other application which logs in to Adaptive Server, always terminate the server process first using the kill command before attempting to terminate the batch at the operating system level.

Version in which this error is raised

All versions

Timestamp Errors

This section contains error messages pertaining to Adaptive Server timestamps.

Error 6901

Severity

21

Message text

```
Overflow on High component of timestamp occurred in
database %d. Database table possibly corrupt.
```

Note This error may be caused by a hardware problem.

Explanation

Each database has one global timestamp which is kept in memory in a structure called a dbtable. The timestamp is not a date and time; rather, it is a sequence number (6 bytes, consisting of a high and low portion) which keeps track of modifications of pages within the database. Each time a page in the database is modified, the page gets the current timestamp and the global timestamp gets incremented. The global timestamp "travels" with a database through dumps, loads, and upgrades. Timestamps are critical during recovery, to ensure that transactions are properly rolled forward or rolled back. The global timestamp is also used for timestamp columns.

The maximum value for timestamps is very high (0xffff 0xffffffff), so it is very unlikely that under normal circumstances a database will reach this limit. However, use of some unsupported database maintenance commands can cause the global timestamp to reach the maximum value.

Each time dbtable is created for the database, Adaptive Server checks the timestamp value. Error 6901 is raised if the timestamp value exceeds the maximum.

Action

Database dump/load is *not* an option to address this problem since that will not reset the timestamp value. The only way to recover from Error 6901 is to drop and re-create the database, using the following steps:

1 bcp out all table data.

2 defncopy out all procedures, rules, defaults and triggers.

3 Use the drop database command to drop the database.

4 Use the create database command to re-create the database.

5 Create all tables.

6 bcp in all table data.

7 Create indexes.

8 defncopy in all procedures, rules, defaults and triggers.

9 Check that the user IDs in sysusers in the restored database are consistent with master..syslogins. Mismatched user IDs in databases can lead to permission problems.

Note Scripts are recommended for reliable re-creation of the database objects.

Additional information

Refer to the glossary for additional information about timestamps and dbtable. Refer to "Developing a Backup and Recovery Plan" in the *System Administration Guide* for complete information about how to safely re-create databases.

Version in which this error is raised

All versions

Error 6902

Severity

21

Message text

```
Page timestamp value falls between the old and new
timestamps from log. Page #=%ld, object id = %ld, page
timestamp=%04x %08lx. Log: old timestamp=%04x %08lx,
new timestamp=%04x %08lx.
```

Note This error may be caused by a hardware problem.

Explanation

The objective of recovery algorithms in Adaptive Server is to bring all the databases to a consistent state: that is, each transaction is either completely applied to the database or no part of it is applied. (Database recovery occurs at Adaptive Server startup as well as when a load database or load transaction command is processed.)

Whenever a page is modified during run time, the page timestamp before the change and the page timestamp after the change are recorded in the log record for that modification. By comparing the timestamp on the page at recovery time with that of the timestamps in the log record, Adaptive Server can tell whether the page is in the state before or after the change was made to it.

Error 6902 occurs during Adaptive Server recovery if the timestamp of the page which is being recovered is neither the old log timestamp nor the new log timestamp. This is a fatal error and causes the recovery process to fail.

Error 6902 is probably a result of a more serious underlying error. Common causes of this error are:

- Data corruption during normal processing
- Abnormal shutdown of Adaptive Server during a transaction update
- Hardware errors.

Action

Recover your database from backups or call Sybase Technical Support.

To prevent this error from occurring in the future, use one of the following sections, depending on what caused the 6902 error.

Shutdown Method

shutdown with nowait or abnormal shutdown of Adaptive Server should be avoided, if possible, after a heavy update period in your databases:

- Use shutdown to shut down Adaptive Server instead of shutdown with nowait after a period of heavy update activity in your databases.

Or:

- checkpoint each database that is being used before shutting down Adaptive Server.

Hardware Errors

Check the Adaptive Server error log to determine whether there are other indications of hardware problems, such as kernel messages reporting I/O errors.

Check the operating system error log or diagnostics utilities for I/O errors.

Using *dump transaction with no_log*

Using dump transaction with no_log can, in some circumstances, result in a 6902 error, especially when used after a heavy update period. Therefore, do not use dump transaction with no_log unless your transaction log is 100 percent full. Try using dump transaction with truncate_only first and use the no_log option only as a last resort.

If you find that the no_log option is used periodically, you might not be performing database or transaction log dumps often enough, or you may not have enough space allocated to the database or transaction log. Check for any occurrences of dump transaction with no_log in any of your scripts or procedures and replace them with dump transaction with truncate_only. For more information about dump transaction with no_log, refer to “dump transaction” in the *Reference Manual* and ”Backing Up and Restoring User Databases” in the *System Administration Guide*for details about the special dump transaction options.

Warning! Use dump transaction with no_log only when there are no active users in the database.

Version in which this error is raised

All versions

Text Manager Errors

This section contains error messages for the Adaptive Server Text Manager.

Error 7101

Severity

17

Message text

```
Unable to allocate new text value, dbid %d.
```

Explanation

When you create a table with text or image columns, the data is stored on its own chain of text/image pages. Error 7101 is raised when you attempt to insert data into one or more rows of a text column in a table, but there is not enough space to allocate the new text pages.

Error 7101 is raised with the following states:

State	Meaning
1	Could not allocate the first page to add a new value in a table's text column.
2	Could not allocate additional pages for an in-progress insert to a text column.

Error 7101 is similar to Error 1105, States 1 and 2, but it is specific to text/image allocation.

Action

1 Check the available space in the database:

```
1> use <database_name>
2> go
1> sp_spaceused
2> go
```

2 Identify the table by examining the query that raised the error. (You can also run sp_helpsegment on the database segments; rows with indid = 255 refer to the tables with text columns.)

3 Check the space usage in the table:

```
1> sp_spaceused <table_name>, 1
2> go
```

To find the space used by text/image columns in this report, look for objects named "t" plus the table name.

4 Use alter database to increase the storage space for the text data.

Additional information

By default, a chain of text pages is placed on the same device as its table. You can specify a separate device for future space allocations to the text pages by using the sp_placeobject system stored procedure. This is desirable when you expect substantial growth in text storage.

See “Placing Text Pages on a Separate Device” in the *System Administration Guide* for details.

Version in which this error is raised

All versions

Error 7105

Severity

22

Message text

```
TEXT/IMAGE page %d does not have a next page, although
it should.
```

Note This error may be caused by a hardware problem.

Explanation

Tables can contain text columns, which are variable-length columns that can hold printable characters, and image columns, which are variable-length columns that can hold binary data. Adaptive Server stores text and image data in a linked list of data pages separate from the rest of the table. Each 2K text or image page stores a maximum of 1800 bytes of data.

Error 7105 occurs when Adaptive Server attempts data retrieval or another operation on a table's text or image data, and the operation fails due to an unexpected termination of the page chain. If the operation involved copying data, the problem could be on the source side (the page chain being read) or the destination side (the page chain being written out). The state of the 7105 message raised can help you determine the location of the problem.

Error 7105 is raised with the following states:

State	Meaning
1	There was an unexpected end to the data page chain while sending text data to the client. At the time of the error, a partial character of a multi-byte character set had been sent.
2	There was an unexpected end to the page chain while sending text data to the client. Unlike State 1, the error occurred while positioned at a whole character boundary.
3	There was an unexpected end to the page chain while preparing to read text data.

State	Meaning
4	There was an unexpected end to the page chain while traversing the chain to convert character counts to byte counts for a multi-byte character set.
5	A source data page could not be found while copying data from a text or image page to a log record.
6	A source data page could not be found while copying data from the client to a text or image page.
7	A destination data page could not be found while copying data from a text or image constant chain to a text or image data page.
8	A destination data page could not be found while copying data from one text or image page to another.
9	There was an unexpected end to the page chain while sending image data to the client.

Error 7105 indicates a bad text or image page chain. Error 2574 is a related error; it is raised when a dbcc check finds problems with a table's text or image page chain.

Action

If Error 7105 is raised when you try to select text or image data from within a delete trigger, and the size of the data exceeds one page, this is due to an Adaptive Server problem. Contact Sybase Technical Support for assistance.

In other cases, the problem is typically due to a small number of bad rows in a table with text or image data, and can be corrected by dropping and re-creating those rows. Use the following steps to correct the problem:

1 Identify the table. Refer to Chapter 2, "How to Find an Object Name from a Page Number" for information.

2 Determine the first page of the text/image chain using dbcc pglinkage, with the starting page being the page number from the error message. For example, assuming dbid=7 and starting page number 531:

```
1> dbcc pglinkage(7, 531, 0, 1, 0, 0)
2> go
Object ID for pages in this chain = 208003772.
Beginning of chain reached.
Page : 531
Page : 532
Page : 530
3 pages scanned. Object ID = 208003772.
Last page in scan = 530.
```

Page 530 is the first page of the text/image chain.

3 Identify the row(s) with the bad data using the following command:

```
1> select column1 from table_name
2> where convert (int, textptr(text_colname))
3> = first_page_of_chain
4> go
```

where *column1* is any column or columns in your table which will uniquely identify the rows.

4 Drop the row(s) in question.

5 Insert the rows back into the table.

6 Run dbcc checktable on the table to verify that the problem has been corrected.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level
- Server error log
- Text of all error messages
- Output of dbcc page and dbcc pglinkage

Version in which this error is raised

All versions

Error 7114

Severity

22

Message text

```
Page %ld is not a valid text page.
```

Note This error may be caused by a hardware problem.

Explanation

Error 7114 occurs when Adaptive Server attempts to access or update a table's text or image data, but the retrieved page is found to be a non-text page.

Text page corruption can give rise to 605 errors ("An attempt was made to fetch logical page '%ld' in database '%.*s' from cache '%.*s'. Page belongs to object with id '%ld', not to object '%.*s'.) when you run dbcc's on the affected data.

This error may be due to hardware or operating system errors, or an Adaptive Server problem.

Action

Contact Sybase Technical Support for assistance when this error is raised. Before calling, collect relevant information using the following steps:

1 Note the page number in the error message. Using the procedure in Chapter 2, "How to Find an Object Name from a Page Number", identify the type of page and the table associated with the page.

2 Check the integrity of the table's text page chains:

```
1> dbcc textalloc(tablename)
2> go
```

Additional information

Have the following information ready when you call Sybase Technical Support:

- Adaptive Server version and SWR rollup level
- Server error log
- Text of all error messages
- Output of dbcc checks.

Version in which this error is raised

All versions

Error 7130

Severity

16

Message text

```
WRITETEXT with no log is not valid at this time.
Database %s must first be configured for this.
```

Explanation

This error occurs when you try to execute the writetext command or dbwritetext routine in a database which does not have the sp_dboption option select into/bulkcopy enabled. Because these operations do not log the changes to your data, the option must be enabled.

Often, writetext and dbwritetext run against a temporary table. The select into/bulkcopy option must be enabled in tempdb before these will run successfully.

Action

Enable the sp_dboption option select into/bulkcopy on all databases affected by the writetext just before you perform this nonlogged operation. To do this, the System Administrator or Database Owner can use the following steps:

1 Enable the option using this procedure:

```
1> use master
```

```
2> go
1> sp_dboption database_name,
2> "select into/bulkcopy", true
3> go

1> use database_name
2> go

1> checkpoint
2> go
```

2 If you want to verify that the change is active, use this procedure:

```
1> sp_helpdb database_name
2> go
```

3 Execute the nonlogged operation.

4 When the nonlogged operation is complete, dump the database using normal system administration procedures.

5 Return the database to its original condition by disabling the select into/bulkcopy option:

```
1> use master
2> go
1> sp_dboption database_name,
2> "select into/bulkcopy", false
3> go

1> use database_name
2> go

1> checkpoint
2> go
```

Repeat step 2 above if you want to verify that the change is active.

Additional information

For more information about writetext and dbwritetext, refer to the *Adaptive Server Reference Manual* and the *Open Client DB-Library Reference Manual*.

For more information about database dumps and recoverability, refer to "Developing a Backup and Recovery Plan" in the *System Administration Guide*.

Version in which this error is raised

All versions

Error 7134

Severity

16

Message text

```
The text table and the table referenced by the text
pointer disagree.
```

Explanation

Adaptive Server stores text and image data in text extents separate from the rest of the table. Each extent contains text data for a given table, and may contain multiple text chains corresponding to text contained in different rows and text columns of the table.

The Adaptive Server Text Manager validates text pages during processing. Error 7134 is raised if the Text Manager finds that the object ID of a text page does not match the object ID of the text extent to which the page belongs.

Action

If Error 7134 is raised when you select text columns within a trigger, note that Adaptive Server does not provide the ability to convert text datatypes within a trigger. You must recode the trigger to resolve this problem.

If the error is raised in some other context not involving triggers, the problem is typically due to a small number of bad rows in a table with text or image data; you may be able to correct it by dropping and re-creating those rows. Use the following steps:

1. Identify the table for which the error was raised. Unless there is a stack trace which provides this information, you will need to know what query caused the error.
2. Run dbcc checktable on the table.
3. Correct any errors detected by dbcc checktable using the sections in this manual that refer to those errors.
4. If the error persists, contact Sybase Technical Support for assistance.

Additional information

When calling Technical Support, have the following information available:

- Server version and SWR Rollup level
- Server error log
- Text of all error messages
- Text of the query that raises the error

Version in which this error is raised

All versions

Distributed Database Network Errors

This section contains error messages for distributed database network code.

Error 7201

Severity

17

Message text

```
Can't open a connection to site '%S_SRVID' because there
are no remote sites available.  Rerun when there are
fewer active remote sites, or contact a user with System
Administrator (SA) role.
```

Explanation

The number of remote sites parameter (called remote sites in pre-11.0 releases) determines the number of site connection handlers that can be active at any one time. Site handlers maintain connections between Adaptive Servers, Backup Servers, and Open Server applications.

Error 7201 is raised if there are no site handlers available when the server attempts to open a connection to another server (named in the error text). This occurs when the number of remote entries in sysservers exceeds the number of remote sites. As an example, the error may be seen when you execute a dump command, since this requires communication with Backup Server, which in turn requires a site handler.

Action

7201 errors are corrected by configuring sufficient remote sites.

1 Check the number of remote sites:

```
1> sp_configure "number of remote sites"
2> go
```

The Run Value in the output shows the number of remote sites that can access the server simultaneously.

2 Increase the value of number of remote sites as necessary. This is a static parameter, so you must shut down and restart the server if you change the value.

Additional information

See "Managing Remote Servers" in the *Security Administration Guide* for more information.

Version in which this error is raised

All versions

Error 7205

Severity 18

Message text
```
Can't open a connection to site '%S_SRVID'. See the
error log file in the Adaptive Server boot directory.
```

Explanation

This error occurs when Adaptive Server is unable to execute a remote procedure call (RPC) on a remote Adaptive Server, Backup Server, or Open Server application.

Note In this writeup, "Server" is used to indicate an Adaptive Server, Backup Server, or Open Server application.

Action

To determine why this error is occurring, perform the following steps:

1 Check whether the remote Server is running by issuing an isql statement similar to the following:

```
% isql -Sname_of_remote_server
```

If you receive the following message:

```
Operating-system error:
Invalid argument
DB_LIBRARY error:
Unable to connect: Adaptive Server is unavailable or does not exist.
```

then the remote Server may not be running. Log in to the remote Server machine and check whether the Server is running. If it is not running, start it.

Check whether you receive the following message after issuing the isql statement:

```
Operating-system error:
No such file or directory
DB_LIBRARY error:
Could not open interface file.
```

If you do, this indicates that your *SYBASE* environment variable (for UNIX, OS/2 and Novell users) or your *SYBASE* logical name (for OpenVMS users) is not set to the correct directory.

The value of the *SYBASE* parameter has to be set to the full path name of the *SYBASE* home directory where the interfaces (or *sql.ini*) file can be found. The value of the *SYBASE* parameter specifies the device as well as the path for *SYBASE* home. If you move the *SYBASE* directory structure to a new device, make sure the device is included in the *SYBASE* environment variable string (UNIX, OS/2, and Novell) or in the *SYBASE* logical name string (OpenVMS).

2 Check the error log for both the sending Server and the receiving Server. If other errors exist, refer to the writeups in this manual for information about recovering from them.

3 If, after resolving any problems found in steps 1 and 2, you still get Error 7205, call Sybase Technical Support with the information listed below.

Additional information

For more information about managing remote Adaptive Servers, refer to "Managing Remote Servers" in the *Security Administration Guide.*

Before calling Sybase Technical Support, have the following information ready:

- Server version and SWR Rollup level
- Server (and Backup Server if appropriate) error log for both the local and remote Servers
- Operating system error log
- sp_helpserver and sp_configure output
- Contents of *interfaces* file
- Text of all error messages

Version in which this error is raised

All versions

Error 7207

Severity

17

Message text

```
Can't open a connection to site '%S_SRVID' because there
are no remote connections available. Rerun when there
are fewer remote users and/or fewer local users making
remote accesses, or contact a user with System
Administrator (SA) role.
```

Explanation

This error occurs when Adaptive Server is unable to open a connection to a remote Adaptive Server, Backup Server, or Open Server application because no remote connections are available.

There may be too many local users with active remote connections. Your System Administrator may also need to configure the server to allow more remote connections.

Action

1 Use the following query to check the remote connections currently in use:

```
1> use master
2> go
1> select count(*) from sysprocesses a, sysservers b
2> where a.hostname=b.srvname
```

2 Do one of the following to resolve the problem:

- wait until some of the presently active connections shut down, then retry your query.
- configure your Adaptive Server to allow more remote connections. The relevant configuration parameters are:
  - number of remote connections
  - number of remote logins

  These are static parameters. You will need to restart Adaptive Server after changing the values.

Refer to the *System Administration Guide* for more information.

Additional information

If you are on version 11.5 or higher, note that each simultaneous connection to XP Server for extended stored procedure execution uses up to one remote connection each.

Version in which this error is raised

All versions

Error 7211

Severity

18

Message text

```
Can't open a connection to site '%S_SRVID' because 'set
up channel' message was not acknowledged.
```

Explanation

This error occurs when Adaptive Server is unable to execute a remote procedure call (RPC) on a remote Adaptive Server, Backup Server, or Open Server application because the remote Server did not acknowledge the connection message sent by the local Server.

One reason for this error is that the physical connection has died before a response was received. The transaction is then aborted.

It may also occur if two access attempts occur within such a short time frame that the first physical connection established has not yet timed out.

Note In this writeup, "Server" is used to indicate an Adaptive Server, Backup Server, or Open Server application.

Action

To find out why this error is occurring, perform the following steps:

1 When automatic timeouts are enabled, a timeout occurs when there is no activity for one minute. If you believe timeouts are occurring, consider disabling automatic timeouts by issuing the following command:

```
1> sp_serveroption server_name "timeouts", false
2> go
```

where *server_name* is the name of the Adaptive Server for which you are disabling automatic timeouts.

2 Check the error log for both the sending server and the receiving server. If other errors exist, refer to the writeups in this manual for information about recovering from them.

3 If, after resolving any problems found in steps 1 and 2, you still get Error 7211, call Sybase Technical Support.

Additional information

For more information about managing remote Adaptive Servers, refer to "Managing Remote Servers" in the *Security Administration Guide.*

Before calling Sybase Technical Support, have the following information available:

- Server version and SWR Rollup level
- Server (and Backup Server if appropriate) error log for both the local and remote Server
- Operating system error log
- sp_helpserver and sp_configure output
- Contents of *interfaces* file

- Text of all error messages

Version in which this error is raised

All versions

Error 7212

Severity

11

Message text

```
Can't close connection %d on site '%S_SRVID' because it
is not active.
```

Explanation

This error occurs when Adaptive Server is unable to close a connection for a remote site on an Adaptive Server, Backup Server, or Open Server application because the remote connection is not active. *%d* is the logical connection channel number and *%S_SRVID* is the remote Adaptive Server name.

One reason for this error is that the physical connection died before a response was received. The transaction is then aborted.

Error 7212 occurs with the following states:

State	Meaning
2	Error 7212 occurs with State 2 when Adaptive Server is attempting to close the logical connection to the remote site and the logical connection is not found.
3	Error 7212 occurs with State 3 when Adaptive Server is sending an attention packet to the remote connection and the logical connection is not found.

