

About This Manual

This manual describes the status codes and messages you may encounter while using components of the Scalable SQL relational data access system or the Btrieve navigational database management system. An explanation or corrective action accompanies each message or code.

The status codes and messages in this manual are not specific to any one platform or product. This manual may refer to utilities, products, or platforms that are not part of your product distribution.

Pervasive Software would appreciate your comments and suggestions about this manual. As a user of our documentation, you are in a unique position to provide ideas that can have a direct impact on future releases of this and other manuals. Please complete the User Comments form that appears on our Web site and fill in part number 100-003038-007.

Send comments to docs@pervasive.com. Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

Who Should Read This Manual?

This manual provides information for developers, testers, and support personnel who install and run Pervasive.SQL 7 server and workstation products. Status codes may be displayed during product installation, while using the utilities, or while using an application.

Pervasive Software would appreciate your comments and suggestions about this manual. As a user of our documentation, you are in a unique position to provide ideas that can have a direct impact on future releases of this and other manuals. Please complete the User Comments form that appears on our Web site and fill in part number 100-003038-007.

Send comments to docs@pervasive.com. Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

Manual Organization

The following list briefly describes each chapter in the manual:

- [Chapter 1—“Status Codes”](#)

This chapter lists and explains each status code you may receive. A brief description follows each code. The status codes are listed in numeric order. Some status codes listed in this chapter refer to engines and utilities for environments other than your own. Status codes are returned by API calls.

- [Chapter 2—“Messages”](#)

This chapter lists and explains the numbered and non-numbered text message you may receive. The messages are categorized by the application or component that generates them. Some messages listed in this chapter refer to engines and utilities for environments other than your own.

- [Appendix A—“Changes Since Btrieve 6.15”](#)

This appendix contains a list of status codes that have been supplemented in Pervasive.SQL with more specific codes.

This manual also includes an index.

Send comments to docs@pervasive.com. Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

Conventions

Unless otherwise noted, command syntax, code, and code examples use the following conventions:

- Case Commands and reserved words typically appear in uppercase letters. Unless the manual states otherwise, you can enter these items using uppercase, lowercase, or both. For example, you can type MYPROG, myprog, or MYprog.

- [] Square brackets enclose optional information, as in [*log_name*]. If information is not enclosed in square brackets, it is required.

- | A vertical bar indicates a choice of information to enter, as in [*filename* | @*filename*].

- < > Angle brackets enclose multiple choices for a required item, as in /D=<5|6|7>.

- variable* Words appearing in italics are variables that you must replace with appropriate values, as in *filename*.

- ... An ellipsis following information indicates you can repeat the information more than one time, as in [*parameter* ...].

- ::= The symbol ::= means one item is defined in terms of another. For example, a::=b means the item *a* is defined in terms of *b*.

Status Codes

This chapter describes the status codes and messages generated by Scalable SQL, Btrieve, or any of their components.

Status codes are not the same as numbered messages (see [Chapter 2, "Messages"](#)); messages are returned to end users by utilities or specific components and generally begin with a prefix and a number. For example: ["MKDE-16: There is insufficient memory to load the MicroKernel Database Engine."](#)

Status codes are returned to an application making a Btrieve or Scalable SQL API call. The application determines whether or not to display the status code to the user. Applications will often display status codes in addition to helpful messages about a problem that has occurred, but the application developer determines this. Many utilities provided with Btrieve and Scalable SQL use the Btrieve and Scalable SQL APIs, and they may therefore return status codes as well as utility-specific messages.

Some messages include references to status codes.

The status codes appear in numeric order. [Table 1-1](#) lists the numeric ranges for each type of code.

Table 1-1
Status Code Ranges

Range	Type of Code
-199 to -100	Informative Status Codes
0	Successful API Completion Status Code
1 to 199	MicroKernel Database Engine Status Codes
200 to 999	Scalable SQL Engine Status Codes
1000 to 1999	MicroKernel Status Codes for Windows, OS/2, and DOS Workstations
2000 to 2099	Btrieve Requester Status Codes
2100 to 2199	Scalable SQL Interface Component Status Codes
2200 to 2299	Scalable SQL Miscellaneous Status Codes
2300 to	Database Names Functions Status Codes

2399

2900 to 2999 [Inscribe Status Codes](#)

3000 to 3099 [MicroKernel Router Status Codes](#)

3100 to 3199 [Network Services Layer Status Codes](#)

7000 to 7099 [User Count Manager Status Codes](#)

8000 to 8499 [Component Management Status Codes](#)

8500 to 8589 [ECAS API Status Codes](#)

8590 to 8599 [W3DBSMGR Status Codes](#)



Note: This chapter may refer to utilities, products, or platforms that are not part of your product distribution.

-199 to -100

Informative Status Codes

This section lists the informative status codes that Scalable SQL can return. Scalable SQL returns these codes as negative values.

-101: The SET statement completed successfully.

The following statements return this status code when they execute successfully:

SET BINARYNULL	SET BLANK (in v3.01 mode)	SET CHAR
SET DATAPATH	SET DDPATH (in v3.01 mode)	SET DECIMALNULL
SET DEFAULT	SET ISOLATION	SET MASK
SET OPENMODE	SET OWNER	SET RANGE
SET SECURITY	SET STRINGNULL	SET VALUES

Scalable SQL made the change you requested. However, if you issued a SET OPENMODE, SET OWNER, or SET ISOLATION statement during a transaction, the change does not take effect until you start a new transaction.

-102: The INSERT statement completed successfully.

Scalable SQL added the specified rows to the table(s).

-103: The UPDATE statement completed successfully.

Scalable SQL made the specified changes to the table(s).

-104: The DELETE statement completed successfully.

Scalable SQL deleted the specified rows from the table(s).

-105: The CREATE statement completed successfully.

The following statements return this status code when they execute successfully:

CREATE DICTIONARY	CREATE TABLE
CREATE GROUP	CREATE TRIGGER
CREATE INDEX	CREATE VIEW

CREATE PROCEDURE

If you issued a CREATE DICTIONARY statement, Scalable SQL successfully created the X\$File, X\$Index, and X\$Field system tables. Otherwise, Scalable SQL successfully added the group, index, stored procedure, table, trigger, or view to the data dictionary.

-106: The ALTER TABLE statement completed successfully.

Scalable SQL successfully made the specified change to the table's dictionary definition. If you altered a column (including a primary key or foreign key) without specifying the IN DICTIONARY keyword, Scalable SQL also changed the data file.

-107: The DROP statement completed successfully.

The following statements return this status code when they execute successfully:

DROP DICTIONARY	DROP TABLE
DROP GROUP	DROP TRIGGER
DROP INDEX	DROP VIEW
DROP PROCEDURE	

If you specified DROP DICTIONARY, Scalable SQL deleted the dictionary files and all associated data files. Otherwise, Scalable SQL successfully removed the group, index, stored procedure, table, trigger, or view from the dictionary. (Dropping a table also deletes the data file for that table.)

-108: The statement contains unresolved substitution variables and cannot be executed.

The current SQL statement contains substitution variables; you cannot execute the statement until you supply values for each variable. If you are using an application that allows substitution variables, refer to the documentation included with the application to determine how to use them with that application. If you are a developer using the Scalable SQL APIs to execute the statement, perform the following steps:

1. Supply values for the variables using the XQLSubst function.
2. Use the XQLExec function to verify the substitution values and, if the statement is not a SELECT statement, to execute the statement.
3. If the statement is a SELECT statement, use the XQLFetch function to fetch the data.

-109: The view contains no more records.

You have reached either the beginning or the end of the view based on the current XQLFetch or xFetch option code.

–110: The GRANT statement completed successfully.

Scalable SQL assigned the specified rights to the user or group.

–111: The REVOKE statement completed successfully.

Scalable SQL revoked the specified rights from the user or group.

–112: The START TRANSACTION statement completed successfully.

Scalable SQL has begun a transaction. All subsequent statements you issue are part of this transaction until you issue either a COMMIT WORK or a ROLLBACK WORK statement. For a savepoint, the SAVEPOINT label remains in effect until you explicitly release or roll back to that label, or until the end of any outer transaction within which the savepoint is nested.

–113: The COMMIT WORK statement completed successfully.

Scalable SQL committed the changes made by your transaction to the data tables. You can no longer undo the changes with a ROLLBACK WORK statement. For a RELEASE SAVEPOINT statement, any changes made since the savepoint was declared can no longer be rolled back separately. They can only be committed or rolled back as part of an outer transaction.

–114: The ROLLBACK WORK statement completed successfully.

Scalable SQL reversed the changes you made during the transaction *except* for any changes you made with operations that are not affected by transaction processing. If you perform one of the following operations within a transaction, Scalable SQL completes the operation, but you cannot roll back the results:

- Operations that create or change dictionary definitions. Therefore, you cannot roll back the results of the following statements: ALTER TABLE, CREATE DICTIONARY, CREATE GROUP, CREATE INDEX, CREATE PROCEDURE, CREATE TABLE, CREATE TRIGGER, and CREATE VIEW.
- Operations that remove dictionary definitions. Therefore, you cannot roll back the results of the following statements: DROP DICTIONARY, DROP GROUP, DROP INDEX, DROP PROCEDURE, DROP TABLE, DROP TRIGGER, and DROP VIEW.
- Operations that grant or revoke security rights. Therefore, you cannot roll back the results of the following statements: GRANT (access rights), GRANT CREATETAB, GRANT LOGIN, REVOKE (access rights), REVOKE CREATETAB, and REVOKE LOGIN.