Action

To find out why this error is occurring, perform the following steps:

1 When automatic timeouts are enabled, a timeout occurs when there is no activity from a remote user for one minute. If you believe timeouts are occurring, consider disabling automatic timeouts by issuing the following command:

```
1> sp_serveroption server_name, "timeouts", false
2> go
```

where *server_name* is the name of the Adaptive Server for which you are disabling automatic timeouts.

2 Check the error log for both the sending Server and the receiving Server. If other errors have been raised, refer to the writeups in this manual for information about recovering from them.

3 If, after resolving any problems found in steps 1 and 2, you still get Error 7212, call Sybase Technical Support.

Additional information

For more information about managing remote Adaptive Servers, refer to "Managing Remote Servers" in the *Security Administration Guide.*

Before calling Sybase Technical Support, have the following information available:

- Server version and SWR Rollup level
- Server (and Backup Server if appropriate) error log for both the local and remote Server
- Operating system error log
- sp_helpserver and sp_configure output
- Contents of interfaces file
- Text of all error messages

Version in which this error is raised

All versions

Error 7214

Severity

11

Message text

```
No site name received in server login packet. Server
login rejected.
```

Explanation

This error occurs when executing a remote procedure call (RPC) from an Adaptive Server which has not been given a name via the sp_addserver system procedure. The Adaptive Server name referred to here is stored in the *sysservers* system table as well as in the *@@servername* global variable.

Action

1 Use sp_addserver to assign a name to the Adaptive Server initiating the RPC, while using that server:

```
1> sp_addserver sql_server_name, local
2> go
```

Refer to "sp_addserver" in the *Reference Manual* for information about valid Adaptive Server names.

2 Restart Adaptive Server to make the change take effect.

3 Verify that the server initiating the RPC has been given a name:

```
1> select @@servername
2> go
```

You can also check the Adaptive Server start message. The output will show the (local) Adaptive Server initiating the RPC. If the output is NULL, then the local server has not yet been assigned a name.

Additional information

For more information, refer to:

- Server Is Unnamed.
- "Identifiers" in the *Reference Manual*

Version in which this error is raised

All versions

Error 7215

Severity

18

Message text

```
Can't send to site '%S_SRVID'.
```

Explanation

This error occurs when you try to execute a remote procedure call (RPC) on a remote Adaptive Server or Backup Server that is found in the local interfaces file but is not currently running.

Note In this writeup, "Server" is used to indicate an Adaptive Server or a Backup Server.

Action

1 Check whether the remote Server is running by issuing an isql statement similar to the following:

```
% isql -Sname_of_remote_server
```

If you receive the following message:

```
Operating-system error:
Invalid argument
DB_LIBRARY error:
Unable to connect: Adaptive Server is unavailable or does not exist.
```

then the remote Server may not be running. Log in to the remote Server machine and check whether the Server is running. If it is not running, start it.

2 Check whether you receive the following message after issuing the isql statement:

```
Operating-system error:
No such file or directory
```

```
DB_LIBRARY error:
Could not open interface file.
```

If you do, this indicates that your *SYBASE* environment variable (for UNIX, OS/2 and Novell users) or your *SYBASE* logical name (for OpenVMS users) is not set to the correct directory.

The value of the *SYBASE* parameter has to be set to the full path name of the *SYBASE* home directory where the interfaces file can be found. The value of the *SYBASE* parameter specifies the device as well as the path for *SYBASE* home. If you move the *SYBASE* directory structure to a new device, make sure the device is included in the *SYBASE* environment variable string (UNIX, OS/2, and Novell) or in the *SYBASE* logical name string (OpenVMS).

Note There are changes to the ASE directory structure in version 12.0. Refer to your Installation and Configuration guide.

Additional information

For more information about managing remote Adaptive Servers, refer to "Managing Remote Servers" in the *Adaptive Server Enterprise Security Administration Guide.*

Before calling Technical Support, have the following information available:

- Server version and SWR Rollup level
- Server (and Backup Server if appropriate) error log for both the local and remote Server
- Operating system error log
- sp_helpserver and sp_configure output
- Content of *interfaces* file
- Text of all error messages

Version in which this error is raised

All versions

Error 7218

Severity

11

Message text

```
Site '%.*s' not found in sysservers.
```

Explanation

This error occurs when the remote Adaptive Server name in a remote procedure call (RPC) is not defined in the sysservers table on the Adaptive Server initiating the RPC.

Action

Add the remote Adaptive Server name to the sysservers table on the Adaptive Server that initiated the RPC by typing the following:

```
1> sp_addserver server_name
2> go
```

Refer to "sp_addserver" in the *Reference Manual* for information about valid Adaptive Server names.

The easiest way to name an Adaptive Server is by assigning it the name stored in the *DSQUERY* environment variable or the name which is used in the -S option when the server is accessed via isql. For example:

```
% isql -SPRODUCTION
```

In the above example, PRODUCTION is the Adaptive Server name. You can check to determine whether the Server has been added to the *sysservers* table by issuing the following command:

```
1> sp_helpserver
2> go
```

Additional information

For more information, refer to sp_addserver, sp_helpserver, and "Identifiers" in the *Reference Manual*.

Version in which this error is raised

All versions

Error 7220

Severity

11

Message text

```
Site '%s' not found in interfaces file.
```

Explanation

When you execute a remote procedure call (RPC) from one Adaptive Server to another, the local interfaces file of the initiating Adaptive Server has to contain information about the remote Adaptive Server.

Error 7220 occurs in one of the following situations:

- When you try to execute an RPC on an Adaptive Server which is not referenced in the local interfaces file. Its name can, however, be found in the sysservers table.

- When you attempt to use the dump or load commands and the Backup Server is not listed in the *interfaces* file.

Note Throughout this writeup, your network addressing file is referred to as your "*interfaces* file."

Action

Eliminate the 7220 error by following the appropriate instructions below.

The Remote Adaptive Server is Not in the Local *interfaces* File

Use one of the following to check the items in the steps below:

- sqledit (for OS/2)
- sqledit or sybinit (for NT)
- sybinit (for all other platforms)
- Determine whether the local *interfaces* file has the query entry for the remote Adaptive Server.
- Determine whether the Adaptive Server name is spelled correctly in the local *interfaces* file and if the tabs are set correctly at the beginning of each row which contains information about the Adaptive Server ports.
- If you cannot find any information about the remote Adaptive Server in any *interfaces* file, you have two options. You can either:
- Go to a backup copy of the local or remote *interfaces* file, and add the name of the remote Adaptive Server to the local *interfaces* file manually. Add the query entry for the remote Adaptive Server to the local *interfaces* file.

  Or, you can:
- Look for lines like the following in your remote Adaptive Server error log (the example is from a SunOS 4.x system). If you do not have a backup copy of the local *interfaces* file, check the remote Adaptive Server's error log for the network name and port numbers at the last restart (that is, at the end of the error log file). Refer to the *Open Client/Server Installation Guide* for more information about Adaptive Server ports.

```
...
00:93/11/08 17:06:01.01 server Opening Master Database ...
00:93/11/08 17:06:01.65 server  Loading Adaptive Server's default
sort order and character set
00:93/11/08 17:06:01.91 kernel network name baldrick, type ether,
port 2056
00:93/11/08 17:06:02.06 server  Recovering database 'master'
```

. . . .

Look for the port number, the network type, and the machine name.

Add an entry to the *interfaces* file like the following:

```
REL110_PRODUCTION
            query tcp ether baldrick 2056
            master tcp ether baldrick 2056
```

Warning! Before you add Adaptive Server's query, master, and console entries to the local *interfaces* file, check that no other Adaptive Server has the same port numbers assigned to it in the local *interfaces* file.

- After changing the *interfaces* file, try accessing the remote Adaptive Server through isql or, if it is not running, restart it.
- Also check your hardware error log for possible hardware failures, in case you still cannot reboot or access the remote Adaptive Server.

The Backup Server is Not Listed in the *interfaces* File

The entry for the local Backup Server is created when you install Adaptive Server. If you have installed a remote Backup Server on another machine, you can choose one of the following options to resolve this error:

- Copy the *interfaces* file entry for this Backup Server to the interfaces file on the machine where your Adaptive Server runs.
- Locate the *interfaces* file on a file system that is shared by both machines.
- Use sybinit (or the appropriate utility as listed earlier) to update your *interfaces* file on the machine where your Adaptive Server runs. Refer to the Adaptive Server configuration and installation guide for information about sybinit.

Additional information

For more additional information, refer to allow remote access configuration parameter in the *System Administration Guide* and "*interfaces* file" in the Adaptive Server installation and configuration guide.

Version in which this error is raised

All versions

Error 7221

Severity

14

Message text

```
Login to site '%S_SRVID' failed.
```

Explanation

This error occurs when you execute a remote procedure call (RPC) in one of the following situations:

- A name was not assigned, via sp_addserver, to the Adaptive Server which is initiating the RPC.
- A remote login for the Adaptive Server initiating the RPC was not added to the Adaptive Server that is receiving the RPC.
- The Adaptive Server receiving the RPC is not configured for allow remote access or for high enough values for number of remote sites or number of remote connections .
- The Adaptive Server receiving the RPC does not have an entry in the network addressing (*interfaces* or *sql.ini*) file for the Server initiating the RPC (or the other way around).
- The Adaptive Server receiving the RPC is in single-user mode and already had one user at the time the RPC was executed.
- The password for the user initiating the RPC is not the same on both Adaptive Servers.

Action

Eliminate the 7221 error by determining which of the above situations apply to you and following the appropriate steps:

1 A name was not assigned to the Adaptive Server which is initiating the RPC.

Use the select *@@servername* command to check if the Adaptive Server initiating the RPC has a name assigned to it.

If the select *@@servername* command returns the name of the server which is initiating the RPC, then continue your checking with step 2.

If the output returned by select *@@servername* is NULL, then assign a name to the local Adaptive Server via the sp_addserver command:

```
1> sp_addserver name_of_server_initiating_rpc,local
2> go
```

Refer to “sp_addserver” in the *Adaptive Server Reference Manual* for information about valid Adaptive Server names.

Restart Adaptive Server after assigning it a name.

Usually you want to name an Adaptive Server the same name as is stored in the *DSQUERY* environment variable or the name which is used in the -S option when the server is accessed via isql. For example:

```
% isql -SPRODUCTION
```

In the above example, PRODUCTION is the Adaptive Server name.

2 A remote login was not added to the Adaptive Server that is receiving the RPC for the Adaptive Server that is initiating it.

Use sp_helpremotelogin to check if the receiving Adaptive Server has a remote login for the Adaptive Server initiating the RPC. If the output returned by sp_helpremotelogin contains information about the remote Adaptive Server, then continue with step 3.

If the output returned by sp_helpremotelogin does not contain information about the remote Adaptive Server, then add a remote login to the remote Adaptive Server for the Adaptive Server that is initiating the RPC by typing the following (on the remote
Adaptive Server):

```
1> sp_addremotelogin name_of_server_initiating_rpc
2> go
```

3 The Adaptive Server receiving the RPC is not configured for allow remote access or for high enough values for number of remote sites or number of remote connections.

Check the remote Adaptive Server's error log for more information on why the RPC initiated by the local Adaptive Server did not succeed. Refer to "Setting Configuration Parameters" in the *System Administration Guide* for information about how to turn on allow remote access and how use sp_configure to raise the values of number of remote sites or number of remote connections.

4 The Adaptive Server receiving the RPC does not have an entry in the *interfaces* file for the Server initiating the RPC (or the other way around).

Add the required entries for the Adaptive Server initiating the RPC to the *interfaces* file of the receiving Server. Use one of the following utilities or copy the entries from the *interfaces* file of the initiating Server and make sure that they have a tab in front of each entry.:

- sqledit (for OS/2)
- sqledit or sybinit (for NT Adaptive Server)
- sybinit (for all other platforms)

5 The Adaptive Server receiving the RPC is in single-user mode and already has one user at the time the RPC is executed.

Wait until the receiving Adaptive Server is not being used any more and then execute the RPC or restart the receiving Adaptive Server in multi-user mode.

6 The password for the user initiating the RPC is not the same on both Servers.

Make user passwords the same for the remote Adaptive Server.

Additional information

For more information, refer to:

- sp_addremotelogin, sp_addserver, and sp_configure in the *Adaptive Server Reference Manual*
- Server Is Unnamed of this manual
- "allow remote access" in the *System Administration Guide*

Version in which this error is raised

All versions

Error 7223

Severity

16

Message text

```
Cannot %S_MSG site '%S_SRVID'.
```

Explanation

This error occurs when Adaptive Server fails to complete a read or send operation during a remote procedure call (RPC). The failure can occur in a number of states. Each of these states is described below.

Note Adaptive Server uses tokens to establish a dialog request with a client. There are several kinds of tokens to send different signals.

State	Meaning
1	Adaptive Server uses buffering for sending and receiving packets to and from remote connections. If an error occurs while Adaptive Server is trying to read the next packet from the remote connection while Adaptive Server is being run in debug mode, Error 7223 is raised with State 1. *%S_MSG* is replaced by "read from" and *%S_SRVID* is the remote Adaptive Server name.
2	If Adaptive Server confirms that no token was read from the remote connection, then Error 7223 is raised with State 2. *%S_MSG* is replaced by "read from" and *%S_SRVID* is the remote Adaptive Server name.

State	Meaning
3	If an error occurs while Adaptive Server is sending a token to a remote connection, Error 7223 is raised with State 3. *%S_MSG* is replaced by "send to" and *%S_SRVID* is the remote Adaptive Server name.
4	If an error occurs while Adaptive Server is sending data to a remote connection, Error 7223 is raised with State 4. *%S_MSG* is replaced by "send to" and *%S_SRVID* is the remote Adaptive Server name.
5	While Adaptive Server is establishing a network connection to a remote Adaptive Server, it translates the remote Server name to a Server ID, opens the logical connection, and then logs in. Error 7223 is raised with State 5 when Adaptive Server fails during one of these operations. *%S_MSG* is replaced by "send to" and *%S_SRVID* is the remote Adaptive Server name.
6	Adaptive Server builds a remote procedure call request message in the remote connection send buffer and then sends it. Error 7223 occurs with State 6 when this send fails. *%S_MSG* is replaced by "send to" and *%S_SRVID* is the remote Adaptive Server name.
7	If Adaptive Server confirms that no data was read from the remote connection, then Error 7223 is raised with State 7. *%S_MSG* is replaced by "read from" and *%S_SRVID* is the remote Adaptive Server name.

Action

1. Check the network connection between the locations of your local and remote Adaptive Servers. Make sure packets of data can be sent and received. You may need to ask your network administrator for help.
2. If you did not find a problem in step 1, call Sybase Technical Support.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR Rollup level
- Server error log for both the local and remote Adaptive Servers
- Text of all error messages
- sp_configure output

Version in which this error is raised

All versions

Error 7227

Severity

17

Message text

```
Can't open a connection to site '%S_SRVID' because
```

```
Adaptive Server is not configured for remote access. Ask
a user with System Administrator (SA) role to
reconfigure Adaptive Server to allow remote access.
```

Explanation

This error occurs when you execute a remote procedure call (RPC) and the sending or receiving server is not configured for remote access.

Action

Configure the local or remote Adaptive Server for remote access by changing the allow remote access configuration parameter via sp_configure.

```
1> sp_configure "allow remote access", 1
2> go
```

After enabling the allow remote access configuration parameter and restarting Adaptive Server, the number of remote logins, number of remote sites, number of remote connections, and remote server pre-read packets configuration parameters will be set to their default values.

If you determine that the default values for number of remote logins, number of remote sites, number of remote connections, and remote server pre-read packets are not adequate for your needs, use sp_configure to change these configuration parameters as well.

Restart Adaptive Server to make the configuration change effective.

Additional information

Refer to “allow remote access” in the *System Administration Guide* and “sp_configure” in the *Reference Manual*.

Version in which this error is raised

All versions

Error 7234

Severity

```
10
```

Message text

```
Failed to connect to remote server %s %s %s %s %s for
site %s. Trying the next query entry of this server, if
any.
```

Explanation

Communication between an Adaptive Server and another server requires a site handler to manage connections between the servers. When you log on to Server A, and issue a command that necessitates communication with Server B (such as Backup Server for dumps/loads), Server A uses a site handler to open a connection to Server B, sending the query and receiving results from Server B.

Error 7234 is raised when Adaptive Server is unable to open a connection to another server. The remote server information in the message text refers to:

- the server name
- the network type
- the network name
- the machine name, and
- the port number.

Processing continues after the error is raised, with the server re-trying the connection using the next entry, if any, in the *interfaces* file.

Error 7234 is an informational message. However, if all attempts to open a connection using the available *interfaces* entries fail, Adaptive Server raises Error 7235 ("No other query entry found for the remote server %s - site %s. Aborting this site handler, as no remote server responded. Please make sure the remote server is accepting connections.").

Action

- Check that the remote server is up and running. If you see the error after issuing a dump command, for example, check that Backup Server is running.
- Check the remote server's network name in sysservers:

```
1> select srvname, srvnetname from master..sysservers
2> go
```

  The remote server's srvnetname entry must match the name appearing in the *interfaces* file.

Version in which this error is raised

All versions

Error 7235

Severity

17

Message text

```
No other query entry found for the remote server %s -
site %s. Aborting this site handler, as no remote server
responded. Please make sure the remote server is
accepting connections.
```

Explanation

Communication between an Adaptive Server and another server requires a site handler to manage connections between the servers. When you log on to Server A, and issue a command that necessitates communication with Server B (such as Backup Server for dumps and loads), Server A uses a site handler to open a connection to Server B, sending the query and receiving results from Server B.

When the first attempt to open a site handler connection fails, Adaptive Server raises Error 7234 ("Failed to connect to remote server %s %s %s %s %s for site %s. Trying the next query entry of this server, if any."). The server then re-tries the connection using the next entry, if any, in the *interfaces* file. Error 7235 is raised after all such connection attempts fail.

Action

- Check that the remote server is up and running. If you see the error after issuing a dump command, for example, type the following:

```
1> exec SYB_BACKUP...sp_who
2> go
```

  If Backup Server is running, you will get a meaningful response to this query.

- Check the remote server's network name in sysservers:

```
1> select srvname, srvnetname from
master..sysservers
2> go
```

  The remote server's srvnetname entry must match the name appearing in the *interfaces* file. (An incorrect srvnetname can prevent the server from reaching the correct remote server.)

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level
- Server error log
- Text of all error messages.

Version in which this error is raised

All versions

Parser Errors (continued)

This section contains error messages for the Adaptive Server Parser.

Error 7364

Severity

16

Message text

```
Tables '%.*s' and '%.*s' have some exposed names. Use
correlation names to distinguish them.
```

Explanation

An exposed name is an unqualified object name which assumes the default owner and database. For example, table1 (exposed name) and master.dbo.table1 (qualified name) may *not* be the same table but
Adaptive Server will think you mean the same table if you reference just table1 as dbo in the master database. Explicit correlation names allow you to make distinctions between references in a query.

In ANSI 89, the following construction is not allowed:

```
select * from table1, table1 [where_clause]
```

Using this syntax returns Error 7364.

Action

The following syntax, which uses *t1* and *t2* as table correlation names, allows you to perform a self-join:

```
select * from table1 t1, table1 t2 [where_clause]
```

Version in which this error is raised

All versions

Error 7380

Severity

10

Message text

```
WARNING: Line %d. Invalid value for '%S_MSG' specified.
This value will be ignored, default value will be used
instead.
```

Explanation

You can specify a prefetch size on select, update, and delete statements for tables bound to caches with large I/O pools. Specify the prefetch size as the I/O size in kilobytes. Valid values for prefetch size are 2K, 4K, 8K, and 16K. The procedure sp_helpcache shows pools and their sizes for the cache to which an object is bound, or for the default cache.

	Error 7380 occurs when the prefetch size you specify is invalid. When this error occurs, Adaptive Server continues processing and the optimizer substitutes the “best” prefetch size available.
Action	Make sure the prefetch size you specify is 2K, 4K, 8K, or 16K. In addition, if the object is bound to a cache, make sure you use a valid prefetch size for that cache. Use sp_helpcache to see pools and their sizes for the cache to which the object is bound.
Additional information	Refer to the *Reference Manual* and the *Performance and Tuning Guide* for information about using prefetch sizes.
Version in which this error is raised	All versions

Threshold Errors

This section contains error messages pertaining to Adaptive Server thresholds. Refer to "Managing Free Space with Thresholds" in the *Adaptive Server Enterprise System Administration Guide* for information about thresholds.

Error 7401

Severity

10

Message text

```
Could not allocate memory for the threshold task's
argument list; threshold task cannot be spawned for
database %d, segment %d, free space %ld.
```

Explanation

This error occurs when a stored procedure associated with a threshold was not initiated because there was not enough memory available in procedure cache.

Action

Although Error 7401 indicates that the threshold task did not run because not enough memory was available in procedure cache at that time, the threshold task may have run later when enough memory did become available. Before performing the following steps, determine whether the threshold task ran at a later time. If the threshold task did run, the following steps are not necessary.

To perform the threshold task that did not run because of this error:

1 Select the database:

```
1> use database_name
2> go
```

where *database_name* is the database for which the error occurred.

2 Run sp_helpthreshold segment_name for the segment named in the error message. This will display the stored procedure associated with the threshold.

3 Run the stored procedure manually.

To prevent this error from occurring in the future, make sure there is enough memory to run stored procedures associated with a threshold by testing them during the time when your system is most heavily loaded. If there is not enough memory, do one of the following:

- Use sp_configure to increase the Adaptive Server procedure cache size configuration parameter (increases the percentage of total cache space that is used as procedure cache).

- Use sp_configure to increase the value of the total memory configuration parameter.
- Rewrite the stored procedure so that it will execute in the procedure cache available for the Adaptive Server.

Additional information

Refer to the *Performance and Tuning Guide* and "Configuring Memory" in the *System Administration Guide* for information about Adaptive Server memory use.

Refer to "Managing Free Space with Thresholds" in the *System Administration Guide* for information about Adaptive Server thresholds.

Version in which this error is raised

All versions

Error 7402

Severity

10

Message text

```
Could not spawn the threshold handler procedure for
database %d, segment %d, when only %ld free pages were
left in the segment.
```

Explanation

Thresholds access Adaptive Server as a detached user. This means that, although a threshold does not access Adaptive Server through the network, it still requires a user connection. Error 7402 occurs when there are no available user connections when an Adaptive Server process reaches a threshold and Adaptive Server tries to execute the associated threshold procedure.

Action

To perform the threshold task that did not run because of this error:

1 Select the database:

```
1> use database_name
2> go
```

where *database_name* is the database for which the error occurred.

2 Run sp_helpthreshold segment_name for the segment named in the error message. This will display the stored procedure associated with the threshold.

3 Run the stored procedure manually.

To prevent this error from occurring in the future, use sp_configure to make sure that there are enough user connections to allow for all clients and the threshold manager. Refer to "Setting Configuration Parameters" in the *Adaptive Server Enterprise System Administration Guide* for information about changing the number of user connections configuration parameter.

Additional information

Refer to "Managing Free Space with Thresholds" in the *System Administration Guide* for information about Adaptive Server thresholds.

Version in which this error is raised

All versions

Error 7403

Severity

```
10
```

Message text

```
Threshold task could not use database %d, and so cannot
execute the threshold procedure for segment %d, free
space %ld.
```

Explanation

A threshold is "owned" by the user who created it, or the last user to modify it with sp_modifythreshold (refer to sp_addthreshold and sp_modifythreshold in the *Reference Manual* for details). Error 7403 occurs when the Threshold Manager tries to use a database and execute a stored procedure and access is denied. The procedure is not executed.

Action

To perform the threshold task that did not run because of this error:

1 Select the database:

```
1> use database_name
2> go
```

where *database_name* is the database for which the error occurred.

2 Run sp_helpthreshold segment_name for the segment named in the error message. This will display the stored procedure associated with the threshold.

3 Run the stored procedure manually using the account of the user that owns it.

To determine the owner of the procedure and user access to the database, type:

```
1> sp_help procedure_name
2> go
```

```
1> sp_helpuser database_name
2> go
```

```
1> sp_helprotect procedure_name
2> go
```

To prevent this error from occurring in the future, make sure that the owner of each threshold has access to the database to which the threshold applies, and that the user has execute permission on the stored procedure associated with the threshold.

Version in which this error is raised

All versions

Error 7404

Severity

```
10
```

Message text

```
Could not find row in systhresholds for database %.*s
describing segment %d, free space %ld.
```

Explanation

This error can occur for either of the following reasons:

- The row in database.dbo.systhresholds describing the segment, free space, and procedure name for the segment named in the error message is incorrect. Therefore, the threshold manager was not able to execute the stored procedure associated with the threshold.
- The clustered index on systhresholds is corrupt and Adaptive Server was not able to retrieve the row.

Action

Run dbcc checktable (systhresholds) and dbcc tablealloc (systhresholds) to check consistency of the system table. If you have corruption on systhresholds, contact Sybase Technical Support for further instructions.

Perform the threshold task that did not run because of this error:

1 Select the database:

```
1> use database_name
2> go
```

where *database_name* is the database for which the error occurred.

2 Run sp_helpthreshold segment_name for the segment named in the error message. This will display the stored procedure associated with the threshold.

3 Run the stored procedure manually.

If the stored procedure is not found or there are incorrect entries for segment or free space, modify the row in systhresholds with sp_modifythreshold to add a stored procedure (threshold action). Refer to the *Reference Manual* for information about sp_modifythreshold.

Additional information

Refer to the material in this manual for errors reported by dbcc.

When calling Technical Support, have the following information available:

- Output of select @@*version*
- Output of all error messages
- Output of dbcc commands
- Output of sp_helpthreshold command

Version in which this error is raised

All versions

Error 7405

Severity

10

Message text

```
Threshold procedure name is missing from the
systhresholds row in database %.*s for segment %d, free
space %ld.
```

Explanation

This error occurs when a row is found in database.dbo.systhresholds for the specified segment, but no procedure name is available, and the threshold task is not able to take appropriate action.

Action

Run dbcc checktable (systhresholds) and dbcc tablealloc (systhresholds) to check consistency of the system table. If you have corruption on systhresholds, contact Sybase Technical Support for further instructions.

If there is no corruption, modify the row in systhresholds with sp_modifythreshold to add a stored procedure (threshold action). Refer to the *Reference Manual* for information about sp_modifythreshold.

Then perform the threshold task that did not run because of this error by running the stored procedure manually.

Additional information

Refer to the writeups in this chapter for errors reported by dbcc.

When calling Technical Support, have the following information available:

- Output of select @@*version*

- Output of all error messages
- Output of dbcc commands
- sp_helpthreshold output

Version in which this error is raised

All versions

Error 7406

Severity

16

Message text

```
%s: unknown option: %.*s
```

Explanation

lct_admin is a function used to manage the last-chance threshold (lct) on the log segment of a database. Error 7406 occurs when an invalid option is passed to lct_admin.

Action

Use one of the valid options listed below. Note that the options are all strings and require quotation marks.

- lct_admin ("reserve", number_of_log_pages) – returns the number of pages required to do a successful dump transaction given the number of pages in a log segment.
- lct_admin ("lastchance", database_ID) – creates a last-chance threshold for a database with log on a separate device. This can be used to create a last-chance threshold for databases that were upgraded from a pre-10.0 SQL server.

  Note You should not have to use the lastchance option of lct_admin.

- lct_admin ("logfull", database_ID) – determines if the last-chance threshold has been reached in a specified database. The procedure will return “0” if last-chance has not been reached, and “1” if it has.
- lct_admin ("unsuspend", database_ID) – unsuspends all tasks that were put to sleep by activation of last-chance threshold. Use this option with caution, as you may run out of space in the log and end up having to do a dump transaction with no_log and then you will not be able to do a dump transaction to a device. Always try to do a dump transaction to a device or a dump transaction with truncate_only before using the unsuspend option of lct_admin.

Additional information

Refer to the *Reference Manual* for a description of valid lct_admin options for your version.

Version in which this error is raised

All versions

Error 7407

Severity

10

Message text

```
Failed to execute the threshold procedure for database
%d, segment %d, free space %ld.
```

Explanation

This error occurs when an attempt to execute a stored procedure defined in *systhresholds* fails.

Action

To perform the threshold task that did not run because of this error:

1 Select the database:

```
1> use database_name
2> go
```

where *database_name* is the database for which the error occurred.

2 Run sp_helpthreshold segment_name for the segment named in the error message. This will display the stored procedure associated with the threshold.

3 Run the stored procedure manually.

To determine whether you can prevent this error from occurring in the future, review the Adaptive Server error log for additional messages regarding this error. If there is specific information as to what caused the procedure to fail, take whatever steps are necessary to prevent future occurrences.

Version in which this error is raised

All versions

Error 7408

Severity

20

Message text

```
Could not find a dbtable for database %d.
```

Explanation

Information about each database is kept in memory in a structure called a dbtable. Error 7408 occurs when an attempt to execute the built-in function **lct_admin** fails because the database table (database ID) information was either not available or incorrect.

Action

Verify that you have entered the correct database name (or database ID) in the threshold procedure. If it is not correct, modify the procedure.

This error can occur because the database is not open because the value of the number of open databases configuration parameter is too low. If the value for number of open databases is too low, use sp_configure to raise it.

If the database name is correct, the value of the number of open databases configuration parameter is large enough to accommodate all the system and user databases, and the command still fails with this message, check the Adaptive Server error log for messages.

To perform the threshold task that did not run because of this error:

1 Select the database:

```
1> use database_name
2> go
```

where *database_name* is the database for which the error occurred.

2 Run sp_helpthreshold segment_name for the segment for which the error occurred. This will display the stored procedure associated with the threshold.

3 Run the stored procedure manually.

Version in which this error is raised

All versions

Error 7409

Severity

20

Message text

```
This task tried to sleep waiting for space to become
available in the log, but could not.  The current
command is being aborted because no further database
changes can be logged.
```

Explanation

The last-chance threshold has been reached in a database. All transactions trying to write to the transaction log should be put to sleep until space is free in the transaction log. Adaptive Server was unable to put the transaction to sleep, so the transaction has been aborted.

Action

Verify that the threshold action defined for the last-chance threshold has executed successfully. If it did not complete, manually execute a dump transaction command to create free space in the log. Then rerun the transaction that was aborted. If the transaction still fails, check the Adaptive Server error log for additional information.

Refer to "Managing Free Space with Thresholds" in the *System Administration Guide* for information about last-chance thresholds.

Version in which this error is raised

All versions

Error 7410

Severity

26

Message text

```
%s: function called with invalid argument.
```

Explanation

This error occurs when an attempt to execute a threshold action fails because of a problem in a memory structure.

Action

If there are additional messages in the server error log that tell you which threshold was initiated, do the following steps:

1 Select the database:

```
1> use database_name
2> go
```

where *database_name* is the database which contains the failed threshold.

2 Determine the stored procedure that is associated with the segment for which the error occurred:

```
1> sp_helpthreshold segment_name
2> go
```

3 Run the stored procedure manually.

If you were not able to determine which threshold was initiated by looking at the error log output, run the following curunreservedpgs query in each database to determine which threshold action did not execute properly:

```
1> use database_name
```

```
2> go

1> select db_name(dbid), d.name,
2> curunreservedpgs (dbid, lstart, unreservedpgs)
3> from sysuages u, sysdevices d
4> where d.low <= u.size + vstart
5> and d.high >= u.size + vstart - 1
6> and d.status &2 = 2
7> go
```

This query returns the database name, device name, and the number of unreserved pages for each device fragment. If you can determine which threshold was reached (compare the output of the curunreservedpgs query with the free space for the segment), execute the threshold action manually. Refer to "Transact-SQL Functions" in the *Reference Manual* for information about curunreservedpgs.

To prevent this error from occurring in the future, review the Adaptive Server error log and look for other messages associated with this error. There may be other errors that will tell you why the threshold failed. If more information is available, take the necessary steps to prevent the error from recurring.

Version in which this error is raised

All versions

Error 7411

Severity

21

Message text

```
Table 'systhresholds' in database '%.*s' is not in its
correct sort order.  Either the clustered index is
missing or there is data corruption in the table.
```

Explanation

This error occurs when an attempt to execute a threshold action fails because of possible corruption on system table systhresholds or its clustered index.

Action

Run dbcc checktable (systhresholds) and dbcc tablealloc (systhresholds) to check consistency of the system table. If you find corruption in systhresholds, call Sybase Technical Support for further instructions.

Have the following information available when you call Sybase Technical Support:

- Output of select *@@version*
- Output of all error messages

- Output of dbcc commands

Version in which this error is raised: All versions

Error 7412

Severity: 10

Message text:

```
Space available in the log segment has fallen critically
low in database '%.*s'.  All future modifications to
this database will be %S_MSG until the log is
successfully dumped and space becomes available.
```

Explanation: This message notifies users that the last-chance threshold has been reached. All transactions requiring access to the transaction log will either be suspended or aborted until the appropriate action is taken to clear the transaction log.

Action:

1. Ensure that the action defined by the last-chance threshold procedure was able to execute successfully. If the procedure failed for any reason, manually run dump transaction to create free space in the transaction log.
2. Look for errors in the Adaptive Server error log to determine why the procedure failed.
3. Restart any transactions that have been aborted.

Additional information: Refer to “Managing Free Space with Thresholds” in the *System Administration Guide* for information about last-chance thresholds.

Version in which this error is raised: All versions

Error 7413

Severity: 10

Message text:

```
%d task(s) are sleeping waiting for space to become
available in the log segment for database %.*s.
```

Explanation: This message is written to the Adaptive Server error log by the checkpoint process when the last-chance threshold has been exceeded, and the specified number of processes have been put to sleep.

Action	Ensure that the appropriate last-chance threshold action is executed to create free space in the log. If the procedure did not run for any reason, manually create free space in the transaction log with a dump transaction command and check on the cause of the failure by referring to the documentation in this manual.
Additional information	Refer to “Managing Free Space with Thresholds” in the *Adaptive Server Enterprise System Administration Guide* for information about last-chance thresholds.
Version in which this error is raised	All versions

Error 7414

Severity	10
Message text	`Threshold procedure '%.*s' returned an error indication (@status=%d) when called for database '%.*s', segment '%.*s', free_space %ld.`
Explanation	This error occurs when a threshold procedure is activated but is not able to execute successfully.

Action	Check the Adaptive Server error log for additional information on the reason for the failure. Correct the situation that caused the failure.
	
	Execute the procedure that failed:
	
	1 Select the database:
	
	`1> use database_name`
	`2> go`
	
	where *database_name* is the database which contains the failed threshold.
	
	2 Determine the stored procedure that is associated with the segment for which the error occurred:
	
	`1> sp_helpthreshold segment_name`
	`2> go`
	
	3 Run the stored procedure manually.

Version in which this error is raised	All versions

Error 7415

Severity

10

Message text

```
The transaction log in database %.*s is almost full.
Your transaction is being %S_MSG until space is made
available in the log.
```

Explanation

This error occurs when the last-chance threshold has been exceeded for the specified database. Each transaction accessing the transaction log will receive this message, indicating whether the transaction will be put to sleep or aborted until log space is available.

Action

This message does not ordinarily require action if the transaction was put to sleep and the threshold procedure frees log space. In that circumstance, the sleeping processes will awaken and continue. If the process was aborted, it will be necessary to restart it. If the process is not able to continue, determine why it cannot continue by looking at the error log and manually dump the transaction log.

Version in which this error is raised

All versions

Auditing Errors

This section contains error messages pertaining to Adaptive Server auditing. Refer to the *Security Administration Guide* for more information about auditing.

Error 7618

Severity

20

Message text

```
The audit process is not currently active. Contact a
user with System Administrator (SA) or System Security
Officer (SSO) role.
```

Explanation

When auditing is required, an audit record is allocated on the stack and filled in with information which varies depending upon the type of event. Adaptive Server formats the record into an audit buffer (the audit queue) and places this buffer in the audit process mailbox.

If the audit process is not active when auditing is required, Adaptive Server terminates the user's process because it should have audited whatever was happening but could not do so. A message is sent to the error log indicating that an audit has been missed, and Error 7618 is raised.

Although there can be a number of reasons why the audit process is not active, the most common reason is that the sysaudits table has run out of space and the audit device is full. This can happen if you do not regularly archive your audit data.

Action

If the audit process died abnormally, there will be audit error messages in the Adaptive Server error log. Look for audit errors in your error log. If it looks like the audit process died abnormally, shut down and restart Adaptive Server, and then run sp_who to determine whether the audit process is active.

If it looks like the audit process did not die abnormally, determine whether the sybsecurity database is full (if that is the case, you will probably see 1105 errors in the error log):

```
1> sp_spaceused sybsecurity
2> go
```

If the sybsecurity database is full, the System Security Officer should do the following:

1 Archive and truncate the sysaudits table.

2 Shut down Adaptive Server.

3 Have the System Administrator restart Adaptive Server.

If the sybsecurity database is not full and the audit process did not die abnormally or does not restart when you restart Adaptive Server, call Sybase Technical Support.

Additional information

Refer to the *Adaptive Server Enterprise Security Administration Guide* for information about auditing, including how to archive audit data.

Version in which this error is raised

All versions

Sequencer Errors (continued)

This section includes error messages for the Adaptive Server Sequencer, which is the subsystem that controls command execution. The Sequencer handles the processing between steps in commands and also handles the procedural steps between commands, such as if, goto and while.

Error 7783

Severity

10

Message text

```
Internal error -- server failed to do garbage collection
on this procedure, id = %ld.  Please save the query tree
and procedure text and inform the technical support.
This is not a critical error, so server continues
processing.
```

Explanation

Remapping a stored procedure, view or trigger generally causes the object to grow. After remapping, Adaptive Server compresses the object. Error 7783 occurs when Adaptive Server is unable to reclaim space after remapping the query tree. It occurs:

- during stored procedure or trigger execution.
- during recovery, usually following a load database.

This error is due to an Adaptive Server problem and is frequently accompanied by Error 7961 ("Remapping utility - Wrong kind of node passed to tree remapping. This is an internal system error.")

Although the server is unable to reclaim space, error 7783 does not cause any severe problems and the procedure or trigger continues processing.

Action

Take the following steps:

1 Identify the stored procedure, view or trigger:

```
1> select object_name(object_id)
2> go
```

where *object_id* is the id named in the error message.

2 Check that the object text is available. It is important to verify this before you continue with the next step.

3 Drop and re-create the object. Note that you are re-creating the stored procedure, view or trigger, and no changes are needed to the underlying table(s).

4 If the error results in a stack trace and/or is accompanied by the 7961 error, you may need to upgrade to a version in which the problem is resolved. Call Sybase Technical Support for assistance.

Additional information

Have the following information available before calling Sybase Technical Support:

- Server version and SWR version level
- Server error log
- Text of all error messages
- sp_configure, sp_who, and dbcc lock output

Version in which this error is raised

All versions

Error 7788

Severity

10

Message text

```
WARNING: An explicit or implicit HOLDLOCK or NOHOLDLOCK
was specified for table '%.*s'; however, since this is
an isolation level 0 scan, the HOLDLOCK or NOHOLDLOCK
is ignored.
```

Explanation

Isolation level 0 allows transactions to read uncommitted data (*dirty reads*). You can specify isolation level 0 for your queries using the at isolation syntax, or you can specify level 0 for your session as part of the transaction isolation level option of the set command. Adaptive Server's default transaction isolation level is 1.

The holdlock parameter makes a shared lock on a specified table or view more restrictive by holding it until the completion of a transaction (instead of releasing the lock as soon as the required data is no longer needed, whether or not the transaction has been completed). Setting transaction isolation level 3 with a set command implicitly applies a holdlock.

The noholdlock parameter prevents Adaptive Server from holding any locks acquired, regardless of the transaction level currently in effect.

Error 7788 is a warning message that is raised during query execution when level 0 scans are active and the query includes an explicit or implicit **holdlock** or noholdlock parameter. The query is executed and the transaction isolation level 0 option of the set command takes precedence over the holdlock or noholdlock parameter of the query.

Action

If you want to use holdlock or noholdlock in your query, use the transaction isolation level option of the set command to change the isolation level to a value other than 0 for your session:

```
1> set transaction isolation level n
2> go
```

where *n* is the value to which you want to set the isolation level.

Then re-enter your query.

Additional information

Refer to the *Transact-SQL User's Guide* for information about selecting an isolation mode for a query or for your session.

Refer to the *Reference Manual* and the *Performance and Tuning Guide* for information about holdlock and noholdlock.

Version in which this error is raised

All versions

dbcc Errors (continued)

This section contains error messages for dbcc (database consistency checker) commands.