For a ROLLBACK TO SAVEPOINT statement, any changes made since the savepoint was declared are rolled back.

–115: Scalable SQL successfully recalled the stored SQL statement.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

Scalable SQL successfully recalled the specified stored statement in the XQLCompile statement buffer, indicating that you can now execute the stored statement. Perform the following steps:

1. Execute the statement using the XQLExec function.

Scalable SQL executes each statement within the stored statement until it finds either a SELECT statement or no more statements.

2. If Scalable SQL finds a SELECT statement, use the XQLFetch function to fetch the data.

3. If you just retrieved data with XQLFetch, use XQLExec to continue executing the statements in the stored statement.

–116: Scalable SQL successfully executed the stored SQL statement.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode. In Scalable SQL v4.0 and later, stored procedures are used instead of stored statements.

Scalable SQL executed the stored statement that you submitted for execution.

–118: The DECLARE statement completed successfully.

The following statements return this status code when they execute successfully:

```
DECLARE CURSOR           DECLARE VARIABLE
                           DECLARE CONDITION
```

If you issued a DECLARE statement, Scalable SQL successfully created the cursor, variable, or condition you defined.

–119: The SET assignment statement completed successfully.

The SET *variable* = *expression* statement returns this status code upon successful completion.

–120: The OPEN statement completed successfully.

The OPEN CURSOR statement returns this status code upon successful completion.

–121: The CLOSE statement completed successfully.

The CLOSE CURSOR statement returns this status code upon successful completion.

–122: The FETCH statement completed successfully.

The FETCH statement returns this status code upon successful completion.

–123: The positioned UPDATE statement completed successfully.

The UPDATE [*table_reference*] SET *set_clause* WHERE CURRENT OF statement returns this status code upon successful completion.

124: The positioned DELETE statement completed successfully.

The DELETE [FROM *table_reference*] WHERE CURRENT OF *cursor_name* statement returns this status code upon successful completion.

–125: The CALL stored procedure statement completed successfully.

The CALL statement returns this status code upon successful completion. This status code indicates the successful

completion condition for the stored procedure.

Send comments to docs@pervasive.com. Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

0

Successful API Completion Status Code

Any API call can return the following status code.

0: The API completed successfully.

This status code is returned for any API call that completes successfully. If an operation is not successful, a positive, nonzero status code is returned.

Send comments to docs@pervasive.com. Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

1 to 199

MicroKernel Database Engine Status Codes

This section describes status codes that the MicroKernel returns.

1: The operation parameter is invalid.

The specified Btrieve operation does not exist or is not valid for this version of the MicroKernel.

2: The application encountered an I/O error.

An error occurred while reading from or writing to the disk. One of the following has occurred:

- The file is damaged, and you must recover it. Refer to the *Pervasive.SQL User's Guide* for more information about recovering a file.
- In NetWare, at least one MicroKernel data file is flagged as Shareable. All MicroKernel data files should be flagged as Non-Shareable. The MicroKernel is the only user that accesses the files (on behalf of many users). As the only user, the MicroKernel can control the integrity of the data files. If you flag your MicroKernel data files as Shareable, data corruption can occur.
- For pre-v6.0 data files, there is a large pre-image file inside a transaction, and there is not enough disk space for a write to the pre-image file.
- For pre-v6.0 data files, there is one pre-image file for multiple data files. For example, if you name the data files CUSTOMER.ONE and CUSTOMER.TWO, both files have pre-image files named CUSTOMER.PRE.
- For pre-v6.0 data files that are larger than 768 MB, there is a conflict among locking mechanisms. The file has not been corrupted. Your application can retry the operation until the conflict is resolved (when the competing application releases the lock your application requires).
- A pre-v6.0 Btrieve engine attempted to open a v6.x or later MicroKernel file.

3: The file is not open.

The operation cannot execute because the file is not open. The application must perform a successful Open operation before the MicroKernel can process any other operations. The MicroKernel also returns this status code if the application passed an invalid position block for the file, or if the application passed a position block with a client ID other than the client ID used to open the file.

This status code has been supplemented in Pervasive.SQL 7. For a list of the new status codes, see "[Status Code 2003 replacement](#)".

4: The application cannot find the key value.

The MicroKernel cannot find the specified key value in the index path. When you receive this status code on an Update or Delete operation, it usually means that the file is damaged, and you must recreate it. Occasionally, a corrupt key can cause this status code. Drop the key, then add it again. Also, pre-v6.x server engines return this status code if two separate files have different ACSs, but those ACSs have the same name. Never use the same name for different ACSs, regardless of the version you are using.

5: The record has a key field containing a duplicate key value.

The MicroKernel cannot add or update a record because the record has a key field that contains a duplicate key value for an index that does not allow duplicate values. The MicroKernel also returns this status code when it cannot

create an index that does not allow duplicate key values because a duplicate key value already exists.

In pre-v6.0 MicroKernels, this status code indicates that you added an AUTOINCREMENT key and the field being indexed contains more than one zero.

6: The key number parameter is invalid.

The value stored in the key number parameter is not valid for the file being accessed. The key number must correspond to one of the keys defined for the file. Valid key numbers are 0 through 118.

7: The key number has changed.

The key number parameter changed before a Get Next, Get Next Extended, Get Previous, or Get Previous Extended operation. The operation requires the same key number parameter as the previous operation, because the MicroKernel uses positioning information relative to the previous key number.

In a related situation, the MicroKernel returns this status code when an application performs a Delete or Update operation immediately following a Get operation. If the application changes the value of the key number in the Delete or Update operation (from the value used with the preceding Get operation), the MicroKernel deletes or updates the record as requested and does *not* return this status code, at least not at this point. However, the MicroKernel *does* return this status code on the very first Get Next, Get Next Extended, Get Previous, or Get Previous Extended operation performed after the deletion or update, even if that Get operation uses the same key value the application passed to the Delete or Update operation.

If you need to change key numbers between consecutive Get Next, Get Next Extended, Get Previous, or Get Previous Extended operations (or in Delete or Update operations as described in the preceding paragraph), use a Get Position operation followed by a Get Direct/Record operation to reestablish positioning for the new index path.

8: The current positioning is invalid.

You must establish the current position in order to update or delete a record. Perform a Get or Step operation to establish the current position. The MicroKernel also returns this status code if the application passed an invalid position block for the file.

9: The operation encountered the end-of-file.

The MicroKernel returns this status code in one of the following situations:

- The operation encountered an end-of-file boundary or tried to read past a file boundary (end-of-file or start-of-file).
- In a Get Next Extended, Get Previous Extended, Step Next Extended, or Step Previous Extended operation, the number of records satisfying the filtering condition is less than the number of specified records to be returned, and the reject count or filter limit has not been reached.
- When reading a file in ascending order according to an index path, the MicroKernel has already returned the last record in that index path. When reading a file in descending order according to an index path, the MicroKernel has already returned the first record in the index path.
- When using the Get By Percentage operation, either the value supplied for the percentage is too high—it exceeds 10,000 decimal (0x2710)—or the file contains no records.

10: The key field is not modifiable.

During an Update operation, the application attempted to modify a key field that is defined as nonmodifiable.

11: The specified pre-imaging filename is invalid.

This status code indicates either that the specified filename does not conform to file naming conventions, or that the pathname is invalid. Make sure the filename or pathname is valid for the environment. This status code can also have the following meanings in client/server versions of the MicroKernel:

- The application attempted to open a file that has .^^^ as its extension. This extension is reserved for the MicroKernel to use during continuous operation. (Only server engines can use continuous operation.)
- The data buffer for a Begin or End continuous operation is not set up correctly.
- You attempted to load a remote file when your Communications Requester settings of Local and Requester are incorrectly set to Yes and No, respectively. To resolve this condition, run the appropriate Setup utility and change the Requester setting to Yes. For more information on the Setup utility, refer to the *Pervasive.SQL User's Guide*.

12: The MicroKernel cannot find the specified file.

Check the key buffer parameter to make sure the pathname is terminated with a blank or a binary zero. Also, ensure that the file exists. When accessing a file on a server, ensure that you have FILE SCAN rights to the directory in which the file resides. Occasionally, a corrupt key can cause this status code. Drop the key, then add it again.

13: The MicroKernel could not open the extension file for an extended file.

The MicroKernel could not open the extension file for an extended file that the application tried to open. For v7.0 and later MicroKernels, an extended file can consist of a base file and up to 15 extension files. Extension files must remain in the same volume and directory as their base files. The MicroKernel returns this status code if you delete, move, or rename the extension files.

For pre-v6.0 MicroKernels, extension files must be loaded on the logical disk drive specified when the extension was created. Both the primary file and its extension file must be on-line to access an extended file.

14: The MicroKernel cannot create or open the pre-image file.

The MicroKernel uses pre-image files only for pre-v6.0 data files.

The MicroKernel returns this status code in one of the following situations:

- The MicroKernel cannot create a new pre-image file because the disk directory is full. The MicroKernel must be able to create a pre-image file.
- The MicroKernel cannot open the pre-image file to restore file integrity. If the pre-image file is erased or damaged, the MicroKernel cannot restore the file's integrity. Refer to the *Pervasive.SQL User's Guide* for more information about recovering damaged files.
- The workstation MicroKernel cannot assign a handle to the pre-image file because the MicroKernel was not started by a user with access rights to the pre-image file.
- The file structure of a pre-image file created by a v6.0 or later MicroKernel is different from the file structure of a pre-image file created by a v5.x MicroKernel. If you have an extraneous .PRE file in v5.x format and you are using a v6.0 or later MicroKernel, the MicroKernel returns this status code when you try to open the data file to which the .PRE file belongs.