Messages that ordinarily have severities greater than 16 will show severity 16 when they are raised by dbcc checktable and dbcc checkalloc so that checks will continue with the next object.

Error 7901

Severity

16

Message text

```
Page %ld was expected to be the first page of a
TEXT/IMAGE value.
```

Explanation

Adaptive Server stores text and image data in a linked list of pages separate from the rest of the table. Error 7901 is raised when a dbcc utility is processing a text/image page chain, but finds that the first page in this chain is not a text or image page. The error is due to a corrupt page header in the first page, or a bad text pointer in a row in the table.

Error 7901 is raised with the following states:

State	Meaning
1	dbcc checktable detected the first-page error while checking a table's text/image chains.
2	dbcc fix_text detected the first-page error while recalculating the statistics for a table's text pages.
3	dbcc rebuild_text detected the first-page error while rebuilding the internal data structures for text/image data. (This state is raised in version 12.0 and higher).

Action

Take the steps described below to resolve the error.

1 Identify the row(s) with the bad data using the following command:

```
1> select <unique_col> from table_name
2> where convert (int, textptr(<text_colname>))
3> = <text_page>
4> go
```

where *unique_col* is any column or columns in your table which will uniquely identify the rows, *text_colname* is the text column, and *text_page* is the page ID in the message text.

2 Drop the row(s) in question.

3 Run dbcc checktable on the table to verify that the problem has been corrected.

As an alternative to dropping row(s), you can create a new table using select into, specifying all rows from the original table except the affected row(s) identified in Step 1. Refer to "How to Rescue Data from a Corrupted Table"in Chapter 2 for details.

Version in which this error is raised

All versions

Error 7902

Severity

16

Message text

```
Data size mismatch occurred while checking TEXT/IMAGE
values. The first page for this value is: %ld. There
were %ld bytes found, which is different from the
expected data length of %ld bytes.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server detects a mismatch between the expected length of the text/image data of a text column and its actual length. The length of the text/image data is stored on the page header of the text/image data page.

This error usually occurs in one of the following situations:

- When you run dbcc checktable on a table whose text or image column was updated to NULL and then updated again to a non-NULL value via a writetext command which did not commit.
- When you run dbcc checktable or dbcc checkdb and your image or text data is corrupted. If this is the case, check your hardware error log and your operating system error log and determine the cause of the 7902 error.

In order to eliminate the 7902 error, perform the following steps for each table encountering it:

1 Record the page numbers reported by Error 7902 for your table.

2 Select the offending rows from your table in a separate temporary table, giving the name of the text column for the textptr function:

```
1> select * into #newtable from mytable
2> where convert (int, textptr(text_colname))
3> in (list_of_page#'s)
4> go
```

The list of page numbers in the above query is the enumeration of all the page numbers specified by the 7902 errors. text_colname is the name of the text column.

3 Update the corrupted text/image data from the old table with the text/image data from the new table:

```
1> begin transaction
2> go
1> update mytable
2> set mytable.text_col = #newtable.text_col
3> from #newtable
4> where mytable.other_col = #newtable.other_col
5> go
```

Make sure that only the required number of rows are updated by the above query and then commit the transaction. other_col in the above query is the name of another column in your table which uniquely identifies a row (like a primary key). If you do not have any primary keys on your old table, run the above query with several search conditions in the where clause:

```
1> update mytable
2> set mytable.text_col = #newtable.text_col
3> from #newtable
4> where mytable.col1 = #newtable.col1
5> and mytable.col2 = #newtable.col2
6> go
```

4 Make sure that only the required number of rows are updated by the above query and then commit the transaction:

```
1> commit transaction
2> go
```

Example

You will get the following messages during dbcc checktable:

```
Checking mytable
Msg 7902, Level 16, State 1:
Server 'PRODUCTION', Line 1:
Data size mismatch occurred while checking
```

```
TEXT/IMAGE values. The first page for this value
is: 321. There were 12 bytes found, which is
different from the expected data length of 0 bytes.

Msg 7902, Level 16, State 1:

Server 'PRODUCTION', Line 1:

Data size mismatch occurred while checking
 TEXT/IMAGE values. The first page for this value
is: 291. There were 17491 bytes found, which is
different from the expected data length
 of 19291 bytes.
```

Run the following query in order to select the erroneous rows from mytable into #newtable:

```
1> select * into #newtable from mytable
2> where convert (int, textptr(text_colname))
3> in (321, 291)
4> go
```

Another Method to Eliminate Error 7902

An alternate method is to select all data from the old table into a new table.

If the index ID is 0 or 255 and you do not have clean backups of your database, copy the data from the corrupted table into a new (dummy) table or into a file. Then rename your old, corrupted table and copy the data back into a new table using the original name. For more information about doing this, refer to Chapter 2, "How to Rescue Data from a Corrupted Table".

Warning! Some data on this page might be lost if you recover your table using bcp or select into (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

Before dumping your database, make sure it is working correctly. More specifically, the following commands should be run prior to each dump:

1 dbcc checkdb.

2 dbcc checkalloc or dbcc checkalloc with the fix option. Refer to "dbcc" in the *Reference Manual* for information about dbcc checkalloc.

Prevention

Consider using writetext with log or update instead of writetext with no log. This way, you will not have unlogged changes to your database and will not get 7902 errors if writetext with no log did not commit.

Additional information

For more information, refer to select into, convert, and writetext in the *Adaptive Server Reference Manual* and in the *Transact-SQL User's Guide.*

Version in which this error is raised

All versions

Error 7928

Severity

16

Message text

```
Index %.*s is not consistent; found %ld leaf rows but
%ld data rows.  Drop and recreate the index.
```

Explanation

This error is returned by dbcc checktable or dbcc checkdb (which calls dbcc checktable). When checking the integrity of a nonclustered index, checktable compares the leaf row count (the total number of rows in the leaf pages for the index) with the count of data rows. Error 7928 is raised when the leaf row count does not match the data row count, and indicates an index problem.

Action

Correct the problem by re-creating the index as follows:

1. Check the output of the dbcc query to identify the table associated with the index.
2. If the table is a system table, refer to "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index. Then go to Step 4.
3. If the table is a user table, take the following steps:
   - Run sp_helpindex on the index to ensure that the information needed to re-create the index is available.
   - Drop the index.
   - Re-create the index.
4. Run dbcc checktable on the table to verify that the corruption is gone. If corruption still exists, call Sybase Technical Support.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server version and SWR rollup level

- Text of all error messages
- All transaction logs dating back to the last clean dbcc run.

Version in which this error is raised

All versions

Error 7930

Severity

16

Message text

Version 11.0.2 and Later

```
Table Corrupt: key(s) in the left child is not less than
the parent key; check left child page %ld of index id
%d for table '%.*s' in database '%.*s'.
```

Version 11.0.1 and Earlier

```
Table Corrupt: keys in left child is not less than the
parent key; check left child page %ld.
```

Explanation

This serious error is returned by dbcc checktable or dbcc checkdb (which calls dbcc checktable). dbcc checktable checks a table and its indexes to determine whether:

- Pages are linked correctly
- Indexes are properly sorted
- Pointers are consistent
- Page offsets are consistent
- sysindexes pages are correct
- The data row count is equal to the leaf row count for nonclustered indexes

During index level checking, for each key in the parent page, dbcc checktable checks whether the last key in the left child page is less than the parent key, and whether the first key in the right child page is greater than or equal to the parent key. If the last key in the left child is not less than the parent key, Error 7930 occurs.

Action

1 For Version 11.0.1 and earlier, use the procedure in Chapter 2, "How to Find an Object Name from a Page Number" to identify which table and index correspond to the page number from the error message text.

2 If the object encountering the error is *not* a system table (a system table's object ID is less than 100), continue with step 3.

   If the object with the error is a system table and the index ID is *not* 0, refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index. Then go to step 5.

   If the object with the error is a system table and the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption but it may be necessary to restore from clean backups.

3 If the object with the error is a user table and the index ID is *not* 0, determine whether there is a clustered index on the table:

```
1> sp_helpindex table_name
2> go
```

   where *table_name* is the name you determined in step 1.

   If there is a clustered index on the table, rebuild it. Then go to step 5.

4 If there is a nonclustered index on the table, rebuild the index using the following steps.

   Translate the index ID into an index name:

```
1> use database_name
2> go
```

```
1> select name from sysindexes
2> where id = object_ID and indid = index_ID
3> go
```

   To ensure that the information needed to re-create the index is available, run sp_helpindex on the index prior to dropping it.

   Drop the index.

   Re-create the index. This clears the corruption in most cases.

5 Run dbcc checktable on the table to verify that the corruption is gone. If corruption still exists, call Sybase Technical Support.

Additional information

Refer to drop index and create index in the *Reference Manual* for information about dropping and re-creating indexes.

Version in which this error is raised

All versions

Error 7939

Severity 16

Message text

```
Table Corrupt: The entry is missing from the OAM for
object id %ld indid %d for allocation page %ld.
```

Explanation

This error is raised when Adaptive Server detects that the allocation pages associated with an object are not recorded in the object allocation map (OAM) for the object. This problem is detected by dbcc checkalloc, dbcc tablealloc, and dbcc indexalloc.

Usually, Error 7939 will not cause any operations to fail at run time, so it is acceptable to wait to correct this problem until non-peak hours. However, do not drop a table with the 7939 error, as subsequent dbcc checkalloc commands may produce 2540 and 2546 errors.

Note The instructions below are for fixing 7939 errors once they have occurred. Two easy-to-use strategies exist for detecting this error sooner in the future. Refer to Chapter 2, "Detecting Allocation Errors as Early as Possible" for information about these strategies.

Action

Occasionally dbcc checkalloc reports this error when no real error condition exists. You can either check to see if these errors are real, or continue with this section and take action to correct them, whether or not they reflect a real allocation error.

Because the process used to discover whether or not the errors are real can be time-consuming, you may want to go directly to the "Error Resolution" section now.

Verifying That the Error Is Real

Run dbcc checkalloc in single-user mode if you suspect the 7939 error messages are incorrect. If the error is in master, use the section Chapter 2, "How to Start Adaptive Server in Single-User Mode" to start Adaptive Server in single-user mode. Refer to "Checking Database Consistency" in the *System Administration Guide* for information about dbcc checkalloc.

Error Resolution

If many of these errors are occurring, it is possible to clear them all at once by using the dbcc checkalloc with fix option command. Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about using dbcc checkalloc.

Make Sure Object Exists

Execute the following query to make sure the object exists and is correctly referred to in the system catalog:

```
1> use database_name
2> go

1> select object_name ( objid_from_error_msg)
2> go
```

If an object name is returned, then the error is on an existing, correctly referenced object. If this is the case, use the procedure in Chapter 2, "Fixing Allocation Errors when Object ID and Index ID are Known".

If a number, or something other than an object name is returned, use the procedure in Chapter 2, "How to Fix and Prevent Allocation Errors".

Version in which this error is raised

All versions

Error 7940

Severity

22

Message text

```
The counts in the OAM are incorrect.  This implies that
there are entries missing.  Run tablealloc utility with
the FIX option on the table with the inaccurate OAM
counts.
```

Explanation

This error is raised when Adaptive Server detects that the total page count in the object allocation map (OAM) is different from the actual number of pages in the chain. This problem is detected by dbcc checkalloc, dbcc tablealloc, and dbcc indexalloc.

This error can result in various run-time failures and should be corrected; however, if no other errors are occurring, you can wait to correct the problem until non-peak hours.

Note The instructions below are for fixing 7940 errors once they have occurred. Two easy-to-use strategies exist for detecting this error sooner in the future. Refer to Chapter 2, "Detecting Allocation Errors as Early as Possible" for information about these strategies.

Action

Occasionally dbcc checkalloc reports this error when no real error condition exists. You can either check to see if these errors are real, or continue with this section and take action to correct them, whether or not they reflect a real allocation error.

Because the process used to discover whether or not the errors are real can be time-consuming, you may want to go directly to the "Error Resolution" section now.

Verifying That the Error Is Real

Run dbcc checkalloc in single-user mode if you suspect the 7940 error messages are incorrect. If the error is in master, use the section Chapter 2, "How to Start Adaptive Server in Single-User Mode" for instructions about how to invoke Adaptive Server in single-user mode. Refer to "Checking Database Consistency" in the *System Administration Guide* for information about dbcc checkalloc.

Error Resolution

If many of these errors are occurring, it is possible to clear them all at once by using the dbcc checkalloc with fix option command. Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about using dbcc checkalloc.

Identify Table: User or System Table

Get the table name that the error occurred on from the original output which indicated this error. If it is a user table, continue with "Action for User Tables". If it is a system table, it requires a different procedure described in the section "Action for System Tables".

Action for User Tables

Run dbcc tablealloc. This command can correct the error only when run in the full or optimized mode, and with the nofix option not specified (the default for user tables):

```
1> dbcc tablealloc (table_name)
2> go
```

Refer to "dbcc" in the *Reference Manual* and "Checking Database Consistency" in the *System Administration Guide* for information about dbcc tablealloc.

Action for System Tables

1 Put the affected database in single-user mode:

- If the database is master, use the procedure in Chapter 2, "How to Start Adaptive Server in Single-User Mode", and then go to step 2.
- If the database is not master, use the sp_dboption stored procedure to put the affected database in single-user mode:

```
1> use master
2> go
1> sp_dboption database_name, single, true
2> go

1> use database_name
2> go

1> checkpoint
2> go
```

2 Run dbcc tablealloc. This command corrects the error only when run in the full or optimized mode, with the fix option specified, because the default value is nofix on system tables:

```
1> dbcc tablealloc (table_name, full, fix)
2> go
```

3 Turn off single-user mode in the database:

- If the database is master, use Chapter 2, "Returning Adaptive Server to Multiuser Mode".
- If the database is not master, use the following procedure:

```
1> use master
2> go
1> sp_dboption database_name, single, false
2> go

1> use database_name
2> go

1> checkpoint
2> go
```

Refer to "dbcc" in the *Reference Manual* and "Checking Database Consistency" in the *System Administration Guide* for information about dbcc tablealloc.

Version in which this error is raised

All versions

Error 7948

Severity	16
Message text	`Page %ld, object %ld, index %d, level %d, was found with fixed-row status improperly set. That status has been cleared by DBCC.`
Explanation	This informational message is displayed by dbcc checktable or dbcc checkdb (which calls dbcc checktable) when a page status bit in a page header is set on when it should not be. This bit is normally set on when there are no variable-length columns for this object (the rows are all fixed-length). If dbcc checktable finds that this bit is set on and there are variable display rows, the bit is turned off and the 7948 message is displayed.
Action	dbcc checktable (or dbcc checkdb) corrects this problem so no further action is required.
Version in which this error is raised	All versions

Error 7949

Severity	16
Message text	`The number of pages used and unused for %ld index %d on allocation page %ld do not match the counts in the OAM entry.`
Explanation	This error occurs when Adaptive Server detects that the allocation page count in the Object Allocation Map (OAM) is different from the actual number of pages in the chain. This problem is detected by dbcc checkalloc, dbcc **tablealloc**, and dbcc indexalloc. The object ID is the value immediately before the index value in the error text. Usually this error will not cause any operations to fail at run time. Therefore, it is acceptable to wait to correct this problem until non-peak hours.

Note The instructions below are for fixing 7949 errors once they have occurred. Two easy-to-use strategies exist for detecting this error sooner in the future. Refer to Chapter 2, "Detecting Allocation Errors as Early as Possible" for information about these strategies.

Action

Occasionally dbcc checkalloc reports this error when no real error condition exists. You can either check to see if these errors are real, or continue with this section and take action to correct them, whether or not they reflect a real allocation error.

Because the process used to discover whether or not the errors are real can be time-consuming, you may want to go directly to the "Error Resolution" section now.

Verifying That the Error Is Real

Run dbcc checkalloc in single-user mode if you suspect the 7949 error messages are incorrect. If the error is in master, use the section Chapter 2, "How to Start Adaptive Server in Single-User Mode" for instructions about how to invoke Adaptive Server in single-user mode. Refer to "Checking Database Consistency" in the *System Administration Guide* for information about dbcc checkalloc.

Error Resolution

If many of these errors are occurring, it is possible to clear them all at once by using the dbcc checkalloc command with the fix option. Refer to Chapter 2, "How to Fix and Prevent Allocation Errors" for information about using dbcc checkalloc.

Make Sure Object Exists

Execute the following query to make sure the object exists and is correctly referred to in the system catalog:

```
1> use database_name
2> go

1> select object_name ( objid_from_error_msg)
2> go
```

If an object name is returned, then the error is on an existing, correctly referenced object. If this is the case, use the procedure in Chapter 2, "Fixing Allocation Errors when Object ID and Index ID are Known".

If a number, or something other than an object name is returned, use the procedure in Chapter 2, "How to Fix and Prevent Allocation Errors".

Version in which this error is raised

All versions

Error 7989

Severity	16
Message text	`The serial allocation flag was found to be improperly set in allocation unit %ld, at extent ID %ld, with alloc map = %d and objid = %ld.  This flag has been cleared by DBCC.`
Explanation	When you initialize a database device, the space is divided into allocation units of 256 2K pages each. When an object requires space, Adaptive Server allocates a block of 8 2K pages within the unit, known as an extent, to the object. The first page of each allocation unit is an allocation page, which stores information about each extent contained in the allocation unit. The allocation page also contains a serial allocation flag, which indicates whether extents in the unit are allocated in serial, unbroken order. This flag enables more efficient processing: dumps and loads, for example, can process a set of transaction log pages as a group without having to traverse individual log pages. dbcc checkalloc examines each allocation page for consistency. When it finds an allocation page with serial allocation turned on, checkalloc verifies that all extents in that allocation unit are indeed allocated serially. If this is not the case, checkalloc raises Error 7989.
Action	Error 7989 is informational. No action is needed. checkalloc resets the flag to indicate that the extents are not serially allocated.
Version in which this error is raised	All versions

RPC Errors

This section contains error messages for Adaptive Server Remote Procedure Calls (RPC).

Error 8006

Severity 20

Message text

```
SQL Server error detected by %S_MSG.  Refer to the
preceding %S_MSG message.
```

Explanation

Error Message 8006 is a general purpose message that Adaptive Server can raise in a number of situations when working with another Sybase server. The '%S_MSG' in the message text is the name of the server that generated the error.

Most commonly, this error indicates that Backup Server encountered a problem in running the dump database or load database commands. When you issue a dump or load command from Adaptive Server, the server interprets the command and sends remote procedure calls (RPCs) to the Backup Server. As dumps and loads execute, Adaptive Server and Backup Server use RPCs to exchange instructions and status messages. Thus, an 8006 error during a dump or load is typically accompanied by additional messages originating from the Backup Server.

Action

Note the date-time stamp of the 8006 error in the Adaptive Server error log. Check the other server's error log (for example, the Backup Server's error log) for error messages corresponding to that time.

If there are any Backup Server errors, refer to Chapter 1, "Backup Server Error Messages," and look up the error description for assistance. Here are some common errors:

- 3.17.2 means that Adaptive Server was started from a directory other than its default working directory. See the writeup for Error 3240 in this manual for details.
- 4.76.3 means that the actual master device size does not match the value in *sysdevices.high*. See the writeup for Error 5013in this manual for a description of how to reset the high value.
- 4.82.2 means that there was an error in reading from or writing to a dump device. Check device and directory permissions.

If no other errors accompanied the 8006 message, there may be a problem with Backup Server. Take the following steps:

1 Use the showserver utility to check that Backup Server is running. Refer to the Adaptive Server utility commands manual for your platform for assistance.

2 Log into Adaptive Server and verify that you can communicate with Backup Server, using the following command:

```
1> exec SYB_BACKUP...sp_who
2> go
```

Under normal situations, you should see a few lines of sp_who output.

Version in which this error is raised

All versions.

Error 8009

Severity

16

Message text