15: The application encountered an I/O error during pre-imaging.

The MicroKernel uses pre-image files only for pre-v6.0 data files.

This status code indicates that either the disk is full or the pre-image file is damaged. If you receive this status code, proceed as follows:

- If the pre-image file is damaged, the integrity of the data file cannot be ensured. Refer to the *Pervasive.SQL User's Guide* for more information about recovering damaged files.
- If the disk is full, erase any unnecessary files.

16: The application encountered an expansion error.

This status code is obsolete in MicroKernel versions 6.0 and later.

The MicroKernel returns this status code when it encounters an error while writing the directory structure to disk prior to creating the expanded file partition. Either the MicroKernel cannot close the file, or a new page was added to the file and the MicroKernel cannot close and reopen the file to update the directory structure. Check for a disk hardware failure.

17: The application encountered a close error.

This status code is obsolete in MicroKernel versions 6.0 and later.

The MicroKernel returns this status code when it encounters an error while writing the directory structure to disk prior to closing the file. Either the MicroKernel cannot close the file, or a new page was added to the file, and the MicroKernel cannot close and reopen the file to update the directory structure. Check for a disk hardware failure. The MicroKernel also returns this status code if the application passed an invalid position block for the file.

18: The disk is full.

The MicroKernel can return this status code in the following situations:

- The disk is full, and the MicroKernel cannot expand the file to accommodate additional records. Erase any unnecessary files. If using a pre-v6.0 MicroKernel, you can possibly extend the file to gain additional disk space.
- There is not enough space to append a new page to the data file.
- The pre-image file is out of disk space. If your files are in pre-v6.0 format and you are in a transaction, the pre-image file size increases for the duration of the transaction. If you receive this status code, either reduce the number of operations in the transaction, or obtain more disk space.
- For files located on a NetWare server, the NetWare owner name for the file is no longer valid, and your application tried to insert or update records in the file, thus causing the file to expand. In this case, the MicroKernel returns this status code when it needs to add a page to the file, regardless of how much disk space is available. To check for an owner name, use the NetWare utility NDIR. To add an owner name, use either FILER (a NetWare text utility) or the NetWare Administrator graphical utility.
- In some environments, you can restrict the amount of disk space available to each user. This status code indicates that the application attempted to expand a data file beyond the amount of disk space allocated to the file's owner.
- You tried to read or modify a file which was not closed properly after a disk full error. Make sure that every application that used the file at the time of the disk full error closed the file successfully.
- If a client connected to a Pervasive server encounters Status 18, other clients performing read-only operations from the same disk may also receive a non-zero status (18 or 2).

19: The application encountered an unrecoverable error.

To ensure file integrity, recover the file as described in the *Pervasive.SQL User's Guide*.

20: The MicroKernel or Btrieve Requester is inactive.

Verify that the IPX/SPX or TCP/IP protocol is properly installed at the client machine.

Verify that no two machines on the network have the same Internal Network Number.

In the DOS and NetWare environments, you must load the MicroKernel and, if applicable, the Btrieve Requester before generating any requests. Also for NetWare, ensure that the appropriate communications modules are loaded at the server.

If you are running an application in a client/server environment and also need to access files located on a local drive, ensure that you enable both Local and Server access and that the local MicroKernel is available.

By default in the Windows and Windows NT environments, the Setup utility enables the Local Engine Usage option. To access a server engine but not a workstation engine, turn the Local option off and turn the Server option on. To access both workstation and server MicroKernels, enable both the Server option and the Local option.

Also, in the Windows environment, ensure that the Btrieve for Windows DLLs and MicroKernel executable are in your path or in the Windows system directory. Ensure that you do not have multiple copies of WBTICOMM.DLL or WBTRCALL.DLL installed on the same machine.

In the Windows NT server environment, you must start the MicroKernel before generating any requests. Ensure that the Windows NT DLLs are in your path. Also ensure that the appropriate communications modules are loaded at the server.

In the OS/2 environment, ensure that the necessary DLLs are in your LIBPATH directory.

This status code has been supplemented in Pervasive.SQL 7. For a list of the new status codes, see ["Status Code 20 replacements"](#).

21: The key buffer parameter is too short.

The key buffer parameter is not long enough to accommodate the key field for the requested index path. Verify that the length of the key buffer equals the defined length of the key specified in the key number parameter.

22: The data buffer parameter is too short.

The data buffer parameter specified by the application was not large enough to accommodate either the minimum length of the record for an Insert or Update operation, or the entire record length for a Get or Step operation. Also, the data buffer may not be large enough to accommodate the length of data required for operations such as Create, Create Index, Stat, Get By Percentage, Find Percentage, or Version.

- For Get or Step operations, the MicroKernel returns as much data as it can and this status code, indicating that it cannot return the entire record.
- For an Insert operation, the MicroKernel does not insert the record if the data buffer is shorter than the fixed-length portion of the record.
- For an Update operation, if the data buffer is too short to contain the fixed-length portion of a record, the MicroKernel does not update the record.
- For the Create, Stat, and Create Index operations, the data buffer is not long enough to contain all the file specifications, the key specifications, and (if specified) the ACS definition.
- For the Get by Percentage or Find Percentage operation, the data buffer length is less than 4 bytes.

- For the Version operation, the data buffer length is less than 5 bytes.

This status code can also indicate a corrupt file if the file allows variable-length records and you receive this status code on a Get or Step operation. In such a corrupt file, you can receive Status Code 54 when you use Get or Step operations to read other records. Recover the file according to the instructions in the *Pervasive.SQL User's Guide*.

23: The position block parameter is not 128 bytes long.

This status code is obsolete in Btrieve language interface versions 6.0 and later.

The position block parameter must be exactly 128 bytes long.

24: The page size or data buffer size is invalid.

The MicroKernel returns this status code in one of the following situations:

- The page size you specified when creating a file is invalid. The page size must be a multiple of 512 bytes and cannot exceed 4096 bytes.
- During a Create operation, the page size is the first file specification the MicroKernel checks. A Status Code 24 at this point can indicate an invalid data buffer parameter.

In pre-v6.1 MicroKernels, the Open operation can return this status code. In this case, the MicroKernel cannot open the file because the file's page size exceeds the Largest Page Size configuration option. To successfully open the file, you must increase the value of the Largest Page Size configuration option and then reload the MicroKernel. The v6.1 or later MicroKernel does not return this status code from the Open operation.

25: The application cannot create the specified file.

The MicroKernel can return this status code if an application attempts to create a data file, but the disk directory or the disk itself is full. If the application is creating a file over an existing file, the MicroKernel returns this status code when the existing file is open or when the operating system prevents the operation for some other reason (for example, because the file is flagged transactional in NetWare).

Sometimes pre-v6.0 MicroKernels can return this status code if the HOLD parameter in NET.CFG or SHELL.CFG is set to ON and the application attempts to create a data file on a network drive. (The HOLD parameter is set to OFF by default.) Creating the file on a local drive is successful regardless of the value of the HOLD parameter.

26: The number of keys specified is invalid.

The number of keys specified for the page size is invalid. The number of key segments must be within the following limits:

Page Size	512	1024	1536	2048	2560	3072	3584	4096
Max. No. Key Segments	8	23	24	54	54	54	54	119

If the MicroKernel is configured to create files in v5.x format, the maximum number of key segments is 8 for files using 512 byte page sizes; the maximum number of key segments for all other v5.x files is 24.

27: The key position is invalid.

The specified key field position is less than 1 or exceeds the defined record length for the file. Either the key position is greater than the record length or the key position plus the key length exceeds the record length.

28: The record length is invalid.

The specified record length (plus overhead for duplicates, record usage count, variable record pointers, key pointers, and blank truncation information) must be less than or equal to the page size minus 6 bytes, and greater than or equal to 6 bytes.

For key-only files, the maximum record length is 253 bytes (255 bytes for a pre-v6.0 file).

29: The key length is invalid.

The specified key length must be greater than 0 but cannot exceed 255 bytes. The length of a key segment must agree with its key type if the key type implies a length (for example, an integer key must have a length evenly divisible by two). Each key page in the file must be large enough to hold at least eight keys. If the page size is too small to accommodate eight occurrences of the specified key length (plus overhead), either increase the file's page size, or decrease the key length.

30: The file specified is not a MicroKernel file.

Either the MicroKernel did not create the file, or a pre-v3.x MicroKernel created it. You can also receive this status code from earlier versions of Btrieve when you open a file created by a later version, if the file has a format incompatible with the earlier version. This status code can also indicate that the first page of the file is damaged. Use a backup copy of your data file. If you receive this status code and you suspect that the header page of the source file is damaged, recover the file as described in the *Pervasive.SQL User's Guide*.

31: The file is already extended.

This status code is obsolete in MicroKernel versions 6.0 and later.

The MicroKernel returns this status code if an application tries to extend a file that has already been extended; you can only extend a file once.

32: The file cannot be extended.