```
Error encountered by %S_MSG.  Please refer to %S_MSG
messages for details.
```

Explanation

Error 8009 is a general purpose message that Adaptive Server can raise in a number of situations when working with another Sybase server. The `%S_MSG' in the message text is the name of the server that generated the error. Most commonly, this error indicates that Backup Server encountered a problem in running the dump database or load database commands.

When you issue a dump or load command from Adaptive Server, the server interprets the command and sends remote procedure calls (RPCs) to the Backup Server. As dumps and loads execute, Adaptive Server and Backup Server use RPCs to exchange instructions and status messages. Thus, an 8009 error during a dump or load is typically accompanied by additional messages originating from the Backup Server.

Action

Check the Backup Server and Adaptive Server error logs for additional messages. If there are any Backup Server errors, refer to Chapter 4, Backup Server Error Messages, for assistance in resolving the error.

If no other errors accompanied the 8009 message, there may be a problem with Backup Server. Take the following steps:

1 Use the showserver utility to check that Backup Server is running. Refer to the Adaptive Server utility commands manual for your platform for assistance.

2 Log into Adaptive Server and verify that you can communicate with Backup Server, using the following command:

```
1> exec SYB_BACKUP...sp_who
2> go
```

Under normal situations, you should see a few lines of sp_who output.

If Backup Server is down or is not communicating, restart the server and repeat the above steps. If the problem persists, contact Sybase Technical Support.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Adaptive Server version and SWR rollup level
- Adaptive Server and Backup Server error logs
- Text of all error messages.

Version in which this error is raised

All versions

Descriptor Manager Errors

This section contains error messages for the Adaptive Server Descriptor Manager.

Error 8201

Severity

26

Message text

Version 11.9.2

```
Keep count of descriptor (objid=%ld, dbid=%d) was
expected to be %d. Instead %d was found.
```

Version 11.5 and Earlier

```
Keep count of descriptor was expected to be %d. Instead
%d was found
```

Explanation

Session descriptors are the data structures used to manage access to Adaptive Server objects. The Descriptor Manager controls and synchronizes access to descriptors. It uses a keep count to determine whether a descriptor is in use. This guarantees that a descriptor is protected while it is in use. Error 8201 occurs if the Descriptor Manager determines that the keep count is incorrect.

Error 8201 is caused by an Adaptive Server problem.

Action

Shut down and restart Adaptive Server. This should clear Error 8201.

If the 8201 error persists, contact Sybase Technical Support.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server Version and SWR rollup level
- Server error log
- Text of all error messages

Version in which this error is raised

All versions

Error 8203

Severity

26

Message text

```
Expected to find the descriptor for object %d in
database %d in %s state.
```

Explanation

Session descriptors are the data structures used to manage access to Adaptive Server objects. As it is used, a descriptor passes through a number of different states:

- *free* – The descriptor is on the free resource list. It can be used only by the Descriptor Manager which can turn it into an active descriptor.
- *create* – This is the stage of a descriptor between free and active, while the fields are being filled in.
- *active* – The descriptor is accessible from outside the Descriptor Manager. This state does not necessarily mean that the descriptor is being used.
- *drop* – An object is in the process of being dropped.
- *destroy* – A descriptor is being versioned so it can be used for another object.
- *clean* – The descriptor is between the active and free states of the chain; in this state the descriptor cannot be found. The next state after this is either free or create.

When the Adaptive Server Descriptor Manager tries to access a descriptor and the descriptor is in an incorrect state, Error 8203 occurs. This error is caused by an Adaptive Server problem.

Error 8203 occurs with the following states:

State	Meaning
1	Adaptive Server expected a descriptor it was trying to get to be in state “active,” but it was not.
2	When Adaptive Server prepared to clean up a descriptor, it expected the descriptor to be in state “destroy,” but it was not.
3	When Adaptive Server prepared to return a descriptor from the resource free list, it expected the descriptor to be in state “free,” but it was not.
4	Adaptive Server prepared to move a descriptor from state “drop” to state “destroy,” but the descriptor was not in state “drop.”
5	Adaptive Server expected a descriptor it was creating for use to be in state “create, ” but it was not.

State	Meaning
6	Adaptive Server expected a descriptor it was creating for use to be in state “create.” However, Adaptive Server received an invalid status and could not determine what state the descriptor was in. Therefore, even though it expected the descriptor to be in state “create,” the message shows it expected the state to be “Invalid Stat Return.”
7	When Adaptive Server prepared to move a descriptor to state “free,” it expected the descriptor to be in state “clean,” but it was not.
8	When Adaptive Server needed to mark a descriptor as “active” and it had previously marked it for “drop,” it expected the descriptor to be in state “drop” and it was not.
9	When Adaptive Server prepared to move a descriptor from “active” to “drop”, it checked to make sure the descriptor was in state “active,” but it was not.
10	When Adaptive Server prepared to store information about a new descriptor, it expected the descriptor to be in state “create,” but it was not.
11	When Adaptive Server tried to initialize the lock promotion information in the descriptor, it expected the descriptor to be in state “create,” but it was not. This state is only raised by diagserver.

Action

If Error 8203 occurs with a state other than 6, shut down and restart Adaptive Server to clear the cache corruption. If the 8203 error occurs again, call Sybase Technical Support.

If Error 8203 occurs with State 6, call Sybase Technical Support.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server Version and SWR rollup level
- Server error log
- Text of all error messages

Version in which this error is raised

All versions

Error 8204

Severity 20

Message text Expected to receive system catalog id. Instead %d was

```
received.
```

Explanation

Session descriptors are the data structures used to manage access to Adaptive Server objects. Adaptive Server tracks individual descriptors by database table and system catalog ID. Error 8204 occurs if an invalid system catalog ID is returned while Adaptive Server is trying to locate a descriptor.

Error 8204 is caused by incorrect system catalogs or an Adaptive Server problem.

Action

Call Sybase Technical Support.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server Version and SWR rollup level
- Server error log
- Text of all error messages
- Output of dbcc checkcatalog if you know which database had the error

Version in which this error is raised

All versions

Error 8207

Severity

```
17
```

Message text

```
Couldn't find an available descriptor for an object.
Configure your system for more descriptors.
```

Explanation

This message is displayed when Adaptive Server is configured for fewer open objects (also known as descriptors) than are currently needed.

The configuration parameter number of open objects sets the maximum number of objects that can be open at one time in Adaptive Server.

Action

1 Use the system procedure sp_configure in order to display the run value of the configuration parameter you need to change:

```
1> sp_configure "number of open objects"
2> go
```

Increasing the value of number of open objects does not have a significant impact on performance or storage requirements. However, be careful when changing Adaptive Server configuration parameters. Refer to "Setting Configuration Parameters" in the *Adaptive Server Enterprise System Administration Guide* for information about selecting an appropriate value for number of open objects.

2 Increase the value of the number of open objects parameter to a higher value than the current run value for that parameter:

```
1> sp_configure "number of open objects", new_value
2> go
```

3 Restart Adaptive Server so that the configuration change will take effect.

Additional information

For more information, refer to:

- "number of open objects" in the *System Administration Guide*
- Open Objects Parameter May Be Too Low in this manual

Version in which this error is raised

All versions

Error 8210

Severity

21

Message text

```
Duplicate of work table descriptor found in the id hash
table for temporary object %d.
```

Explanation

Adaptive Server uses work tables to temporarily store query results. Work table descriptors are the data structures used to manage access to work tables. For each task, the server maintains a list of the descriptors for work tables in use by that task. Error 8210 is raised when a problem is detected while placing a descriptor into a task's work table descriptor list; the error results in a stack trace.

Error 8210 is raised with the following states:

State	Meaning
1	The descriptor was found to belong to an object that is not a work table.
2	There is already a descriptor for this work table in the task's work table descriptor list.

Error 8210 is due to an Adaptive Server problem with batch scripts or stored procedures that use cursors and employ referential integrity checks.

Action

Contact Sybase Technical Support to upgrade to a server version in which the problem is resolved.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Adaptive Server version and SWR rollup level
- Server error log

Note The error log contents are especially critical for analysis by Sybase Technical Support. Additional messages that may appear in the error log, along with the 8210 message, are important clues to the underlying problem.

Version in which this error is raised

All versions.

Error 8211

Severity

26

Message text

```
Mismatch found between the name and id descriptor hash
table. Descriptor hashed by name = 0x%0x and hashed by
id = 0x%0x.
```

Explanation

Session descriptors are the data structures used to manage access to Adaptive Server objects. ID and name hash tables are used to maintain session descriptors. Every session descriptor should have a corresponding entry in the ID and name hash tables. Error 8211 occurs when Adaptive Server finds two descriptors for one object or when a descriptor for an object exists in only one of the hash tables (ID or name table).

Error 8211 is caused by an Adaptive Server problem.

Action

Shut down and restart Adaptive Server to clear the cache corruption.

If the 8211 error persists, call Sybase Technical Support.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Server Version and SWR rollup level
- Server error log

- Text of all error messages

Version in which this error is raised

All versions

Error 8219

Severity

18

Message text

```
Couldn't compute lock promotion level in DES for object
%d in database %d. Using default level = %d.
```

Explanation

You can use the following configuration parameters to control how Adaptive Server handles lock promotion:

- lock promotion hwm sets the maximum number of locks allowed on a table. When the number of locks that an individual command acquires exceeds this number, Adaptive Server attempts to acquire a table lock.
- lock promotion lwm sets the minimum number of locks that are allowed on a table before Adaptive Server attempts to acquire a table lock. Adaptive Server does not attempt to acquire a table lock until the number of locks on a table is equal to the lock promotion lwm.
- lock promotion pct sets the percentage of locks (based on the table size) above which Adaptive Server attempts to acquire a table lock when the number of locks is between the lock promotion hwm and the lock promotion lwm.

If a command requires more locks than the number specified by lock promotion lwm but less than the number specified by lock promotion hwm, Adaptive Server compares the number of locks to the percentage of locks on the table. If the number is greater than the percentage, Adaptive Server attempts to issue a table lock.

Adaptive Server computes the lock promotion threshold as a percentage of the table size. If Adaptive Server is unable to get the table size, Error 8219, an informational message, is displayed and Adaptive Server uses the value of lock promotion lwm to determine the lock promotion level.

Action

Error 8219 can occur if the table size is 0 (no rows) or the number of pages for the table is 0. Run dbcc checktable and dbcc checkalloc for the table to determine if this is the case.

If this is not the case and the 8219 error occurs again, call Sybase Technical Support.

Additional information

Refer to the *Adaptive Server Enterprise Performance and Tuning Guide*, the *Reference Manual*, and the *System Administration Guide* for information about Adaptive Server lock promotion.

Before calling Technical Support, have the following information available:

- Server Version and SWR version level
- Server error log
- Text of all error messages
- Output of sp_configure (or a copy of the configuration file), sp_helpdb, sp_help on the table involved, and sp_spaceused on the table involved

Version in which this error is raised

All versions

sysindexes Manager Errors

This section contains writeups for sysindexes Manager errors.

Error 8402

Severity

21

Message text

```
Index row for object %ld (index id %d) was not found in
database %d.
```

Explanation

Each database has a sysindexes table. The sysindexes Manager maintains and provides access to the rows of the sysindexes table. It handles inserting and deleting of sysindexes rows and manages the abort and recovery of individual operations performed on the rows. It also provides read/write access to the database information structure.

sysindexes contains one row for each:

- Clustered index
- Nonclustered index
- Table that has no clustered index
- Table that contains text or image columns

Error 8402 occurs when one of these rows is expected but not found for the object specified in the error message.

Action

1 Determine the name of the database mentioned in the error message:

```
1> use master
2> go
1> select db_name(database_ID)
2> go
```

where *database_ID* is the database ID in the error message.

2 Search the sysindexes table for the object displayed in the error message:

```
1> use database_name
2> go
```

```
1> select * from sysindexes
2> where id = object_ID
3> go
```

where *object_ID* is the object ID in the error message.

3 If the object is found in sysindexes, continue with step 4.

   If the object is *not* found in sysindexes, call Sybase Technical Support.

4 If the object specified in the error message is *not* a system table (a system table's object ID is less than 100), continue with step 5.

   If the object specified in the error message is a system table, refer to Chapter 2, "How to Fix a Corrupted Index on System Tables" for instructions on how to repair the system table index. Then go to step 7.

5 Drop the index for the object specified in the error message.

6 Re-create the index. This clears the corruption in most cases.

7 Run dbcc checktable on the table to verify that the corruption is gone. If corruption still exists, call Sybase Technical Support.

Additional information

Refer to drop index and create index in the *Reference Manual* for information about dropping and re-creating indexes.

Version in which this error is raised

All versions

Error 8419

Severity

21

Message text

```
Could not find index descriptor for objid %d, indid %d
in dbid %d.
```

Explanation

Each index accessed in Adaptive Server requires a descriptor which is a "handle" for referencing the index. When the descriptor is no longer needed (for example, the query completes or the index is dropped), the server frees up the descriptor so that it can be assigned to another object as needed.

Error 8419 prevents a stored procedure or other compiled object from being recompiled after an index or a table column is dropped; consequently the server attempts to process the query via a stale plan, using the missing index or column. When the descriptor containing sysindexes information for the index is found to be missing from cache, Error 8419 is raised. Here is a typical scenario:

```
create table t1
create procedure p1 as insert into t1 values(...)
go
begin transaction
```

```
go
create clustered index on t1(..)
execute procedure p1 (query plan uses the new index)
rollback transaction (index dropped)
go
execute procedure p1 (Error 8419 raised as the server
                      attempts to use the old plan)
```

Error 8419 is due to an Adaptive Server problem and is raised with the following states:

State	Meaning
1	Index descriptor row not found in cache when attempting to make a copy of the table's first index structure.
2	Index descriptor row not found in cache when attempting to copy the next index structure for the table.
3	Could not retrieve a pointer to the first index descriptor row for the table in the cache.
4	Could not retrieve a pointer to the next index for the table in the cache as expected.
5	Could not retrieve a specific field within an index descriptor row using a pointer to the table's first index row.
6	Could not retrieve a specific field within an index descriptor row using a pointer to the table's next index row.
7	Could not retrieve a pointer to an index descriptor row in the cache after dropping the index.
8	Could not retrieve a pointer to an index descriptor row in the cache after dropping a column of the table.
9	Could not retrieve a pointer to an index descriptor row in the cache while accessing the root page of the index. This state is raised only by diagserver.
10	Could not retrieve a pointer to an index descriptor row in the cache while releasing locks associated with the row.
11	Could not find the start of the index descriptor chain.
12	Could not update a specific field within an index descriptor row using a pointer to the index row.

Action

Call Sybase Technical Support to upgrade to a version in which the problem is resolved.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Adaptive Server version and SWR rollup level

- Server error log
- Text of the query which causes the error.

Version in which this error is raised

All versions

Process Status Structure Errors

This section contains error messages for the Adaptive Server Process Status Structure (PSS).

Error 8601

Severity

20

Message text

```
Process %d tried to keep the PSS for process %d when it
already had the PSS for process %d kept.
```

Explanation

A process is a task being carried out by Adaptive Server. Processes can be in various states such as running or waiting for resources. Processes cease to exist when they complete normally or are terminated with the kill command. For each user process and some system processes, Adaptive Server maintains a Process Status Structure (PSS) to track the state of the process and its use of resources.

Typically, each process is concerned with its own PSS. However tasks such as diagnostic commands, system procedures like sp_who, and the kill command require that a process interact with the PSS for another process. For example a process performing a kill must temporarily attach ("keep") the PSS of the target process to its own PSS, read or write necessary information in the target PSS, and "unkeep" it. A process can only keep one PSS (besides its own) at a time.

Error 8601 is raised when a process is already keeping the PSS of one process, and tried to keep the PSS of another process at the same time. The error is sometimes seen when the System Administrator or other user is attempting a kill operation or running a diagnostic command.

Action

Do the following to correct this error:

- shut down and restart the Server to clear the PSS structure.
- defer kill operations and diagnostic commands to a later time.

Version in which this error is raised

All versions

Site Buffer Manager Errors

This section contains error messages for the Adaptive Server Site Buffer Manager.

Error 8704

Severity

26

Message text

```
Could not get the address lock. address=0x%1x,
status=%d.
```

Explanation

Address locks control memory addresses. They are used for:

- Object Allocation Map (OAM) pages
- Allocation pages
- Updating the disk maps in the database table memory structure
- Updating values for descriptors and site buffers

The Site Buffer Manager is one of the managers that control Server-to-Server communications.

Error 8704 occurs when the Site Buffer Manager fails to get an address lock for an active site buffer and fails to attach a logical connection to it.

This error is caused by an Adaptive Server problem.

Action

Shut down and restart Adaptive Server.

If the 8704 error persists, call Sybase Technical Support.

Additional information

Have the following information available before calling Sybase Technical Support:

- Server version and SWR version level
- Server error log
- Text of all error messages
- sp_configure, sp_who, and dbcc lock output

Version in which this error is raised

All versions

Disk Manager Errors

This section contains writeups of common Adaptive Server Enterprise Disk Manager errors.

Error 9004

Severity

20

Message text

```
I/O error detected on disk buffer 0x%X for page request
%ld, dbid %d.
```

Explanation

When you issue dbcc commands such as dbcc checkalloc or dbcc checkdb, Adaptive Server reads a specified number of pages from disk into disk buffers so that consistency checks can be carried out on the pages. Error 9004 is raised when the server attempts a synchronous disk read during dbcc processing, but the request fails with an I/O error.

Action

Check the server error log for similar errors from other operations, such as an 823 error when running a query ("I/O error detected during %S_MSG for %S_BUF.") or *sddone* errors ("sddone:write error on virtual disk %d"). Refer to the writeups for those errors elsewhere in this manual for more information. Since Adaptive Server reports 9004 and similar errors when its I/O operation gets an unsuccessful return value, these errors can be an early warning of disk problems.

Check your disks as follows:

- Note the database ID shown in the 9004 message text.
- Identify the affected device(s) using the steps shown in "How to Determine Which Physical Devices a Database is On" in Chapter 2.
- Examine the availability and condition of the device(s) using your standard operating system procedures.

Additional information

See "Checking the Operating System Error Log" in Chapter 2 if you need assistance in locating the system log.

Version in which this error is raised

All versions

Log Transfer Errors

This section contains writeups for common Adaptive Server Enterprise Log Transfer Utility errors.

Error 9122

Severity

21

Message text

```
The log transfer process lost its connection to the LTM
for database '%.*s'. Once the LTM re-establishes the
connection, replication will continue with no data loss.
```

Explanation

Through version 11.0.3 of Replication Server, the Log Transfer Manager (LTM), a replication system component, reads the transaction log of the Adaptive Server database being replicated and transfers this information to the Replication Server.

Error 9122 is raised if Replication Server and/or LTM are shut down improperly; if the shutdown occurs while sending data, Adaptive Server is unable to disconnect the LTM connection normally, resulting in a 9122 error and stack trace.

Action

Error 9122 is not a serious error. Once you restart Replication Server and LTM, the log transfer process will resume normally.

To prevent future occurrences of this error, be sure to perform a log suspend before shutting down LTM.

Version in which this error is raised

11.0.3 and earlier

Remote Create Errors

This section contains error messages for the functions that create remote objects.

Error 11202

Severity

16

Message text

```
Unable to create table '%.*s' because it has no remote
storage location defined. Creation of local user tables
is allowed only in the master and tempdb databases.
```

Explanation

Component Integration Services (CIS) is a feature that extends Adaptive Server capabilities and allows you to access and manipulate tables in remote servers as if they were local tables. When you install Adaptive Server, files needed to interact with CIS are placed in the installation directory.

Error 11202 is raised when you attempt to create a table, but Adaptive Server is unable to locate or otherwise access one of these files. The problem can occur even if you are creating a local table and do not utilize CIS.

Action

Use the following checklist to ensure that the necessary file is accessible:

1. Check that the *SYBASE* environment variable is properly set. If you recently upgraded, make sure that this variable was updated. Restart the server after checking this variable.
2. Make sure that you start the server as user *sybase*.
3. Check that the files *dsopt1* (necessary for local tables) and *dsopt2* (necessary for remote tables) exist in your installation directory. For example, in Unix sites the file *$SYBASE/install/dsopt1* must exist with read permissions before you can create local tables. Restore any missing files and restart Adaptive Server.

Note The ASE directory structure has changed with 12.0. Check your Installation and Configuration guide for details.

Version in which this error is raised

11.5 and later

reorg utility Errors

This section contains error messages for the Adaptive Server reorg utility.

Error 11903

Severity

16

Message text