The MicroKernel returns this status code when it must create an extension file to accommodate a file which is growing larger than the operating system's file size limit, but the MicroKernel encounters an error from the operating system when it tries to create and open the new extension file. Possible causes for receiving this status code include the following: the directory is full, the disk is full, or the MicroKernel has not been granted sufficient rights by the operating system.

33: The MicroKernel cannot unload.

In the DOS environment, the MicroKernel returns this status code for one of the following two reasons:

- You attempted to unload the MicroKernel when you have loaded another terminate and stay resident (TSR) program after you loaded the MicroKernel. Unload the other TSR before unloading the MicroKernel.
- You attempted to unload the MicroKernel from a 32-bit application that uses the BSTUB interface with the DOS/4G extender.

34: The specified extension name is invalid.

This status code is obsolete in MicroKernel versions 6.0 and later.

The MicroKernel returns this status code if an application specified an invalid filename for the extended partition. Check the validity of the filename.

35: The application encountered a directory error.

Either a Get Directory operation specified a drive that does not exist, or a Set Directory operation specified an invalid pathname. Check the validity of both the drive and the pathname.

36: The application encountered a transaction error.

Version 7.0 and later MicroKernels return this status code if Scalable SQL defines too many transaction nesting levels. The v7.0 MicroKernel allows each transaction a total of 255 internal nesting levels. However, Scalable SQL v4.0 uses some of these levels internally to enforce atomicity on INSERT, UPDATE, and DELETE statements. Therefore, a session can effectively define no more than 253 savepoints to be active at one time. This limit may be further reduced by triggers that contain additional INSERT, UPDATE, and DELETE statements. If your operation reaches this limit, you must reduce the number of savepoints or the number of atomic statements contained within it.

Version 6.x MicroKernels return this status code if the MicroKernel tried to perform a Begin Transaction operation, but it was not configured to allow transactions. Use the Setup utility to specify a higher value for the Number of Transactions configuration setting. Next, stop and then restart the MicroKernel so that your changes take effect.

- On a machine that is running both a workstation MicroKernel and a Btrieve Requester accessing a client/server MicroKernel, ensure that both the workstation MicroKernel and the client/server MicroKernel are configured for transactions.
- In a client/server environment, all servers running MicroKernels to which the client machine is attached must be configured for transactions, even if the files involved in the transaction are only located on one of the servers.

37: Another transaction is active.

The application issued a Begin Transaction (19 or 1019) operation while another transaction was active by the same user or task; the active transaction can be nested or non-nested. This status code often indicates a problem in nested transactions within your application.

38: The MicroKernel encountered a transaction control file I/O error.

This status code is obsolete in MicroKernel versions 7.0 and later.

The MicroKernel returns this status code if an error occurs when the MicroKernel tries to write to the transaction control file. Possible causes for receiving this status code are that the disk is full, the disk is write protected, the transaction control file (BTRIEVE.TRN) that is created when you load the MicroKernel has been deleted, or the transaction control file is flagged read-only or is corrupt.

39: A Begin Transaction operation must precede an End/Abort Transaction operation.

The application issued an End Transaction (20), End Nested Transaction (2020), Abort Transaction (21), or Abort Nested Transaction (2021) operation without a corresponding Begin Transaction (19 or 1019) or Begin Nested Transaction (2019 or 3019) operation. Make sure that each End or Abort Transaction operation in your program is executed only after a successful Begin Transaction operation.

40: The file access request exceeds the maximum number of files allowed.

This status code is obsolete in MicroKernel versions 6.0 and later.

The application tried to access more than the maximum number of files allowed within a transaction. You set the maximum number of different files that you can access during a logical transaction when you configure the MicroKernel.

41: The MicroKernel does not allow the attempted operation.

The MicroKernel returns this status code for one of the following reasons:

- The application tried to perform an operation that is not allowed at this time. The MicroKernel does not allow some operations under certain operating conditions. For example, the MicroKernel returns this status code if the application attempts to perform a Step operation on a key-only file.
- If using a server engine, the key number parameter of a continuous operation MicroKernel call is not valid.

Also, the MicroKernel prohibits certain operations during transactions because they have too great an effect on the file or on performance. These operations include Set Owner, Clear Owner, Create Index, and Drop Index.

42: A file previously opened in Accelerated mode was not closed.

This status code is obsolete in MicroKernel versions 6.0 and later.

The MicroKernel returns this status code for one of the following reasons:

- Either the application tried to open a v5.x data file that was previously accessed in Accelerated mode by a v5.x MicroKernel and never successfully closed, or the application tried to open a file for which a v6.0 or later MicroKernel encountered an unrecoverable error during a Set or Clear Owner operation. The file's integrity cannot be ensured. Refer to the *Pervasive.SQL User's Guide* for information about recovering damaged files.
- Your application tried to open a file in MicroKernel v5.x format using a v5.x MicroKernel; however, that same file was previously accessed by a v6.0 or later MicroKernel, which failed to close the file successfully and left a pre-image file on the disk. Version 5.x MicroKernels cannot read pre-image files created in v6.0 or later format.

43: The specified record address is invalid.

The MicroKernel returns this status code for one of the following reasons:

- The record address specified for a Get Direct operation is invalid. Either the address is outside the file's boundaries, or it is not on a record boundary within or on a data page, or the record at the specified address has been deleted. For a Get Direct operation, specify the 4-byte address obtained by a Get Position operation.
- If the records' file is in v5.x format, this status code can indicate a file access conflict. For example, task 1 has a file locked in an exclusive transaction. Task 2 is reading records from the same file and tries to update a record that the transaction inserted. If task 2 reads the record and then task 1 aborts the transaction, task 2 receives this status code when issuing the Update operation.
- For a Find Percentage operation that is seeking a percentage based on a record's physical location within the file, the specified record address is invalid.
- The file may be corrupt, and you must recover it. Refer to the *Pervasive.SQL User's Guide* for information about recovering damaged files.

44: The specified key path is invalid.

The application tried to use the Get Direct/Record operation to establish positioning on an index path for a key whose

value is null in the corresponding record. The MicroKernel cannot establish positioning based on a null key value. Pre-v6.x MicroKernels return Status Code 82 in this situation; therefore, write your application to check for both of these status codes.

45: The specified key flags are invalid.

The key flags specification on a Create operation is inconsistent. If a key has multiple segments, the duplicate, modifiable, and null attributes should be the same for each segment in the key. Also, you cannot use the null or manual key attributes in a key-only file. The MicroKernel also returns this status code if the application attempted to specify a different ACS for two or more segments of a segmented key.

46: Access to the requested file is denied.

The MicroKernel returns this status code for one of the following reasons:

- The application opened a file in Read-only mode and tried to perform a Write operation on that file.
- The application attempted to perform a Write operation on a file that is flagged read-only by the operating system.
- When the application opened the file, it did not correctly specify the owner name required for updates.
- (Workstations Only) If a workstation engine user or task opens a file that client machine has opened using a server MicroKernel, the MicroKernel returns this status code if the workstation engine user attempts to write to the file.

47: The number of files opened exceeds the maximum allowed.

This status code is obsolete in MicroKernel versions 6.0 and later.

Pre-v6.0 workstation MicroKernels return this status code when the number of files opened in Accelerated mode exceeded the number of buffers available in the MicroKernel's cache. When a file is opened in Accelerated mode, the MicroKernel reserves one of its cache buffers for the file. It always reserves five empty buffers for index manipulation. Reconfigure Btrieve with both a smaller /P configuration option (to allocate more buffers) and a larger /M option (to increase the cache allocation).

48: The alternate collating sequence definition is invalid.

The MicroKernel returns this status code for one of the following reasons:

- The first byte of an ACS definition (the identification byte) does not contain the hexadecimal value *AC* (for user-defined ACSs), *AD* (for locale-specific ACSs), or *AE* (for international sorting rules support). Make sure that the first byte contains the appropriate value.
- You are running a pre-v6.0 MicroKernel, and you attempted to add an index with a locale-specific ACS to a file in v5.x format. Pre-v6.0 files do not support locale-specific ACSs.
- You are running the v6.0 or later MicroKernel with the Create File Version option set to v5.x, and you attempted to create a file that contains a key with a locale-specific ACS. Pre-v6.0 files do not support locale-specific ACSs.

49: The extended key type is invalid.

The MicroKernel returns this status code for one of the following reasons:

- You tried to create a file or an index with an invalid extended key type, or you are trying to assign an ACS to a BINARY key or key segment. You can assign an ACS only to a STRING, LSTRING, or ZSTRING key type.

- You defined an index requiring an ACS, but no ACS definition exists (either in the file or in the key definition passed in the data buffer).
- You attempted to create a key segment with both the Case Insensitivity and the Alternate Collating Sequence flags set, and the MicroKernel is configured to create files in v5.x format. This combination is invalid for v5.x files.
- You attempted to create a file that contains multiple ACSs, but you are running a pre-v6.1 MicroKernel. Only v6.1 and later MicroKernels support multiple ACSs.
- You are running the v6.x or later MicroKernel with the Create File Version option set to v5.x, and you attempted to create a file with a NUMERICSA or NUMERICSTS key. Pre-v6.0 files do not support these key types.
- You set the default file creation format to v6.x, but you are using one of the new Pervasive.SQL 7 data types, such as CURRENCY or TIMESTAMP, which require the v7.x file format.

50: The file owner is already set.

The application tried to perform a Set Owner operation on a file that already has an owner. Use the Clear Owner operation to remove the previous owner before specifying a new one.