```
You cannot run REORG on a table which uses allpages
locking.
```

Explanation

As of version 11.9, Adaptive Server provides two new locking schemes:

- Datapages locking, which locks only the data pages
- Datarows locking, which locks only the data rows.

Since neither scheme locks index pages, they are referred to together as the data-only locking scheme.

The pre-11.9 locking scheme, known as allpages locking, continues to be available with 11.9. This scheme locks the data and index pages affected by queries, and is the default locking scheme.

Since tables that use data-only locking have unique space management needs, Version 11.9 introduces the reorg utility, which reorganizes a table to reclaim unused space and improve clustering. Error 11903 is raised when you attempt to use reorg on an allpages locked table, since reorg applies exclusively to tables that use the data-only locking scheme.

Action

Do not use the reorg utility on an allpages locked table.

Additional information

Use sp_help to verify a table's locking scheme:

```
1> sp_help table_name
2> go
```

Version in which this error is raised

11.9 and later

Fault Isolation Utility Errors

This section contains error messages for the Adaptive Server Fault Isolation Utility.

Error 12716

Severity

17

Message text

```
Data unavailable: Unable to complete because the
required data is currently unavailable. The command
failed and the transaction aborted. Try again later or
contact your System Administrator.
```

Explanation

Adaptive Server's Recovery Fault Isolation (RFI) feature enables recovery to mark a single page or a set of pages of a user database suspect, while allowing access to the remainder of the database. The bad page(s) are said to be "offline", and may not be used until the problem is corrected. (For a detailed discussion of RFI, refer to Chapter 2, “Online Recovery and Recovery Fault Isolation”).

Error message 12716 results when you issue a query that requires access to an offline page. This error is followed by Error 12717 (Suspect data encountered: Database '%.*s', object '%.*s', index ID '%ld', page ID '%ld'.).

Action

The offline page(s) must be repaired before they can be used. Responsibility for repairing offline pages belongs to a user with sa_role. The System Administrator can bring a page online to make it accessible to sa_role:

```
1> sp_forceonline_page database_name, pgid, 'sa_on'
2> go
```

where pgid is the page ID shown in the accompanying 12717 error.

Warning! When forcing pages online, note that Adaptive Server does not do any consistency checks before onlining pages.

Use existing methods to diagnose and repair problems, checking the server error log for any errors related to the offline page and determining appropriate corrective action. Before making repaired pages available to all users, check that repairs have been made correctly by running dbcc commands or queries with known results.

Additional information

Refer to Chapter 2, "Online Recovery and Recovery Fault Isolation" for a detailed discussion of RFI.

Version in which this error is raised

11.9.2 and later

Error 12717

Severity

17

Message text

```
Suspect data encountered: Database '%.*s', object
'%.*s', index ID '%ld', page ID '%ld'
```

Explanation

Adaptive Server's Recovery Fault Isolation (RFI) feature enables recovery to mark a single page or a set of pages of a user database suspect, while allowing access to the remainder of the database. The bad page(s) are said to be "offline", and may not be used until the problem is corrected. (For a detailed discussion of RFI, refer to Chapter 2, "Online Recovery and Recovery Fault Isolation").

Error message 12717 results when you issue a query that requires access to an offline page. This error is preceded by Error 12716 (Data unavailable: Unable to complete because the required data is currently unavailable. The command failed and the transaction aborted. Try again later or contact your System Administrator.).

Action

The offline page(s) must be repaired before they can be used. Responsibility for repairing offline pages belongs to a user with sa_role. The System Administrator can bring a page online to make it accessible to sa_role:

```
1> sp_forceonline_page database_name, pgid, 'sa_on'
2> go
```

where pgid is the page ID shown in the 12717 error.

Warning! When forcing pages online, note that Adaptive Server does not do any consistency checks before onlining pages.

Use existing methods to diagnose and repair problems, checking the server error log for any errors related to the offline page and determining appropriate corrective action. Before making repaired pages available to all users, check that repairs have been made correctly by running dbcc commands or queries with known results.

Additional information

Refer to Chapter 2, "Online Recovery and Recovery Fault Isolation" for a detailed discussion of RFI.

Version in which this error is raised	11.9.2 and later

Create Utilities Errors (continued)

This section contains additional writeups for Adaptive Server Enterprise Create Utility errors.

Error 12818

Severity

16

Message text

```
Cannot create an object of this type in system-created
proxy database.
```

Explanation

Adaptive Server's High Availability (HA) feature allows two nodes in a cluster to take over each other's workload in the event of a failure. The server that fails is called the primary, and the server that takes over the workload is the secondary. Movement of data access services from a failed primary server to a companion secondary server is known as failover.

In the HA architecture, the secondary companion is configured with proxy databases which act as placeholders; in the event of failover, proxy databases are dropped and replaced with databases mounted and recovered from the primary server.

Error 12818 is raised when you attempt to execute certain Data Definition Language (DDL) commands in a proxy database, including:

- create/drop procedure
- create/drop view
- create/drop trigger

These commands are not available in a proxy database, which is a system database reserved solely for HA use in the event of failover.

Action

No action is needed. User objects should not be placed in a proxy database.

For details about the DDL commands supported in primary and secondary servers in an HA environment, see *Using Sybase Failover in a High Availability System*.

Version in which this error is raised

12.0 and later

System Procedure Errors

This section contains error messages raised by Adaptive Server system procedures.

When system procedure error messages are raised, the error number is displayed as 0 and the global variable @@error returns 0. This section includes non-zero error numbers along with the error message text. If you call Sybase Technical Support about an error in this section, include the error number that is shown here even though that number was not displayed when the error was raised.

Error 17417

Severity

0

Message text

```
Maximum file descriptors or FILLM process quota too low
to support requested number of user connections.
Configuration variable 'user connections' will not be
modified.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when you attempt to increase the value of the configuration parameter number of user connections using sp_configure and the new value is too large for your current operating system parameters.

Although most of the open file descriptors are available for user connections, a few are used by Adaptive Server for opening files and devices. The number of user connections allowed by Adaptive Server is limited by the number of open file descriptors allowed by the operating system. The number allowed is platform- and version-specific. For example:

- *SunOS 4.1.1* (Sun Solaris 1.x) has 256 file descriptors available. It has 1,024 available if DBE, the Database Excelerator, is installed.

  SunOS 4.1.3 has both “hard” and “soft” limits as settable parameters. The soft limit determines the number of open file descriptors available to Adaptive Server up to the maximum of 1,024 on a 4.1.3 operating system. A system running 4.1.3 has 256 file descriptors available, 2,048 if the Sun Database Excelerator (DBE) is installed.

- *SunOS 5.x* (Sun Solaris 2.x) sets both “hard” and “soft” limits: the soft limit can be increased by the user up to the hard limit, but the hard limit can be changed only by someone with “superuser” permissions. The soft limit determines the number of open file descriptors available to Adaptive Server, up to a maximum of 10,000 on a SunOS 5.x operating system. The SunOS 5.x operating system has a maximum of 3,072 file descriptors on a 4.9.1 SQL Server.

  Note To set the hard limit to a value greater than 1024, you must be logged in as “root.”

To verify the maximum number of user connections you are allowed to configure, run the following query in isql:

```
1> select @@max_connections
2> go
```

The value returned is the maximum number of user connections configurable, based on the number of available file descriptors. If the number of connections is not sufficient for your needs, alter the kernel on your operating system.

Action

1 Verify the configured descriptor soft and hard limits:

Type	C Shell	Bourne Shell
Soft	% limit descriptors	$ ulimit -n
Hard	% limit -h descriptors	$ ulimit -Hn

2 To increase the soft limit, use one of the following commands:

C Shell	Bourne Shell
% limit descriptors *x*	$ ulimit -Sn *x*

where *x* is the value to which you want to increase the soft limit.

These commands can be used in your runserver file to increase the soft limit every time Adaptive Server is started. Because the runserver file is a Bourne Shell script in the Sun Solaris operating system, use the Bourne Shell version of these commands in the runserver file.

3 Refer to the Adaptive Server installation and configuration guides for instructions on changing the hard limit.

> **Note** For SunOS 4.1.3, since the number of user connections cannot exceed the number of file descriptors available, you cannot increase user connections to a higher value than the hard-coded number of file descriptors.

Recalculate the necessary user connections. Refer to your Adaptive Server installation and configuration guides and recalculate your operating system process quotas.

Additional information

For information about user connection limits for your operating system, refer to your Adaptive Server installation and configuration guides.

Version in which this error is raised

All versions

Error 17461

Severity

```
0
```

Message text

```
Object does not exist in this database.
```

Explanation

This error occurs when the object you specified does not exist in the current database when you run any of the following stored procedures:

- sp_auditobject
- sp_auditsproc
- sp_depends
- sp_estspace
- sp_foreignkey
- sp_help
- sp_helpartition
- sp_helpconstraint
- sp_helpindex
- sp_helptext
- sp_primarykey

- sp_procqmode
- sp_remap
- sp_spaceused

Action

Make sure you are in the database in which the object resides before entering the stored procedure command.

Use the owner name in your object name specification if you are not the object owner.

Version in which this error is raised

All versions

Error 17715

Severity

0

Message text

```
The last-chance threshold for database %1! is now %2!
pages.
```

Explanation

This informational message is displayed when a last-chance threshold has been established on a newly created log device (sp_logdevice), or an existing one has been adjusted after a procedure which alters the size of a log device was executed (sp_dropsegment, sp_extendsegment). The threshold is established automatically by Adaptive Server.

Action

No action is necessary.

Version in which this error is raised

All versions

Error 17716

Severity

0

Message text

```
Could not update the last-chance threshold for database
%1!
```

Explanation

Procedure lct_admin is run automatically when any attempt is made to add or modify the log segment of a database with sp_logdevice, sp_dropsegment, or sp_extendsegment. Error 17716 occurs when an attempt to run such a command has failed.

Action

Check other messages displayed near the attempted command. They may give you additional information regarding the failure of lct_admin or may indicate the original command was not successful. If the command appeared to run successfully, except for this error message, contact Sybase Technical Support about evaluating the current status of the last-chance threshold in this database.

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log
- Text of all error messages

Version in which this error is raised

All versions

Error 17737

Severity

```
0
```

Message text

```
Partitioned objects cannot be moved.
```

Explanation

Adaptive Server does not allow you to use sp_placeobject on partitioned objects. Error 17737 occurs when you try to use sp_placeobject to change future space allocations for a partitioned table.

Action

If you want to use sp_placeobject on the table, use the unpartition clause of the alter table command to concatenate all partitions:

```
1> alter table table_name unpartition
2> go
```

Then run your sp_placeobject command again.

Additional information

Refer to the *Reference Manual* for information about alter table and sp_placeobject.

Version in which this error is raised

All versions

Error 17870

Severity

```
0
```

Message text

```
Table '%1!' does not exist in database '%2!' -- cannot
```

```
add thresholds.
```

Explanation

systhresholds should exist in every database in Adaptive Server. Error 17870 occurs when the system table systhresholds does not exist in the current database, and sp_addthreshold or sp_modifythreshold was not able to add or modify the requested threshold.

Action

Run dbcc checkalloc and dbcc checkdb to determine if the database is corrupted.

Contact Sybase Technical Support if systhresholds is missing from a database.

Have the following information ready when you call Technical Support:

- Output of select *@@version*
- Output of all error messages
- Output of dbcc commands

Additional information

Refer to “sp_addthreshold” in the *Reference Manual* for additional information.

Version in which this error is raised

All versions

Error 17871

Severity

```
0
```

Message text

```
There is no segment named '%1!'.
```

Explanation

This error occurs when an attempt is made to add, drop, or modify (using **sp_addthreshold**, sp_dropthreshold, or sp_modifythreshold) a segment that does not exist in the current database.

Action

Run sp_helpsegment to display segment information about the current database. If the segment does exist, check that you have spelled it correctly in your procedure. If the segment does not exist and you believe it has been dropped, use sp_addsegment to add it back.

If the requested segment does not exist and you are certain no one has dropped it, run dbcc commands to check the integrity of the database.

Additional information

Refer to “sp_addthreshold” in the *Reference Manual* for additional information.

Version in which this error is raised

All versions

Error 17872

Severity	0
Message text	This threshold is too close to one or more existing thresholds. Thresholds must be no closer than 128 pages to each other.
Explanation	This error occurs when an attempt is made to add or modify thresholds **sp_addthreshold** or (sp_modifythreshold) that are too close together.
Action	Thresholds for a database cannot be closer than (*size_of_threshold* + (2 * *@@thresh_hysterisis*)). Ensure that the free space allocation allows enough pages. Rerun the command with the appropriate value for free space. Refer to "Managing Free Space with Thresholds" in the *Adaptive Server Enterprise System Administration Guide* for information about threshold size and hysteresis value.
Version in which this error is raised	All versions

Error 17873

Severity	0
Message text	Adding threshold for segment '%1!' at '%2!' pages.
Explanation	This message is displayed after the successful completion of a sp_addthreshold command or a sp_modifythreshold command that adds a threshold.
Action	No action is required.
Version in which this error is raised	All versions

Error 17874

Severity	0
Message text	A threshold at %1! pages is logically impossible for segment '%2!'. Choose a value between %3! and %4! pages.
Explanation	This error occurs when an attempt is made to either add or modify a threshold (sp_addthreshold or sp_modifythreshold) for a number of pages that is greater than the segment named in the command.

Action	Run sp_helpdb *database_name* to display the device fragments used by the specified database and the size of the device fragment. Then rerun the command with an appropriate page value for *free_space*.
Version in which this error is raised	All versions

Error 17875

Severity	0
Message text	`Only the DBO of database '%1!' or a user with System Administrator (SA) role may add, delete, or modify thresholds in that database.`
Explanation	This error occurs when you try to run sp_addthreshold, sp_dropthreshold, or sp_modifythreshold and you are not the owner of the current database and do not have "sa" role.
Action	If you are not the database owner, run sp_displaylogin to display roles currently assigned to you (you must have sso_role or sa_role to display this information about another user). Ask your system administrator to give you sa_role or ask another user to execute the command.
Version in which this error is raised	All versions

Error 17903

Severity	0
Message text	`Table '%1!' does not exist in database '%2!' -- cannot drop thresholds.`
Explanation	systhresholds should exist in every database in Adaptive Server. Error 17903 occurs when the system table systhresholds does not exist in the current database, and sp_dropthreshold is unable to drop the requested threshold.
Action	Run dbcc checkalloc and dbcc checkdb to determine if there is corruption in the database. Contact Sybase Technical Support if systhresholds is missing from a database.

Additional information	Have the following information ready when you call Technical Support: • Output of select *@@version* • Output of all error messages • Output of dbcc commands
Version in which this error is raised	All versions

Error 17904

Severity	0
Message text	`Segment '%1!' does not have a threshold at '%2!' pages.`
Explanation	This error occurs when an attempt is made to either drop or modify a threshold **(sp_dropthreshold** or sp_modifythreshold) that does not exist for the free space page allocation given in the command.
Action	Run sp_helpthreshold without any options to display all the thresholds and their free page allocations for the current database. Then rerun the stored procedure with corrected parameters.
Version in which this error is raised	All versions

Error 17905

Severity	0
Message text	`You may not drop the log's last-chance threshold.`
Explanation	The last-chance threshold is created and maintained by the Adaptive Server and cannot be dropped by a user. Error 17905 occurs when an attempt is made to drop the database last-chance threshold with sp_dropthreshold.
Action	No action is required.
Version in which this error is raised	All versions

Error 17906

Severity	0
Message text	Dropping threshold for segment '%1!' at '%2!' pages.
Explanation	This message is issued after successful completion of an sp_dropthreshold command or an sp_modifythreshold command that drops a threshold.
Action	No action is required.
Version in which this error is raised	All versions

Error 17910

Severity	0
Message text	Database '%1!' has no thresholds -- table '%2!' does not exist.
Explanation	systhresholds should exist in every database in Adaptive Server. Error 17910 occurs when the system table systhresholds does not exist in the current database and sp_helpthreshold is unable to output information from that table.
Action	Run dbcc checkalloc and dbcc checkdb to determine if there is database corruption.
Additional information	Have the following information ready if you call Sybase Technical Support: • Output of select *@@version* • Output of all error messages • Output of dbcc commands
Version in which this error is raised	All versions

Error 18031

Severity	0
Message text	This procedure can only affect thresholds in the current database. Say 'USE %1!', then run this procedure again.

Explanation

Thresholds for a specified database may only be created or modified while in the database. Error 18031 occurs when you try to run sp_addthreshold, **sp_dropthreshold**, or sp_modifythreshold when you are not in the database for which you are running the command.

Action

1 Determine the name of the current database:

```
1> select db_name()
2> go
```

2 Select the correct database:

```
1> use database_name
2> go
```

where *database_name* is the name of the database for which you want to run the threshold command.

3 Rerun the sp_addthreshold, sp_dropthreshold, or sp_modifythreshold command.

Version in which this error is raised

All versions

Error 18032

Severity

0

Message text

```
You may not alter the free space or segment name of the
log's last-chance threshold.
```

Explanation

The only parameter you can change for the last-chance threshold is the name of the procedure that will execute when the threshold is crossed. Error 18032 occurs when you try to modify the last-chance threshold name or free space allocation with sp_modifythreshold.

Action

No action is required.

Version in which this error is raised

All versions

Error 18033

Severity

0

Message text

```
You cannot drop user because user '%1!' owns thresholds
```

```
in database.
```

Explanation

The user who created a threshold or last modified a threshold is the threshold owner. A user who owns a threshold cannot be dropped. Error 18033 occurs when you use sp_dropuser to try to drop a user who owns a threshold.

Action

You have three options:

- Drop the threshold procedure owned by the user. Then you can drop the user.
- Have another user with appropriate privileges modify the threshold. Then you can drop the user.
- If you want to inactivate the user Adaptive Server login while leaving the threshold active, use sp_locklogin to prevent use of the login.

Version in which this error is raised

All versions

Error 18097

Severity

```
0
```

Message text

```
Object must be a trigger, view or stored procedure.
```

Explanation

The procedure sp_procqmode displays the query processing mode of a stored procedure, view, or trigger. The *object name* parameter is the name of the stored procedure, view, or trigger whose query processing mode you are examining. Error 18097 occurs when the object name you specify is not a trigger, view, or stored procedure.

Action

Check the name of your object and make sure it is the name of a trigger, view, or stored procedure.

Additional information

Refer to the *Reference Manual* for information about sp_procqmode.

Version in which this error is raised

All versions

Error 18123

Severity

```
0
```

Message text

```
Configuration option changed. The SQL Server must be
rebooted before the change in effect since the option
```

```
is static.
```

Explanation

When you use sp_configure to modify a static parameter:

- The configuration value is updated
- The configuration file is updated
- The change takes effect when you restart Adaptive Server

Error 18123 is an informational message letting you know that a static parameter has been changed and that the change will not take effect until Adaptive Server is restarted.

Action

Shut down and restart Adaptive Server to make the parameter change take effect.

Additional information

Refer to the *Reference Manual* for information about sp_configure.

Version in which this error is raised

All versions

Error 18145

Severity

```
0
```

Message text

```
Less memory moved than requested in cache '%1!'.
Requested size = %2! Kb: from pool = %3!, to pool = %4!,
actual memory moved = %5! Kb.
```

Explanation

The Adaptive Server command sp_poolconfig allows you to create, drop, resize, and get information about memory pools within named data caches. You can use sp_poolconfig to move memory from one pool to another.

Error 18145 is an informational message that is displayed when the source pool you specified has less memory available then you requested to have moved to the destination pool. Adaptive Server moves as much memory as it can of the amount you requested and specifies in the message output how much was actually moved.

Action

If the amount of memory that was moved to the destination pool is not enough:

1 If necessary, use sp_cacheconfig to enlarge the cache size.

2 Use sp_poolconfig to reconfigure the source pool to have a larger amount of memory available.

Additional information

Refer to the *Reference Manual* for information about sp_poolconfig and sp_cacheconfig.

Version in which this error is raised

All versions

Error 18750

Severity

0

Message text

```
The command '%1!' can not be run in this server. The
server is currently configured with mode '%2!'.
```

Explanation

In Adaptive Server's High Availability (HA) cluster configuration, two nodes in a cluster can take over each other's workload in the event of a failure; once the problem with the primary is resolved, you can use failback to move the primary companion's shared disks from the secondary node back to the primary node and start the primary companion on the primary node.

Error 18750 is raised if you try to assign a (primary or secondary) server to a mode that is incompatible with its current state. For example, the error is raised if you issue sp_companion 'prepare_failback' on the secondary companion, but the secondary is not in the correct state for failback.

Action

Take the following steps:

1 Check the companion mode of your servers by issuing the following on both primary and secondary servers:

```
1> select @@cmpstate
2> go
```

This command returns a value denoting the current companion mode. Refer to "Determining the Companion's Mode" in *Using Sybase Failover in a High-Availability System* for a description of the return values.