51: The owner name is invalid.

The MicroKernel returns this status code for one of the following reasons:

- If the application received this status code from a Set Owner operation, the owner names specified in the key buffer and data buffer do not match.
- If this status code occurred during an Open operation or a DROP TABLE statement, the application attempted to open a file that has an owner name assigned to it. The application must specify the correct owner name in the data buffer. Ensure that the owner name is null-terminated in the data buffer and that the data buffer length is set long enough to include the owner name plus the null terminator.
- If an application received this status code when removing a file from Continuous Operation mode, then the client ID of the calling application differs from the client ID of the application that originally put the file into Continuous Operation mode. (Version 7.0 and later MicroKernels return Status Code 88 in this instance.)

52: An error occurred while writing to the cache.

This status code is obsolete in MicroKernel versions 6.0 and later.

While trying to make a cache buffer available, the MicroKernel attempted to write data to a disk from a file that was previously opened in Accelerated mode. The operating system returned an I/O error during the write. This generally indicates a hardware memory problem. Unload and reload Btrieve before you continue.

53: The language interface version is invalid.

An application tried to access a file containing variable-length records with a language interface from Btrieve v3.15 or earlier. To access files with variable-length records, you must use a v4.x or later interface. The MicroKernel can also return this status code when an application calls BTRVID using a pre-v6.0 MicroKernel with the DOS interface. Pre-v6.0 MicroKernels do not support BTRVID.

54: The variable-length portion of the record is corrupt.

During a Get or Step operation, the MicroKernel could not read all or part of the variable-length portion of a record. The MicroKernel returns as much data as possible to the application. This status code usually indicates that one or

more pages used to store variable-length records are corrupt. Check the data buffer length the MicroKernel returns to see how much of the record was returned. Recover the damaged file as described in the *Pervasive.SQL User's Guide*.

55: The application specified an invalid attribute for an AUTOINCREMENT key.

The application tried to specify either the segmented or duplicate attribute for an AUTOINCREMENT key type. An AUTOINCREMENT key can be part of a segmented key only if the key number of the AUTOINCREMENT key is less than the key number of the segmented key. An AUTOINCREMENT key cannot allow duplicates.

56: An index is incomplete.

An index can be damaged if a Create Index operation (31) or a Drop Index operation (32) is interrupted before it runs to completion. Perform a Drop Index operation to completely remove the damaged index from the file, then rebuild the index with the Create Index operation, if desired.

57: An expanded memory error occurred.

This status code is obsolete in MicroKernel versions 6.0 and later.

Retrieve for DOS returns this status code if it receives an error from the Expanded Memory Manager. This status code usually means that the MicroKernel was unable to save or restore the memory mapping register context, indicating an incompatibility with another application that uses expanded memory.

58: The compression buffer length is too short.

The application tried to read or write a record that is longer than the value specified for the size of the compression buffer. Reconfigure the MicroKernel using the Setup utility, specifying a higher value for the Largest Compressed Record Size option.

59: The specified file already exists.

This status code is returned for the Create operation if the application specified -1 in the key number parameter and the name of an existing file in the key buffer parameter. To overwrite the existing file, remove the -1 from the key number parameter. To preserve the existing file, alter the filename specified in the key buffer parameter.

60: The specified reject count has been reached.

The MicroKernel rejected the number of records specified by the reject count before a Get Next Extended, Get Previous Extended, Step Next Extended, or Step Previous Extended operation found the requested number of records that satisfy the filtering condition. Check the first two bytes returned in the data buffer for the number of records that were retrieved.

61: The work space is too small.

The Get Next Extended, Get Previous Extended, Step Next Extended, and Step Previous Extended operations use a buffer as work space. This status code indicates that the work space (set by default to 16 KB) is not large enough to hold the filtering data buffer structure and the largest record to be received. When using v6.15.2 and later MicroKernels, you receive Status Code 0 if the work space is large enough to hold the filter/extraction expression and enough of the record to include all of the fields to be filtered or extracted.

62: The descriptor is incorrect.

The descriptor (data buffer structure), which is passed for a Get Next Extended, Get Previous Extended, Step Next

Extended, or Step Previous Extended operation, is incorrect. The descriptor length (the first two bytes of the data buffer) on the extended operation call must be the exact length of the descriptor. This requirement does not apply to the data buffer length option, which can still be declared longer than necessary.

On a Stat Extended operation, the signature field in the data buffer is not set to 0x74537845, the subfunction field is not set to 0x00000001, or the namespace field is not set to 0x00000000.

On a Get Direct/Chunk or Update Chunk operation, the descriptor structure in the data buffer is incorrect, or it is inconsistent (either internally or with respect to the data buffer length).

63: The data buffer parameter specified on an Insert Extended operation is invalid.

An Insert Extended operation provided an invalid buffer. Either the buffer length is less than 5 bytes, or the number of records specified is 0. Correct the buffer length or the number of records.

64: The filter limit has been reached.

The MicroKernel returns this status code for any of the following reasons:

- During a Get Next Extended, Get Previous Extended, Step Next Extended, or Step Previous Extended operation, a rejected record was reached; no other record can satisfy the given filtering condition, going in the direction that the operation specified. This is applicable only if the first segment of the key that the key number specified is also used as the first term of the filtering field.
- You can receive this status code when the number of records to be retrieved, which is specified in the data buffer of the extended operations, is greater than the number of records present in the Btrieve file that satisfy the filter condition.

65: The field offset is incorrect.

The field offset in the extractor of a Get Next Extended, Get Previous Extended, Step Next Extended, or Step Previous Extended operation is invalid based on the length of the retrieved record. Make sure that the field offset is a valid value (from 0 through the record length minus 1).

66: The maximum number of open databases has been exceeded.

The MicroKernel tried to open files bound to too many Scalable SQL databases. Use the Setup utility to configure a higher value for the Maximum Databases option. Refer to the *Pervasive.SQL Programmer's Guide* for more information about bound files.

67: The MicroKernel cannot open the SQL data dictionary files.

The MicroKernel returns this status code for one of the following reasons:

- An application attempted to use a data file that is bound to a Scalable SQL database, but the MicroKernel could not open one of the Scalable SQL data dictionary files (FILE.DDF or, if the file has RI definitions, RELATE.DDF) or the configuration file (DBNAMES.CFG).
- You attempted to create a file with the Replace option, and a bound Scalable SQL data file with the same name and location already exists. However, the MicroKernel could not open the Scalable SQL data dictionary file FILE.DDF, or the configuration file (DBNAMES.CFG).

If the data file has RI definitions, the DBNAMES.CFG file must be in the location you specified using the Setup utility; you determine the location of DBNAMES.CFG by using the Database Names Directory option for the MicroKernel. Also, ensure that FILE.DDF and RELATE.DDF (if the file has RI definitions) are in the locations specified by the

database's configuration.

Refer to the *Pervasive.SQL Programmer's Guide* for more information about RI.

68: The MicroKernel cannot perform the RI Delete Cascade operation.

The MicroKernel cannot enforce the Delete Cascade rule on a file under RI control because the record that the application attempted to delete has more than 16 levels of descendants. Delete records from the lower levels, and then try again to delete the record that the application was attempting to delete initially. Refer to the *Pervasive.SQL Programmer's Guide* for more information about RI.

69: The Delete operation specified a file that is damaged.

The application encountered an error while the MicroKernel was attempting to enforce the Delete Cascade rule in response to a Delete operation. This status code indicates that the related file has been damaged and must be recreated. Refer to the *Pervasive.SQL User's Guide* for more information about recovering damaged files.

Refer to the *Pervasive.SQL Programmer's Guide* for more information about RI and the Delete Cascade rule.

71: There is a violation of the RI definitions.

If you have attempted an Insert operation on a file under RI control, you can receive this status code if a foreign key value in the record to be inserted does not have a corresponding primary key in the referenced file. If you are performing an Update operation, there are two possible causes for this status code:

- You attempted to change the value of a primary key.
- You attempted to change the value of a foreign key to a value that does not exist for the defined primary key.

If you attempted a Delete operation, the restrict rule is enforced, and a primary key value in the record you are trying to delete references a foreign key in the referenced file. Refer to the *Pervasive.SQL Programmer's Guide* for more information about RI.

72: The MicroKernel cannot open the RI referenced file.

The referenced file cannot be found at the location specified by FILE.DDF and DBNAMES.CFG. Be sure that the referenced file is in one of the data file locations that the DBNAMES.CFG file specifies for the named database.

- If the DBNAMES.CFG file is defined on a server, verify that the file location does not contain a drive letter.
- If the DBNAMES.CFG file is defined for a workstation engine, make sure that the drive letters are the same (and map to the same locations) as specified in DBNAMES.CFG.

Refer to the *Pervasive.SQL Programmer's Guide* for more information about RI.

73: The RI definition is out of sync.

The MicroKernel returns this status code for one of the following reasons:

- You tried to open a data file that is bound to a Scalable SQL database, and the database to which the file is bound was not found in the DBNAMES.CFG file.
- You tried to open a data file with RI definitions that is bound to a Scalable SQL database, and the table to which the file is bound was not found in the database's FILE.DDF file, or the table's location and filename does not match the file's location and filename as configured in the DBNAMES.CFG or FILE.DDF file.

- You attempted to modify a bound file, and the RI definition for that file disagrees with the definition in the RELATE.DDF file.
- You can also receive this status code on an Insert or Delete operation, or on an Update operation that would change a foreign key, if the file related to this file is out of sync (an attempt to open or modify the related file would have returned this same status code).
- You attempted to create a file with the Replace option, and a bound Scalable SQL data file with the same name and location already exists. However, the MicroKernel detected that the existing bound file was out of sync (that is, an attempt to open the existing file would have returned this same status code).

Check the RI constraints on your database. For information about how to do this, refer to the *Pervasive.SQL User's Guide*.

74: The MicroKernel aborted the transaction.

This status code is obsolete in MicroKernel versions 6.0 and later.

This is an informative status code. A NetWare-based MicroKernel replaced an End Transaction operation with an Abort Transaction operation after detecting an error for a Transaction Tracking System (TTS) file inside the transaction. The MicroKernel then executed the Abort Transaction operation.

75: The server routing list is too small.

If you use SET BRQPARMS=/S:1 with the Btrieve for OS/2 Requester, this status code occurs after you try to open a file or get the version number of Btrieve for NetWare. Set the /S value to 2 or higher. Do not set /S so that it defaults to 1.

This status code has been supplemented in Pervasive.SQL 7. For a list of the new status codes, see [Appendix A. "Changes Since Btrieve 6.15"](#).

76: There is a conflict on the referenced file.

An application attempted to perform an Update, Insert, or Delete operation on an RI-controlled file that references another file. The application cannot open the referenced file for RI checking because it is already open in Exclusive mode. Wait until the referenced file is closed or is opened in a mode other than Exclusive, and then retry the operation. Refer to the *Pervasive.SQL Programmer's Guide* for more information about RI.

77: The application encountered a wait error.

This status code is obsolete in MicroKernel versions 7.0 and later.

This is an informative status code. You must retry the operation yourself; the MicroKernel does not automatically retry the operation. A client/server MicroKernel returns this status code in one of the following situations:

- The application specified a wait lock bias for an operation, but another user has locked the requested resource.
- The application is currently processing a wait transaction and tried to access a file that another user has locked.

When you are using the Btrieve Requester to access the MicroKernel, the Requester waits and retries if a requested resource is locked. When a server-based application is accessing the MicroKernel and the requested resource is locked, a wait is also required. In this case, the MicroKernel is expected to perform the wait. Because this would occupy the MicroKernel and lock out other users who might be trying to release the requested resource, the MicroKernel does not perform the wait. Instead, it returns this status code, and the server-based application must retry later.

78: The MicroKernel detected a deadlock condition.

The application should clear all resources (for example, by aborting or ending the transaction or releasing all record locks) before proceeding. This breaks the deadlock, allowing other applications to access the resources for which they are waiting.

79: A programming error occurred.

This status code is obsolete in MicroKernel versions 7.0 and later.

Although very rare, it is possible to receive this status code when there is a malfunction that the MicroKernel cannot specifically detect or from which the MicroKernel cannot recover. Retry the operation. If the error persists, there may be system corruption; try to clear the system by rebooting, and then try the operation again.

80: The MicroKernel encountered a record-level conflict.

The MicroKernel did not perform the Update or Delete operation because of a record-level conflict. For example, station A reads a record, station B reads the same record and updates it, and then station A attempts to update the record. The application should reread the record prior to resending an Update or Delete operation. Alternatively, the application can employ record locks to avoid conflicts.

In key-only files, you receive this status code if the record is moved in the file's b-tree after being read and before being updated or deleted. A record can move as a result of other records being inserted, updated, or deleted.

In pre-v6.0 workstation MicroKernels, a page-level conflict also produces this status code. For example, station A reads a record, station B reads and updates a different record on the same data page, and then station A attempts to update the record it read. (You also receive this status code with v6.0 and later MicroKernels if the files are in v5.x format.)

In pre-v6.0 workstation MicroKernels, a conflict can also occur if station A reads a record, station B reads a different record and updates it, causing duplicate pointers to be updated on record A, and then station A tries to update the record.

81: The MicroKernel encountered a lock error.

The MicroKernel returns this status code in one of the following situations:

- The application tried to unlock a record that is locked with a multiple record lock, but the record position stored in the data buffer does not correspond to any record locked in the associated file.
- The application tried to unlock a single-record lock with a multiple-record lock or vice-versa.
- (Workstation Only) You tried to open a v6.0 or later file in Read-only mode on a NetWare drive on which you only have Read and Filescan rights. The MicroKernel cannot create the .LCK file. You must obtain Create and Write rights or redirect the .LCK files to a directory on which you have Create and Write rights. This directory must be shared by all users of a single datafile. For example, if two users are sharing a data file located at \\example\stats1:\test\file.btr and the lock files are redirected to C:\, then file sharing problems occur due to the users not sharing a single lock file directory.
- (Workstation Only) You tried to open a file in MEFS mode on a local drive, and the DOS SHARE.EXE program is not loaded. Either load SHARE.EXE or open the file in SEFS mode.
- (Workstation Only) You manipulated a file (you opened and then locked all or part of the file) in MEFS mode on a local drive, and the DOS SHARE.EXE program ran out of locks. Reload the SHARE program and specify a higher value for the /L and /F options.
- (Workstation Only) You redirected the lock file to a local drive with the /I option, and the DOS SHARE.EXE program is not loaded.

- (Workstation Only) You are running in MEFS mode, and you ran out of operating system file handles when the MicroKernel attempted to allocate a file handle to the .LCK file. Refer to your operating system documentation for instructions about configuring your operating system to allow more file handles.
- In pre-v7.x MicroKernels, the lock table is full. Decrease the number of locks that the application uses, or use the Setup utility to specify a higher value for the Number of Locks option.

82: The MicroKernel lost positioning.

When performing a Get Next or Get Previous operation on a key with duplicates, the application tried to retrieve a record that was deleted or whose key value was modified by another application. Use a Get Equal or a Get Direct/Record operation to re-establish positioning. (See Status Code ["44: The specified key path is invalid."](#) for a related positioning problem.)

83: The MicroKernel attempted to update or delete a record that was read outside the transaction.

This status code is obsolete in MicroKernel versions 7.0 and later.

The MicroKernel returns this status code if the application tried to update or delete a record within a transaction, but it did not read the record within the transaction. The application must read the record within the transaction before attempting to modify the data.

84: The record or page is locked.

The application tried to apply a no-wait lock on a record that is currently locked by another application, or the application tried to access a file in a no-wait transaction while another application holds one or more active record locks in that file. This status code can also occur if the application tried to update or delete a record locked by another application.

The application can use either of the following recovery methods:

- Retry the operation until it is successful. This can be the simplest and quickest solution for a network with light to moderate use.
- Use the wait option (+100/+300) instead of the no-wait option (in versions that support the wait option).

The MicroKernel can return this status code on an Insert, Update, or Delete operation when it attempts to lock an index page to insert or delete a key value. Have your application check for this status code and retry the operation if the status code is returned.

On Btrieve for NetWare, you can receive status 84 when running an application on a Win95 client if the NetWare operating system runs out of record locks. To solve this, increase the "maximum record locks per connection" and, if necessary, the "maximum record locks" (system wide limit) on the NetWare server.

85: The file is locked.

The MicroKernel returns this status code in one of the following situations:

- The workstation MicroKernel has a file open, and client machine that has the Requester loaded tries to open the same file via the server MicroKernel. The server MicroKernel cannot open the file because it cannot obtain exclusive access. The client machine that has the Requester loaded receives this status code.
- In a workstation engine environment, the MicroKernel can return this status code on an Open, Insert, Update, or Delete operation for a file under heavy usage by multiple users or tasks. The MicroKernel must momentarily have exclusive access to the file during these operations, and it retries the operation several times before returning this status code. In this case, the application can reissue the operation. In addition,

you can reconfigure the workstation MicroKernel with a lower Operation Bundle Limit and Initiation Time Limit to reduce the amount of time the MicroKernel keeps a lock on the file. Refer to the *Pervasive.SQL User's Guide* for more information on how to do this.

- In a workstation environment, a v6.15 or later MicroKernel has a pending modification (Insert, Update, or Delete) as an incomplete system transaction in a file that has been opened in MEFS mode. If multiple users or tasks attempt to access (Get/Step) or modify (Insert, Update, or Delete) the shared file, the MicroKernel returns this status code. An access operation can receive this status code only if the writing phase of the system transaction has started.

Reconfiguring the MicroKernel with a lower Operation Bundle Limit and Initiation Time Limit reduces the occurrences of file contention that produce this status code. Refer to the *Pervasive.SQL User's Guide* for more information about using the setup utility to reconfigure the MicroKernel.

If you are a developer and want more information about system transactions, refer to the *Pervasive.SQL Programmer's Guide*.

- While one user has a file locked in an exclusive transaction, another user attempts to lock all or part of that file.
- When opened by a server MicroKernel, a file is in transition into Continuous Operation mode. Retrying eventually works.
- When opened by a server MicroKernel, two data files have the same filename but different extensions (for example, INVOICE.HDR and INVOICE.DET). One file is open and in Continuous Operation mode, causing the MicroKernel to generate a delta file (for example, INVOICE.^^). The MicroKernel returns this status code when you attempt to open the second file.
- When opened by a Windows NT server MicroKernel using Microsoft File and Print Services for NetWare on behalf of a Win16 Windows client, the file was also opened simultaneously by a Win32 Windows NT or Windows 95 machine. This causes the server MicroKernel to open the same physical file using two different paths.