2 Check that the command which raised the 18750 message (appearing in the message text) is compatible with the server's companion mode. For example, you cannot issue a suspend command to a server currently in failover mode.

At times the 18750 message may require corrective action. For example, if you issued prepare_failback and your primary companion is in normal companion mode but the secondary companion is in secondary failover mode, you will need to recover from this inconsistent state manually. The process involves shutting down both servers, restarting the secondary and repairing any suspect databases on the secondary. The steps for achieving this are platform-specific. Refer to one of the following:

- Troubleshooting Sybase Failover on HP in *Using Sybase Failover in a High-Availability System*
- Troubleshooting Failover on HACMP for AIX in *Using Sybase Failover in a High-Availability System*
- Troubleshooting Failover on TruCluster for Digital Unix in *Using Sybase Failover in a High-Availability System*
- Troubleshooting Sybase Failover on Windows NT in *Using Sybase Failover in a High-Availability System*

Version in which this error is raised: 12.0 and later

Kernel Errors

This section includes error messages generated by the Adaptive Server kernel. They have text descriptions but no error numbers. These errors are displayed only in the Adaptive Server error log. They are presented in alphabetical order.

Buffer Mismatch Error

Message text

```
Bufunhash: Buffer %lx from database '%.*s' has page
number %ld  in the page header and page number %ld in
the buffer header. Buffer will be unhashed, no action
is necessary. Printing headers for diagnostics:
```

Note This error may be caused by a hardware problem.

Explanation

This error only appears in the error log, and it means that a cache management problem occurred. This can be a very serious error because it is often followed by database corruption, such as 605 errors.

Although often a result of hardware failure, this error can also be caused by operating system or Adaptive Server problems.

Action

Note If possible, shut down and restart Adaptive Server immediately after this error occurs. This may prevent the buffer cache error from being flushed to disk.

1. Run complete diagnostics on the machine running Adaptive Server and on all disk drives and controllers attached to that machine.
2. Run complete dbcc checks, including dbcc checkalloc and dbcc checkdb, on any databases involved.
3. Repair or replace faulty hardware.
4. Shut down and restart Adaptive Server.
5. If you do not find any hardware problems, call Sybase Technical Support immediately.

Additional information

Before calling Technical Support for assistance, have the following information available:

- Server version and SWR version level
- Server error log

- Hardware error log
- Text of all the error messages
- Reproducible case (if possible)

Character Set Error

Message text

```
Adaptive Server doesn't support use of
character_set_name character set.
```

Explanation

This error occurs when the client software default character set has not been installed on the Adaptive Server to which the client is trying to connect. This happens most often when the client is running on a different operating system from Adaptive Server.

Action

To allow the client and Adaptive Server to communicate, do the following:

1 Determine whether the client software character set is installed on Adaptive Server by executing this query:

```
1> use master
2> go
1> select name, id, csid from syscharsets
2> go
```

If the client software's default character set appears in the output, have your primary site contact call Sybase Technical Support. If it does not appear, go to step 2.

If the client software's default character set does not appear in the output, you may be able to set the client software to the character set which Adaptive Server is using. For isql, you can do this with the -J flag. Refer to "Configuring Client/Server Character Set Conversions" in the *System Administration Guide* for details. If you cannot set the client software to the character set Adaptive Server is using, go to step 2.

Warning! Familiarize yourself with all the instructions in the Adaptive Server installation and configuration guide before completing step 2, in order to make the character set addition a smooth process.

2 Run the installation utility sybinit (or the equivalent for your platform). Choose the additional character set item from the character set menu. If the particular character set you need is not one of the options listed, you must purchase that character set by purchasing a Language Module from Sybase, and install it with sybinit or its equivalent.

3 Try again to connect to Adaptive Server with the client software. Call Sybase Technical Support if you are still unable to connect to Adaptive Server.

Could Not Create Shared Memory Error

Message text

```
kistartup: could not create shared memory
```

Note This error may be caused by a hardware problem.

Explanation

kistartup is the Adaptive Server kernel initialization routine. It performs the following tasks:

- Reads the configuration parameters
- Allocates kernel tables and buffer pools
- Sizes memory
- Initializes devices
- Spawns the service processes
- Starts the scheduler.

When kistartup is unable to create shared memory for kernel initialization, the error "kistartup: could not create shared memory" is written to the Adaptive Server error log.

kistartup errors occur after os_create_region errors.

Action

Refer to os_create_region Errors for information about recovering from the os_create_region errors which precede the kistartup error. Once those errors are corrected, the kistartup error will not occur.

Interaction with Monitor Server

Shared memory errors may be seen if you shut down Adaptive Server and restart it while Monitor Server is still active. You can correct this type of error by shutting down Monitor Server before restarting Adaptive Server.

The interaction between Adaptive Server and Monitor Server is covered fully in the *Monitor Server User's Guide.*

Current Process Infected with %d

Message text

```
current process (0x%x) infected with %d
```

Note This error may be caused by a hardware problem.

Explanation

Adaptive Server reports this error when it detects a UNIX signal specifying an error. The signal values ("%d") displayed in the above message vary by platforms and Adaptive Server versions; the most common are either 10 or 11.

Current process infected with 10

A value of 10 [SIGBUS] means that the operating system detected an address alignment error or a miscellaneous hardware error (for example, bus timeout).

A timeout can occur when the CPU issues a request across the bus for the contents of a memory location, and that request is not answered within that CPU's timeout period (usually a few nanoseconds).

Current process infected with 11

A value of 11 [SIGSEGV] means that the operating system detected a segment violation error.

Sometimes this error occurs in conjunction with stack overflow or data corruption. For more information on stack overflow, refer to Stack Guardword Error.

The message appears in the Adaptive Server error log followed by a stack trace. The SQL causing error or the *lasterror* displayed in the Server error log may be the underlying cause for this error. But they can also be just the last data Adaptive Server had in its cache space.

In order to identify the *lasterror* (except in the cases where the *lasterror* is 0), get the number displayed by Adaptive Server in the *lasterror* field from the Adaptive Server error log and consult this manual for more information on this error number.

In the following example, the value for *lasterror* is 614.

```
00: 94/02/14 11:32:26.02 kernel: current process (0x1fb001d)
infected with 11
```

```
00: 94/02/14 11:32:26.07 kernel: Address 0x808a6ef
(closetable+0x2f7), siginfo (code, address) = (2, 0x30)
00: 94/02/14 11:32:26.07 kernel: **********************************
00: 94/02/14 11:32:26.07 kernel: SQL causing error : CREATE TRIGGER
00: 94/02/14 11:32:26.07 kernel: curdb = 22 pstat = 0x10018
lasterror = 614
```

Action

In order to correct this error, try to eliminate the *lasterror* first, as this might be one of the causes for this error (except when *lasterror* is 0).

Also, try rerunning the command referenced in the SQL causing error to see if the problem reoccurs.

If the process is infected with 11 and the problem can be reproduced, try and correct it as follows:

1 If the SQL causing error is a compiled object such as a stored procedure, trigger, or view, drop and recreate the object.

2 If the SQL causing error is ad hoc rather than a compiled object, moving the data may fix the problem. Use one of these options:

- Select the table data into a new table, drop the old table, and rename the new table to the old table name.
- Bulk copy the affected table out, drop and re-create the table, and bulk copy back in. This is the most efficient solution for a very large table.

Note If moving the data corrects the problem, the data may have been corrupt. Be aware that moving corrupted data can lead to a data loss.

Check your hardware error log as this error can be caused by hardware failure as well.

Additional information

If this error occurs frequently, contact Sybase Technical Support. It would be useful to have a reproduction case (for example, a query which always causes the “infected” message) before calling Technical Support.

Before calling Technical Support for assistance, have the following available:

- Server version and SWR version level
- Server error log
- Operating system error log
- Text of all the error messages
- Reproducible case (if possible)

dopen Error

Message text

```
dopen: open '%s', %s\n
```

Explanation

A dopen error is raised when Adaptive Server is unable to complete an asynchronous or synchronous disk I/O operation on a unix platform. For example:

```
00:1999/12/26 22:59:10.97 kernel  dopen: open '/dev/vx/rdsk/db01076_m',
No such file or directory
```

dopen errors can occur during the Adaptive Server startup sequence, or during an engine startup. Depending on the nature of the problem, they may be accompanied by other messages such as *dopen* errors or numbered errors. A *dopen* error aborts the startup sequence.

Action

Examine the server error log and note the exact *dopen* message text, the text immediately preceding and following the message, and any subsequent errors in the log. Corrective action depends on the type of problem encountered:

dopen: open '%s', No such file or directory

Make sure that the device, identified by the path shown in the message, exists and is correctly specified in your *RUN_server* file.

If the device no longer exists, you will need to restore from known, clean backups.

dopen: open '%s', Permission denied

Check permissions on the directory or disk file shown in the message. Verify that the account starting the server has the appropriate permissions.

dopen: open '%s', I/O error

Since Adaptive Server reports this error when its I/O operation does not get a successful return value, this type of *dopen* error can be an early warning of disk problems. Examine the availability and condition of the device using your standard operating system procedures.

If you use mirroring, a *kdconfig* error following a dopen I/O error may indicate a mirror problem. For example:

```
00:99/05/26 16:05:29.64 kernel kdconfig: opening secondary master device
00:99/05/26 16:05:29.65 kernel dopen: open '', No such file or directory
00:99/05/26 16:05:29.65 kernel kdconfig: unable to read secondary master
device
```

Refer to udunmirror Errors in this manual for information about resolving mirror problems.

Additional information

Have the following information ready when you call Sybase Technical Support:

- Adaptive Server version and SWR rollup level
- Server error log
- Operating System error log
- *RUN_server* file

Version in which this error is raised

All versions

Dstart I/O Request Repeatedly Delayed

Message text

```
kernel: dstartio: I/O request repeatedly delayed; block
number: %ld, vdn: %ld
```

Note This error may be caused by a hardware problem.

Explanation

This error message is not dangerous and there is no data corruption connected with it unless it is reported together with other kernel messages in the error log which indicate corruption.

Error Message

If this is an error message, it is a sign of inconsistencies or corruption relating to the asynchronous I/O capabilities of Adaptive Server. In this case, the above error message will be reported along with other error messages in the Adaptive Server error log.

If this message is accompanied by other error or kernel messages in the Adaptive Server error log, refer to the appropriate section of this manual (by error number or name) for corrective measures regarding those other messages.

Warning Message (Informational)

This can be an informational message (occurring by itself, any number of times). If this is a warning message, it is displayed by itself in the Adaptive Server error log. In this case, the message is displayed in the Adaptive Server error log when:

- Adaptive Server is configured for a different rate of concurrent asynchronous I/O requests than the operating system is configured for. Every time an I/O request has to be requeued in Adaptive Server's I/O queues, there is a decrease in overall performance due to requeueing overhead.
- Adaptive Server is configured for a lower throughput rate of concurrent asynchronous I/O than it is capable of.

Asynchronous disk I/O allows Adaptive Server to initiate a disk operation to continue processing other data while the disk operation is completed. This can improve performance and recoverability especially when the master device, the user databases, and transaction logs are on raw disk partitions. For more information on raw disk partitions, refer to Chapter 2, "Device Administration Issues".

Action

The following I/O-related Adaptive Server configuration parameters and operating system parameters play a role when this error occurs:

- disk i/o structures – the initial number of block I/O structures at Adaptive Server startup. This is an Adaptive Server configuration parameter.
- max async i/os per server – the number of asynchronous I/O requests outstanding at one time for Adaptive Server. This Adaptive Server configuration parameter applies only to multiprocessor machines.
- max async i/os per engine – the number of asynchronous I/O requests outstanding at one time for one engine (in the case of multiprocessor machines) and the maximum number of asynchronous I/O requests outstanding at one time for Adaptive Server (in the case of single processor machines). This is an Adaptive Server configuration parameter.
- Asynchronous I/O operating system parameters.

Determine the current I/O-related Adaptive Server configuration parameter settings by running the following commands:

```
1> sp_configure "disk i/o structures"
2> go
1> sp_configure "max async i/os per engine"
2> go
1> sp_configure "max async i/os per server"
2> go
```

Use the instructions in the sections below to determine appropriate values and then use sp_configure to change the values of disk i/o structures, max async i/os per server, and max async i/os per engine. In general:

- set max async i/os per server to a value less than or equal to the total allowable I/O count for the system on which the server is running.
- set max async i/os per engine to a value less than or equal to max async i/os per server.

For SunOS 4.1

Check the UNIX kernel */usr/share/sys/sys/asynch.h* file and record the value of the kernel parameter *MAXASYNCHIO* (maximum rate of concurrent asynchronous I/O requests, usually equal to 200 or more).

Check the values of the I/O-related Sybase configuration parameters and adjust them while Adaptive Server is shut down in order to match the value of *MAXASYNCHIO*.

If you are running on SunOS 4.1.3 (Sun Solaris 1.x), the UNIX kernel parameter *perproc_maxasyncio* (per-process limit on concurrent asynchronous I/O, found in */usr/share/sys/conf.common/param.c*) defaults to 50. Increase this parameter to 200, to match the value of the Adaptive Server configuration parameter max async i/os per server. Rebuild the UNIX kernel after making the change to the UNIX kernel parameter.

The above recommendations are made to help you clear the error message. For asynchronous I/O tuning, contact your hardware vendor to find out what are the operating system parameters that improve asynchronous I/O performance on your platform.

For SunOS 5.x (Sun Solaris 2.x)

On older versions (prior to Solaris 2.4), check whether there is a hard-coded value for *MAXASYNCHIO* (usually set to 200 and stored in */usr/include/sys/asynch.h* or in */etc/system*). If there is a hard-coded value for *MAXASYNCHIO*, set the values for Adaptive Server's asynchronous I/O configuration parameters to match the value of *MAXASYNCHIO*.

As of Solaris 2.4, *MAXASYNCHIO* is not configurable. Adaptive Server's asynchronous I/O configuration parameters should not be set higher than 200.

For AIX

Configure the operating system for a higher number of outstanding asynchronous I/Os using smit. Bring up the smit aio screen by typing smit aio at the AIX prompt. Then, go to the “Change/Show Characteristics of Asynchronous I/O” screen and raise the “maximum number of REQUESTS” parameter to 4096.

This change will improve Adaptive Server performance. You will *not* need to reboot the operating system for the change to take effect. For more information on asynchronous I/O, search for “asynchronous” using the info utility.

Also, raise the values of the Adaptive Server I/O-related configuration parameters to 4096 to match the “maximum number of REQUESTS” value in smit.

For Digital UNIX

In Digital UNIX releases before 3.0 (Digital OSF), the number of outstanding asynchronous I/Os was hard-coded in the operating system kernel in the *AIO_MAX* parameter (stored in a header file in */sys*). Its default value was 64. For those releases, the Adaptive Server I/O-related configuration parameters should be set to 64 so that Adaptive Server can match the operating system's rate of concurrent I/O requests.

In 3.0 and later releases, you can change the number of outstanding kernel asynchronous I/Os. Refer to the Adaptive Server installation and configuration guides for details.

Also, raise the values of the Adaptive Server I/O-related configuration parameters to match the value of the kernel's maximum outstanding asynchronous I/O requests.

For HP-UX

Depending on the operating system kernel, the maximum number of outstanding asynchronous I/O requests can be either 50 (on older HP-UX 8.x kernels) or 500 (on patched HP-UX 8.x kernels and default on HP-UX on 9.x kernels). this value is stored in the *MAX_CONCURRENT* parameter in */usr/include/sio/asynch*. Adjust Adaptive Server's I/O-related configuration parameters to match *MAX_CONCURRENT*.

The error log example below shows the Adaptive Server start-up messages associated with the older 8.x asynchronous I/O driver and its limit of 50 asynchronous requests at one time. Once a later asynchronous patch has been applied or the new operating system version has been installed, this limit is increased to 500.

```
00: 92/06/04 14:38:07.10 server: Number of blocks left
for proc headers: 351.
00: 92/06/04 14:38:07.36 kernel: basis_dioinit: Using older version
of asynch I/O driver with 50 concurrent requests
00: 92/06/04 14:38:07.38 server: Opening Master database...
```

Other Operating Systems

For other operating systems, check your operating system manuals for information about asynchronous configuration issues.

If problems persist, refer to the error documentation in this manual or contact Sybase Technical Support.

Additional information

Refer to:

- "raw partition" in the Adaptive Server installation and configuration guides.
- "Device Administration Issues" in Chapter 2 of this manual.
- Your operating system documentation on asynchronous I/O.

Have the following information ready when you call Technical Support:

- Output from sp_configure for I/O-related parameters or listing of your configuration file
- Server version and SWR version level
- Server error log
- Operating system error log
- Text of all the error messages

Failed to Open Device Error

Message text

```
kdopen: failed to open device %s for vdn %d
```

Note This error may be caused by a hardware problem.

Explanation

Before using a database device, Adaptive Server attempts to perform some initialization tasks. When mirroring is being used, if Adaptive Server is unable to open the secondary side of a mirrored device for these initialization tasks, the message "kdopen: failed to open device %s for vdn %d" is written to the error log. *%s* is the physical path to the secondary side of the mirrored device and *%d* is the virtual disk number of the device.

This error can be raised for the following reasons:

- A database device is offline.
- A database device has been removed or renamed.

- A database file or device is inaccessible to Adaptive Server for any other reason such as insufficient permissions.

Action

1 Examine the Adaptive Server error log for other messages that may explain why this error has occurred. Refer to the instructions in this manual for resolving other errors associated with this error.

2 Identify the device involved in the error:

```
1> select name from master..sysdevices
2> where low/power(2,24) = vdevno
```

where *vdevno* is the virtual disk number of the device as shown in the error message.

3 Examine the availability and condition of the device involved using your standard operating system procedures and make sure the device is accessible.

File Already in Use Error

Message text

```
basis_dlock: file '%s' already in use by a SQL Server
```

Note This error may be caused by a hardware problem.

Explanation

During startup, Adaptive Server tries to lock the master device. If Adaptive Server finds that the master device is already locked, the message "basis_dlock: file '%s' already in use by an Adaptive Server" is written to the error log and Adaptive Server does not start. *%s* is the name of the master device file.

The most common reasons for this error are:

- You are trying to start a server with a master device that is already being used by another Adaptive Server.
- The Server you are trying to start is already running.

Action

At the operating system level, check the processes that currently exist and determine whether an Adaptive Server using the master device you specified is already running.

If an Adaptive Server using the master device you specified is already running, and this is not the Server you want to start, determine the correct master device name for the Server you want to start.

If you cannot determine whether the master device is already in use, reboot the operating system.

Note If the master device resides on an operating system file system, you can copy the file to another name, modify your runserver file, and start Adaptive Server with that new file.

Memory Too Fragmented Error

Message text

```
Physical memory on this machine may be too fragmented
```

Note This error may be caused by a hardware problem.

Explanation

At startup, Adaptive Server allocates and initializes the memory needed for its structures and cache. This error occurs when Adaptive Server is not able to initialize memory needed for its structures and static cache.

The most common cause of this error is that the total memory required for the configurable parameters is more than the memory configured for Adaptive Server.

Action

1 At the operating system level, check memory use.

On UNIX machines, there are several commands for checking memory use. Following are examples of two of these commands:

```
% /etc/pstat -s

44076k allocated + 8920k reserved =52996k used, 40960k available
% vmstat
procs   memory           page              disk       faults      cpu
r b w avm fre  re at  pi po  fr de  sr s0 d1 d2 d3  in  sy  cs us sy id
4 0 0   0 336   0  1   0  0   0 16   0  1  0  0  0 181 260 399 84  2 14
```

If processes other than Adaptive Server are leaving memory fragmented, rebooting the operating system may resolve this error.

2 If you still cannot start Adaptive Server, refer to Chapter 1, "Adaptive Server Does Not Start After Altering Configuration" for instructions.

Start your Adaptive Server.

3 Once your Adaptive Server is running, use sp_configure either to reduce the value of the configurable parameters (such as number of user connections and number of locks) or to increase the value of the total memory configuration parameter. The total memory parameter controls the size of memory, in 2K units, that Adaptive Server allocates from the operating system.

   Refer to the *Adaptive Server Enterprise Performance and Tuning Guide* and "Configuring Memory" in the *System Administration Guide* for information about memory-related configuration parameters.