86: The file table is full.

Using the Setup utility, specify a higher value for the Open Files configuration option.

87: The handle table is full.

Either you have attempted to open more handles than the MicroKernel is configured to support, or the MicroKernel attempted to open more files than the operating system allows. To configure the MicroKernel to allow more handles, use the Setup utility and increase the value of the Handles option.

To configure your operating system to allow more handles, refer to your operating system documentation. It is helpful to know the following details regarding the MicroKernel's requirements for handles from the operating system. When the same file is opened multiple times, the MicroKernel uses only one operating system handle. However, if the file is in v6.x or later format and the file is shared via MEFS mode, the MicroKernel opens a second handle for the associated .LCK file.

If the file is in v5.x format, the MicroKernel might request a second handle, for the .PRE file. Also, if the file (in any format) is placed in Continuous Operation mode, the MicroKernel requests another handle for the delta file. If the file is extended, the MicroKernel requests an operating system handle for each of the extension files.

88: The application encountered an incompatible mode error.

The MicroKernel returns this status code in one of the following situations:

- If an application opens a file in Exclusive mode, all other applications receive this status code when they try

to open the same file in any mode.

- If an application opens a file in any mode other than Exclusive, all other applications receive this status code when they try to open the same file in Exclusive mode.
- (Workstation Only) If a user or task opens a file in MEFS mode, all other users or tasks using the same MicroKernel receive this status code when they try to open the same file in SEFS mode.
- (Workstation Only) Similarly, if a user or task opens a file in SEFS mode, all other users or tasks using the same MicroKernel receive this status code when they try to open the same file in MEFS mode.
- (Workstation Only) If a user or task opens a file in Accelerated mode and MEFS is enabled, all other remote users or tasks receive this status code when they try to open the same file in any mode. All other local users or tasks receive this status code when they try to open the same file in any mode other than Accelerated.
- If another client machine has the Requester loaded and has the file open, applications running v6.15 or later MicroKernels receive this status code when they try to open the same file in SEFS mode.

Using the MicroKernel's Continuous Operation mode, this status code can also indicate one of the following situations:

- You attempted to remove a file from continuous operation, but the file is not in Continuous Operation mode.
- You attempted to remove a file from continuous operation, but a different client placed the file into continuous operation.
- You attempted to include two files in continuous operation that have the same name but different extensions.
- You attempted to include a file in continuous operation, but the file is already in continuous operation.

89: A name error occurred.

This status code is obsolete in MicroKernel versions 5.0 and later.

BSERVER was loaded before you specified the short name to which the device was redirected. You must specify all short names that you want to share with the NET SHARE command before you start BSERVER.

90: The redirected device table is full.

This status code is obsolete in MicroKernel versions 6.0 and later.

The DOS Requester's redirection table or server routing table is full. This occurs if you attach to additional servers or map additional drives after loading the Requester. Reload the Requester, specifying a larger value for either the Number of File Servers (/S) option or the Number of Mapped Drives (/R) option. This status code also occurs if you detach from a server and attach to a different server. Once a client has attached to a server, the Requester does not remove its name from the server routing table.

91: The application encountered a server error.

The MicroKernel returns this status code in one of the following situations:

- The Requester cannot establish a session with the server. Either the client/server MicroKernel is not loaded or the server is not active.
- The SPX drivers are not installed or are outdated.
- The value for the Number of Sessions configuration option is too low. Use the Setup utility to specify a higher value for this option.

- An application specified a path for a file and did not include the volume name in the path.
- The Btrieve Message Router has not been loaded, and the following situation has occurred: an application that uses both the Btrieve Message Router and the MicroKernel to make remote calls (and which therefore includes the server and volume name when performing an Open operation) has attempted to open a remote file. Because the Btrieve Message Router does not interpret the server name, the MicroKernel attempts to do so but cannot.
- A communication or network addressing problem exists in your network environment, so the MicroKernel requests never reach their destination server address. Ensure that your client and server network components are up to date and certified for your network environment.
- For NetWare servers only:
 - the Maximum Packet Receive Size configuration option is inappropriate for your environment. For example, the setting should be 1500 for an Ethernet LAN or 4096 for a Token-Ring LAN.
 - The user count limit has been exceeded. Either close a session or upgrade your user count. For more information about purchasing and installing additive user counts, refer to the *Pervasive.SQL User's Guide*.
 - Ensure that the NDS network number is the same as the Internal Network Number viewed by BINDER.EXE output.
 - You ran BUTIL.NLM to roll forward a file using a log filename other than the default, and your BLOG.CFG file did not contain a correct entry such as "\dir\file.ext=vol:\dir\log.ext".

This status code has been supplemented in Pervasive.SQL 7. For a list of the new status codes, see ["Status Code 91 replacements"](#).

92: The transaction table is full.

This status code is obsolete in MicroKernel versions 7.0 and later.

The MicroKernel returns this status code if the application exceeded the maximum number of active transactions. Use the Setup utility to specify a higher value for the Number of Transactions configuration option.

93: The record lock types are incompatible.

The application tried to mix single-record locks (+100/+200) and multiple-record locks (+300/+400) in the same file at the same time. You must release all locks of one type before you can execute a lock of the other type.

94: The application encountered a permission error.

The MicroKernel returns this status code in one of the following situations:

- The application tried to open or create a file in a directory without the proper privileges. The MicroKernel does not override the network privileges assigned to users.
- The designated server is in the server routing table, but your particular client is not logged into that server.
- Both the server-based and workstation MicroKernels (pre-v6.15) are trying to access the same file at the same time.
- A NetWare application tried to access a file using NetWare Runtime support with the given username. Specifically, one of the following situations exists regarding the supplied username:
 - The user is not a valid user on the NetWare Runtime server.

- The user does not have the appropriate rights to access the file.
- The username is ADMIN or SUPERVISOR. For security reasons, the MicroKernel does not enable you to use ADMIN or SUPERVISOR as a username when enabling NetWare Runtime support.
- When using the Win32 Requester from a Windows NT or Windows 95 client machine to a NetWare server, you must use the same username for logging in to both the client machine and the NetWare server. You cannot be logged in to NetWare as SUPERVISOR or ADMINISTRATOR.
- When using the Win32 Requester from a Windows NT or Windows 95 client machine using NetWare emulation to a Windows NT server, the server cannot use Microsoft File and Print Services for NetWare. This causes the requester to attempt authentication as though the server were a NetWare Runtime server.
- It is recommended that you keep the default Requester setting of Yes on FPNW servers running Pervasive.SQL. You may receive a Status Code 94 if you change this setting to No when you are running the Btrieve Interface locally on the FPNW server and are using a local FPNW drive mapping or local FPNW UNC path.

95: The session is no longer valid.

The server MicroKernel returns this status code for one of the following reasons:

- The previously established session is no longer active due to an error at the client machine, at the server, or on the network. Verify that the client machine is still attached to the server, and then unload and reload the Btrieve Requester.
- The server MicroKernel has reached the maximum number of SPX sessions. Use the Monitor utility to check this statistic. Use the Setup utility to specify a higher value for the Number of Sessions configuration option.
- This status code could indicate a time delay problem if the client machine does not receive a response back from the server in an appropriate timeframe or after an appropriate number of retries. Refer to your network configuration documentation for information about increasing timeout and retry parameters. This is often necessary in a WAN environment or a LAN configuration with heavy network traffic.
- For NetWare servers:
 - Verify that the Maximum Packet Receive Size configuration option is appropriate for your environment. For example, the setting should be 1500 for an Ethernet LAN.
 - Ensure that the SPX timeout parameters are set as follows in both the client machine's NET.CFG file and the server's SPXCONFIG.NLM file:

```
SPX VERIFY TIMEOUT=54
```

```
SPX LISTEN TIMEOUT=108
```

```
SPX ABORT TIMEOUT=540
```

These three values must have a 1:2:10 ratio. You can increase these values to at most three times the default. If you continue to receive this status code after increasing these values, the problem is most likely not related to these settings.

- For Windows NT servers, verify that the Maximum Packet Size registry setting is 576 decimal or 240h. The path to the MaxPktSize registry setting is HKEY_LOCAL_MACHINE\System\currentControlSet\Services\NwlnkIPX\NetConfig\MaxPktSize.

If you continue to receive this status code after increasing the network timeout parameters, this status code usually indicates a problem with network communications. Verify that you have up to date network cards and

drivers; for example, incompatible LAN card drivers can also cause this status code to occur. Consult your LAN administrator for network communication troubleshooting.

This status code has been supplemented in Pervasive.SQL 7. For a list of the new status codes, see [“Status Code 95 replacements”](#).

96: A communications environment error occurred.

The MicroKernel returns this status code for one of the following reasons:

- You tried to attach to the MicroKernel on a server, but the SPX connection table or the MicroKernel's client table is full. Use the Setup utility to specify a higher value for the Number of Remote Sessions configuration option.
- An application that calls the MicroKernel can return this status code for the following reasons:
 - Not all of the clients have been properly reset.
 - You loaded the MicroKernel with a value that is too small for its Active Clients setting. Use the Setup utility to increase the Active Clients configuration option.
 - You are using a limited user count version of the MicroKernel, and it has reached the maximum number of users.
 - The DBNAMES.CFG file contains a named database definition specifying a data location on a different server.
- In the Windows environment, the v6.15 and later MicroKernel returns this status code if you try to run a Btrieve-based application, but the client table is full. Use the Setup utility to specify a higher value for the Active Clients configuration option. Refer to the *Pervasive.SQL User's Guide* for information on the Active Clients option.
- You attempted to connect to a Btrieve v6.10 server using a Win32 Btrieve v7.0 requester. This combination is not compatible and is therefore prevented by design.