Additional information

Before calling Technical Support for assistance, have the following available:

- Server version and SWR version level
- Server error log
- Text of all the error messages
- Text of your configuration file or sp_configure output

Memory Usage in Procedure Headers

Message text

```
memory usage in procedure headers does not match memory
usage. count in pss (%d) for server process id %d.
```

Explanation

A process is a task being carried out by Adaptive Server. Processes can be in various states such as running or waiting for resources. For each user process and some system processes, Adaptive Server maintains a Process Status Structure (PSS) to track the state of the process and its use of resources. It also maintains an array of structures in the procedure header, each of which represents a physical page of memory in procedure cache.

When a process completes, Adaptive Server performs cleanup tasks such as freeing up resources used by the task. This error is reported when Adaptive Server cleans up after normal completion of a batch process, and determines that the memory usage tracked by the procedure header does not match the usage total maintained in the PSS.

This error may be due to an Adaptive Server problem.

Action

No action is required, since this is an informational message and does not result in any incorrect behavior. Upon detecting the memory count mismatch, Adaptive Server corrects the appropriate counter, completes the cleanup tasks, and continues processing.

Additional information

Before calling Technical Support for assistance, have the following available:

- Server version and SWR version level
- Server error log
- Text of all the error messages

Version in which this error is raised

11.5, 11.0.3.1 and earlier

Network Information Message

Message text

```
Using '%s' for network information.
```

Explanation

Adaptive Server uses information from the network addressing (*interfaces*) file located in the *$SYBASE* directory when it starts up and when it tries to establish a connection to a remote server. At those times, Adaptive Server displays this message, including the complete path of the interfaces file which was used.

Following are examples of this message for UNIX and OpenVMS:

UNIX

```
kernel: Using '/usr/u/sybase/interfaces' for network information.
```

OpenVMS

```
kernel: Using 'SYBASE_SYSTEM:[SYBASE]interfaces' for network information.
```

If a remote procedure call (RPC) is executed and no connection is currently established for that remote Server, Adaptive Server must first establish a connection to it, using a site handler process, before the RPC can be executed. Each time Adaptive Server establishes a site handler process, this informational message is written to the Server error log.

If the site handler process times out, then this message can occur many times for each remote Server. If the site handler for a remote Server does not time out, then Adaptive Server will write this message for the first RPC to that remote Server, and the message will occur only once for each remote Server.

Use sp_helpserver to check the status of the remote Adaptive Server. If the status column of the sp_helpserver output is “timeouts”, then you can use the following action to keep the network information message to one occurrence.

Action

To limit this message to one occurrence for a remote Server, disable the “timeouts” option to that remote Adaptive Server. Log into isql as “sa” and use sp_serveroption to set the “timeouts” option to “no timeouts”:

```
1> sp_serveroption remote_sql_server_name,
2> timeouts, false
3> go
```

This will keep the connection to the remote Server established. It will be maintained until either the local or remote Server is shut down. If you have few available user connections, remember that one connection per remote Adaptive Server connection will be in constant use.

Refer to the *Reference Manual* for information about sp_serveroption.

No More Alarms Available Error

Message text

```
uasetalarm: no more alarms available
```

Explanation

The Transact-SQL command waitfor defines a specific time, time interval, or an event for the execution of a statement block, stored procedure, or transaction. Adaptive Server uses alarms to carry out waitfor commands correctly.

The number of alarms needed by Adaptive Server is related to the requirements of an application and the number of instances of that application being run simultaneously. The maximum number of alarms available for use by Adaptive Server is controlled by the number of alarms configuration parameter.

When Adaptive Server needs more than the number of alarms available, the message “uasetalarm: no more alarms available” is written to the error log.

Action

1 Determine the current value of the configuration parameter number of alarms:

```
1> sp_configure "number of alarms"
2> go
```

2 Change the value of number of alarms to a value larger than the current value:

```
1> sp_configure "number of alarms", new_value
2> go
```

3 Shut down and restart Adaptive Server.

Note Each alarm structure uses 20 bytes of memory. If you raise the value of number of alarms significantly, readjust Adaptive Server's memory use accordingly.

Open Objects Parameter May Be Too Low

Message text

```
server: Warning: OPEN OBJECTS parameter may be too low;
attempt was made to free up descriptor in desfree().
Run sp_configure to increase parameter value.
```

or

```
server: Warning: OPEN OBJECTS or OPEN DATABASES
parameter may be too low; attempt was made to free up
descriptors in release_dbtable().  Run sp_configure to
increase parameter value.
```

Explanation

This message is a warning indicating that you do not have enough object descriptors or open databases available in the chain containing all the free descriptors.

This message is displayed in the Adaptive Server error log when Adaptive Server reaches the end of its array of allocated object descriptors, tries to recycle previously-used object descriptors and fails. It can also be displayed when Adaptive Server has fewer open databases configured than the number of databases on the Server.

Action

To recover from this error, increase the values of the configuration parameters number of open objects and number of open databases.

The configuration parameter number of open databases sets the maximum number of databases that can exist at one time on Adaptive Server.

The configuration parameter number of open objects sets the maximum number of objects that can be open at one time on Adaptive Server.

Increasing the value for number of open databases or number of open objects does not have a significant impact on performance or storage requirements, but it does require some memory. For more information about how Adaptive Server allocates memory, refer to the *Adaptive Server Enterprise Performance and Tuning Guide* and “Configuring Memory” in the *System Administration Guide*.

Use the system procedure sp_configure to display the run values of the configuration parameters you need to change:

```
1> sp_configure "number of open databases"
2> go
```

If the value for number of open databases is less than the value displayed by select count (*) from sysdatabases, increase the value of number of open databases. Otherwise, increase the value of number of open objects only.

Additional information

In Adaptive Server 11.5 and later , you can use sp_countmetadata to determine the current value of number of open objects and number of open databases. For example:

```
1> sp_countmetadata 'open objects'
2> go
There are 338 user objects in all database(s), requiring
165 Kbytes of memory. The 'open objects' configuration parameter
is currently set to 500.
1> sp_countmetadata 'open databases'
2> go
There are 5 databases, requiring 189 Kbytes of memory. The
'open databases' configuration parameter is currently set to 12.
```

os_attach_region Error

Message text

```
os_attach_region: shmat(%d): %s
```

Note This error may be caused by a hardware problem.

Explanation

Adaptive Server uses the function os_attach_region to attach to a shared memory region based on a shared memory identifier. This error occurs when Adaptive Server fails to attach to a region. In this message, *%d* is the shared memory identifier and *%s* is an operating system error message.

The message “os_attach_region: shmat(%d): %s” is raised on UNIX systems only. Other operating systems raise slightly different errors.

Action

An os_attach_region error can occur when the value of the operating system parameter *shmseg* is too small. For information about configuring shared memory properly, refer to the Adaptive Server installation and configuration guide.

This message can also be raised by Monitor Server. Refer to the *Monitor Server User's Guide* for details.

Additional information

Refer to the operating system man pages for the shmat() system call.

os_create_region Errors

Message text

```
os_create_region: shmget (0x%x): %s

os_create_region: Shared memory segment %d is in the way

os_create_region: uninitialized shared memory
descriptor

os_create_region: shmat (%d): %s

os_create_region: can't allocate %d bytes
```

Note This error may be caused by a hardware problem.

Explanation

Adaptive Server uses the following functions to manage shared memory:

- os_get_shmid – create a shared memory identifier
- os_create_region – create a region based on a shared memory identifier
- os_attach_region – attach to a region based on a shared memory identifier
- os_detach_region – detach from (and delete) the shared region
- os_format_shmid – format a shared memory identifier for printing

When os_create_region errors occur, Adaptive Server will not start.

The message texts shown here apply to UNIX systems only. Other operating systems raise slightly different errors.

os_create_region: shmget (0x%x): %s

This message is written to the error log when Adaptive Server fails to get a shared memory segment. In this message, *%x* is a shared memory key based on the shared memory identifier and *%s* is an operating system error message.

os_create_region: Shared memory segment %d is in the way

This error follows the shmget message and is also written to the Adaptive Server error log. A value of -1 for *%d* means the region does not exist.

os_create_region: uninitialized shared memory descriptor

During creation of a shared memory region, Adaptive Server attempts to validate the descriptor for the memory region. This message is written to the error log if the descriptor is found to be invalid.

os_create_region: shmat (%d): %s

This message is written to the error log when Adaptive Server fails to attach at an address. In this message, *%d* is the shared memory identifier and *%s* is an operating system error message.

os_create_region: can't allocate %d bytes

Adaptive Server was unable to allocate the number of bytes it requested for the shared memory region.

Action

1 At the operating system level, check which shared memory processes are using and whether shared memory segments are being used by Adaptive Server.

To check this on UNIX, run this command as the "sybase" user:

```
% ipcs -m
IPC status from workstation1 as of Fri May 26 14:08:25 1995
T     ID     KEY         MODE     OWNER    GROUP
Shared Memory:
m    256 0x699b7e24 --rw-------   sybase   sybase
m    257 0x699b7e25 --rw-------   sybase   sybase
```

If shared memory segments are being used by Adaptive Server, reboot the operating system to clear shared memory or remove them using the ipcrm operating system command.

Note Before removing the shared memory segments, identify the process that created them using the command "ipcs -ma" to make sure you only remove the appropriate segments.

2 Check the *$SYBASE* directory to determine whether there are any **.krg* or **.srg* files left from an abnormal Adaptive Server exit. If any such files exist, delete them.

3 os_create_region errors can occur when shared memory is not configured properly on your operating system. For information about configuring shared memory properly, refer to the Adaptive Server installation and configuration guide for your platform.

Shared Memory Error on Digital Unix

os_create_region: can't allocate %d bytes indicates that one or more kernel parameters needs to be reset. Logically, resetting shm-max should allow Adaptive Server to configure shared memory. However, other operating system kernel parameters also affect allocation. Consult your operating system documentation for details.

Additional information

Refer to the operating system man pages for the shget() and shmat() system calls.

Refer to the operating system man pages for ipcs and ipcrm.

Read/Write Error

Message text

```
kernel: write error on virtual disk %d, block %ld
```

or:

```
kernel: read error on virtual disk %d, block %ld
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server fails to read or write to the specified location on disk. This failure is usually a result of a physical disk problem.

If mirroring is enabled, you may also see the following errors:

```
udunmirror: i/o error on %s device '%s'
```

or

```
sddone: %s error on virtual device
```

Action

1 Identify the device with the problem by selecting the row from sysdevices which has the virtual disk number indicated in the error message:

```
1> select name, phyname from master..sysdevices
2> where low/power(2,24) = virtual_disk_number
3> go
```

The output from this query should provide the physical name of the damaged disk. Shut down Adaptive Server and examine the disk as soon as feasible and correct any problems found to prevent further corruption. Then restart Adaptive Server.

Note The sp_diskblock procedure described below also identifies the device.

2 Use the procedure described below to determine which database (if there is more than one database on the device), table, or index is affected by this error. Use this information to assess the potential severity of the problem quickly and decide on an action which is appropriate for your operating environment.

Use the sp_diskblock procedure supplied in Chapter 2, "How to Gather Information About Read/Write Errors". The syntax of this procedure is:

```
sp_diskblock virtual_disk, block_number
```

For example, if the read/write error message displays a virtual disk number 4 and a block number 871, the procedure command is:

```
1> sp_diskblock 4, 871
2> go
Virtual disk 4, block 871 corresponds to:
Logical page 1895 in the "production" database (dbid=4) on device "main".
```

Use the dbcc page command to determine which object is using that page. Refer to Chapter 2, "How to Find an Object Name from a Page Number" for details on this procedure.

Additional information

Refer to the writeups for *dopen Errors* and *udunmirror Errors* in this manual for information about dopen, sddone, and udunmirror errors.

Server Is Unnamed

Message text

```
server is unnamed
```

Explanation

This informational message is displayed during Adaptive Server start-up if the Adaptive Server has not been given a name via the sp_addserver system procedure. The Adaptive Server name is stored in the sysservers system table as well as in the *@@servername* global variable. This name is used primarily for managing Server-to-Server communication and does not have to match the Server name used either in the *interfaces* file or in the runserver file.

Action

No action is required for this informational message. However, the Adaptive Server administrator can prevent this message by assigning a name to the Adaptive Server via the sp_addserver procedure. Refer to the *Reference Manual* for information about restrictions on valid Adaptive Server names with sp_addserver.

To eliminate the message, assign a name to Adaptive Server using the sp_addserver command:

```
1> sp_addserver sql_server_name, local
2> go
```

The *@@servername* global variable will not reflect the name until Adaptive Server is restarted. All subsequent starts of the Adaptive Server display the new name during Adaptive Server start-up. For example, for an Adaptive Server named “MY_SERVER”, the message would be:

```
server: server name is 'MY_SERVER'
```

Although this Adaptive Server name is associated with Server-to-Server communication, specifying a name does not in itself allow this communication and therefore does not represent a security risk. Naming all Adaptive Servers in this way makes the Adaptive Server error logs easy to identify, because the Adaptive Server name is contained within the log.

Stack Guardword Error

Message text

```
*** Stack guardword corrupted.
```

Explanation

This error occurs when the stack area used by an Adaptive Server process is corrupted. Adaptive Server usually shuts itself down after this error to avoid problems that could arise from this stack corruption.

Warning! If this error occurs and Adaptive Server remains up, shut it down immediately!

At start-up, Adaptive Server allocates one stack area for every configured user connection. These stacks are in contiguous areas of memory, with a guard at the end of each stack. At the end of each stack guard area is a “guardword,” which is a 4-byte structure with a pattern. Adaptive Server periodically checks this pattern to determine whether it has changed. A change indicates that a process has overflowed its stack guard area. When this occurs, Adaptive Server prints the following messages in the error log and then shuts down:

```
kernel: *** Stack overflow detected: limit: 0x%lx sp: 0x%lx
kernel: *** Stack guardword corrupted
kernel: *** Stack corrupted, server aborting
```

In the first message, *limit* is the address of the end of the stack guard area, and *sp* is the current value of the stack pointer.

In addition, Adaptive Server periodically checks the stack pointer to determine whether it is completely outside both the stack and the stack guard area for the process. If it is, Adaptive Server shuts down, even if the stack guardword is not corrupted, and prints the following messages in the error log:

```
kernel: *** Stack overflow detected: limit: 0x%lx sp: 0x%lx
kernel: *** Stack corrupted, server aborting
```

Stack corruption is usually the result of one process's stack overflowing its boundaries and writing on top of another process's stack. In general, an Adaptive Server process writes past the end of its stack area for one of two reasons:

- The stack size is too small. The information placed on the stack is valid but its total size exceeds the specified stack size. This is often caused by a single, complex query which requires more than the average amount of stack space.
- An incorrectly handled error results in recursive calls which eventually fill the stack, regardless of its size.

Action

If this error occurs because a complex query has run out of stack area, correct the error by using one of the two following methods:

Solution One: Break Up the Complex Query

Break up the query into smaller pieces, if possible.

Solution Two: Increase Stack Size

Use the following information to select an appropriate size for the stack size configuration parameter:

- The stack size should always be an even increment of 2K (2048 bytes). If a stack size is specified that is not divisible by 2K, Adaptive Server rounds it up to the next 2K increment.
- Increase the stack size by no more than 2K at a time. If the stack size is too large, you may be unable to start Adaptive Server.
- Increasing the stack size results in Adaptive Server requiring more memory for each configured user connection and device.

Use sp_configure to increase the stack size configuration parameter as follows:

```
1> sp_configure "stack size", new_value
2> go
```

Once this change is made, shut down and restart Adaptive Server so that the change will take effect.

Warning! If you increase the value of the stack size configuration parameter, you may need to adjust other memory-related parameters such as total memory. Refer to the *Adaptive Server Enterprise Performance and Tuning Guide* and "Configuring Memory" in the *System Administration Guide* for details.

Additional information

Refer to "Setting Configuration Parameters" in the *System Administration Guide* and for information about sp_configure and configuration parameters.

t_rcvconnect Error

Message text

```
nconnect: t_rcvconnect, %s
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when an Adaptive Server using a TLI-based interface tries to establish a client connection and an operating system error message is returned. This can happen when Adaptive Server tries to execute an RPC on a Server that is not running (Adaptive Server, Open Server, or Backup Server).

When this error occurs, the connection is not established and the message "nconnect: t_rcvconnect, %s" is written to the error log. *%s* is the text of the operating system error that caused the failure.

This is not a serious error unless it occurs frequently.

Action

Resolve the problem that caused the operating system error contained in the Adaptive Server error message. You may need assistance from an operating system or network administrator. Refer to the operating system man pages for information about the t_rcvconnect() system call.

Check the Adaptive Server error log. If other errors exist, refer to the writeups in this manual for information about recovering from them.

If, after resolving the operating system problem and the problems found in the error log, you still get the message “nconnect: t_rcvconnect, %s,” call Sybase Technical Support.

Additional information

Before calling Technical Support, have the following information available:

- Server version and SWR version level
- Server error log
- Operating system error log
- sp_helpserver and sp_configure output
- Contents of *interfaces* file

udunmirror Errors

Message text

```
udunmirror: i/o error on primary device '%s'
udunmirror: i/o error on secondary device '%s'
udunmirror: failing over to '%s'
```

Explanation

This error occurs when Adaptive Server detects an I/O error on the primary or secondary device of a mirrored pair of devices. It is usually a result of a physical disk problem.

I/O Error Detected on Primary Device

If an I/O error is detected on the primary device, Adaptive Server disables mirroring. sysdevices reflects information for the device as if a disk unmirror with side = primary and mode = retain was issued. Adaptive Server does this automatically and writes kernel messages like the following to the error log:

```
kernel: udunmirror: i/o error on primary device 'primary_device'
kernel: DataServer i/o to the device will be disabled
kernel: udunmirror: failing over to 'secondary_device''
```

I/O Error Detected on Secondary Device

If an I/O error is detected on the secondary device, Adaptive Server disables mirroring and sysdevices reflects information for the device as if a disk unmirror with side = secondary and mode = retain was issued.
Adaptive Server does this automatically and writes kernel messages like the following to the error log:

```
kernel: udunmirror: i/o error on secondary device  '/usr/u/sybase/test.mir'
kernel: DataServer i/o to the device will be disabled
```

Action

1 Using your operating system procedures, determine whether the device is damaged. If it is damaged, repair or replace it.

If the device is not damaged but it was unavailable for some other reason, make it available again.

2 If you use operating system files for Sybase devices or if the name of the device was changed, drop the mirror and re-create it:

```
1> disk unmirror name="device_name",
2> side={primary|secondary}, mode=remove
3> go

1> disk mirror name="device_name",
2> mirror="full_mirror_name"
3> go
```

where side is the device that was temporarily disabled.

Note On a UNIX system, you can also use the operating system command touch to create the file and then issue a disk remirror command.

3 If you did not need to drop and re-create the mirror, use the disk remirror command to make the device accessible to Adaptive Server:

```
1> disk remirror name = "device_name"
2> go
```

Note You can use the disk remirror command if Adaptive Server automatically disabled mirroring or if you issued a disk unmirror with mode = retain command. If you issued a disk unmirror with mode = remove command, you can only reinstate disk mirroring by issuing the disk mirror command.

Additional information

Refer to the *Reference Manual* for information about the disk mirror, disk remirror, and disk unmirror commands.

ueoffline: no more engines can be offlined

Message text

```
ueoffline: No more engines can be offlined in order to
support %d outstanding network connections. Please try
later when the server has fewer connections.
```

Explanation	As of Version 12.0, Adaptive Server allows SAs to bring down server engines by using the dbcc engine(offline) command. Along with the existing dbcc engine(online) command, this makes it possible for SAs to dynamically optimize the number of engines on a server based on the work load. When you execute dbcc engine(offline), Adaptive Server checks whether the remaining engines can accommodate the outstanding user connections. This error is raised if taking down the engine would make it impossible to support the current connections to the server.

Action	Wait until there are fewer active connections, then re-try the dbcc engine(offline) command.
	
	Use sp_who periodically to determine the number of active user connections on your server. Use the following command to determine the upper limit on the number of file descriptors available for server connections:
	
	`1> select @@max_connections`
	`2> go`
	
	The return value represents the maximum possible descriptors available. Take into account overhead such as site handler connections, and note that overhead increases with the number of engines.

Version in which this error is raised	12.0 and later

Index

Symbols

Numerics

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W