97: The data buffer is too small.

The application either tried to read or write a record that is longer than the current allowed settings for the MicroKernel or the Btrieve Requester, as follows:

- For an Update, Insert, or Create operation, the application receives this status code if the data buffer length it specifies for the record exceeds the message buffer length.
- For a Get, Step, or Stat operation, the application receives this status code if the message buffer is shorter than the length of the data the MicroKernel would return, regardless of the data buffer length specified in the application.
- For a Get Chunk or Update Chunk operation, the total size of the retrieved or updated chunk exceeds the message buffer length.

For information about configuring the message buffer size, refer to the *Pervasive.SQL User's Guide*.

To avoid receiving this status code, perform one of the following steps:

- Increase the size of the message buffer by using the Setup utility to specify a higher value for the Communications Buffer Size option.
 - For local files, the setting is MicroKernel Database Engine/Memory Resources/Communications

Buffer Size.

- For remote files, the setting is Btrieve Communications Manager/Server Communications Configuration/Communications Buffer Size.
- *DOS Requesters only*: Reload the Btrieve Requester and specify a higher value for the message buffer size. This is done using the /D parameter which is documented in *Getting Started With Pervasive.SQL*.
- For Windows NT servers, verify that the Maximum Packet Size registry setting is 576 decimal or 240h. The path to the MaxPktSize registry setting is HKEY_LOCAL_MACHINE\System\currentControlSet\Services\NwlnkIPX\NetConfig\MaxPktSize.

98: The MicroKernel detected an internal transaction error.

This status code is obsolete in MicroKernel versions 6.0 and later.

A NetWare-based MicroKernel detected an error while executing the operation on a NetWare Transaction Tracking System (TTS) file. The application can perform only an Abort Transaction operation at this point.

99: The Btrieve Requester is unable to access the NetWare Runtime server.

The Btrieve Requester returns this status code when you enable NetWare Runtime server support and the Requester either detects no existing connection or cannot find a valid login username. SUPERVISOR and ADMIN are *not* valid usernames, even if supplied with the correct password. If the Requester cannot find a login username other than SUPERVISOR or ADMIN, there is no valid name to pass.

This error is also returned if you try to access a server with the NetWare Runtime server support disabled and you do not have an existing connection to that server.

100: No cache buffers are available.

This status code, returned by v6.0 and later MicroKernels, indicates that the MicroKernel has used all the cache buffers it allocated at load time. Use the Setup utility to increase the value for the Cache Allocation configuration option.

101: Insufficient operating system memory is available.

This status code, returned by v6.0 and later client/server MicroKernels, indicates that there is not enough operating system memory available to perform the requested operation. To fix this problem, perform one or more of the following:

- Use the Setup utility to decrease the value for the Cache Allocation configuration option.
- Add memory to the server.

For more information about the configuration options, see Status Code [“1002: A memory allocation error occurred.”](#), a similar status code that the Windows-based MicroKernel returns in the same situation.

102: Insufficient stack space is available.

This status code, returned as a run-time error by v6.0 and later MicroKernels, indicates that the MicroKernel has run out of stack space. To increase the amount of stack space available to your application, relink the application, setting the stack size to a higher value. The MicroKernel returns this status code only to Windows-based applications that call WBTRCALL.DLL, or applications that call the Btrieve API on the local server.

103: The chunk offset is too big.

The MicroKernel returns this status code in one of the following situations:

- A Get Direct/Chunk operation specified an offset beyond the end of the record, either explicitly or using the next-in-record bias to the subfunction value. Unless the MicroKernel returns this status code while processing the first chunk, the operation was partially successful. Check the data buffer length parameter immediately after the call to see how much data was retrieved (and therefore how many chunks).
- An Update Chunk operation specified an offset that is more than one byte beyond the end of the record. This status code indicates that the MicroKernel has made no changes to the record.
- An Update Chunk operation with an Append subfunction causes a record length to exceed its operating system file size limit. The MicroKernel has made no changes to the record.

104: The MicroKernel does not recognize the locale.

In v6.0 or later MicroKernels, the Create or Create Index operation returns this status code to indicate that the operating system was not able to return a collation table for the country ID and code page specified. Ensure that the application specified the locale's country ID and code page correctly and that the operating system is configured to support the country ID and code page.

105: The file cannot be created with Variable-tail Allocation Tables (VATs).

Version 6.0 and later MicroKernels return this status code when an application tries to create a file with Variable-tail Allocation Tables (VATs) but without variable-length records (a precondition for files to use VATs). This status code applies to key-only files as well as to regular data files.

106: The MicroKernel cannot perform a Get Next Chunk operation.

Version 6.0 and later MicroKernels return this status code when an application calls the Get Direct/Chunk operation to retrieve a chunk from a record and uses the next-in-record bias on the descriptor subfunction, but after the application establishes its positioning in the record (and prior to this call), the target record is deleted.

107: The application attempted to perform a chunk operation on a pre-v6.0 file.

Version 6.0 and later MicroKernels return this status code when an application tries to use either a Get Direct/Chunk operation or an Update Chunk operation on a file in pre-v6.0 format.

109: An unknown error was encountered either creating or accessing a semaphore.

OS/2 and Windows NT workstation MicroKernels return this status code when trying to perform an operation using incompatible versions of the DLLs. Shut down the MicroKernel and make sure that you are using the most recent versions of the DLLs.

This status code has been supplemented in Btrieve 7.0. For a list of the new status codes, see ["Status Code 109 replacements"](#).

110: The MicroKernel cannot access the archival logging configuration file.

The archival logging configuration file (BLOG.CFG) contains entries for the data files on the drive for which you want to perform archival logging. The MicroKernel returns this status code for the following reasons:

- The MicroKernel cannot find the BLOG.CFG file. Ensure that the file is in the \BLOG directory in a real root directory of the physical drive that contains data files you want to log. (That is, do not use a mapped root

directory.) If your files are on multiple volumes, you must create a \BLOG directory on each volume.

- The MicroKernel cannot open the BLOG.CFG file. Either the file is locked or it does not exist.
- The MicroKernel cannot read the BLOG.CFG file. Either the file does not use the correct format or it is corrupt. Refer to the *Pervasive.SQL User's Guide* for information about the format of the BLOG.CFG file.

111: The specified filename was not found in the archival logging configuration file.

The MicroKernel cannot find the specified file in the BLOG.CFG file. The file must be specified in the BLOG.CFG file on the same physical drive.

112: The specified file is in use by another client.

Before the MicroKernel can perform a roll forward, the file must be in the same state it was in when it was last backed up. If another client changes the file, you must restore the file again before rolling forward.

113: The MicroKernel cannot find the archival log for the specified file.

The MicroKernel cannot find the archival log file associated with the specified file. By default, the MicroKernel names the archival log file the same as the logged file, but with a .LOG extension. However, you can specify a different filename for the archival log file in the BLOG.CFG file. Ensure that the BLOG.CFG file indicates the correct filename for the archival log and ensure that the archival log file exists.

114: The archival log for the specified file is invalid.

The archival log associated with the specified file is not a valid archival log file. By default, the MicroKernel names the archival log file the same as the logged file, but with a .LOG extension. However, you can specify a different filename for the archival log file in the BLOG.CFG file. Ensure that the BLOG.CFG file indicates the correct filename for the archival log and ensure that the archival log file exists.

115: The MicroKernel cannot access the archival logging dump file.

The MicroKernel cannot access the archival logging dump file for one of the following reasons:

- The filename indicated for dumping entries in an archival log is not a valid filename. Be sure this filename does not contain a volume specification. The dump file is created on the same volume as the log file.
- The caller does not have access rights to the dump file.
- The MicroKernel cannot open the file because another user has opened the file using an exclusive operating system lock.

130: The MicroKernel ran out of system locks.

This status code is obsolete in MicroKernel versions 6.15 and later.

This status code can indicate a temporary condition in which no system locks are currently available. The following are example cases:

- A single client is performing a very large transaction, in which thousands of records are being modified.
- Many clients are performing large transactions concurrently.

A client can receive this status code whether or not it is in a transaction. In some cases, a client can simply retry the failed operation. If other clients have released system locks in the interim, the retried operation may succeed. If a client in a transaction receives this status code, end or abort the transaction. If the transaction is very large, consider breaking it into multiple, smaller transactions. You can also use the Setup utility to lower the number of system locks devoted to explicit locking. To do so, lower the values assigned to the Number of Locks and/or Number of Sessions configuration options.

132: The file has reached its size limit.

The MicroKernel returns this status code in one of the following situations:

- An operation attempted to allocate more than 16,777,216 pages to a data file.
- A data file has remained in continuous operation for a lengthy period of time, causing its delta file to exceed the operating system limit.
- A pre-v6.15 data file has reached the operating system file size limit.

133: More than 5 concurrent users attempted to access the same data file.

In the Pervasive.SQL Developer Kit for a workstation environment, you attempted to access a data file with more than five MicroKernels at the same time. The Pervasive.SQL Developer Kit for workstation environments limits the number of concurrent users of a file to five engines.

134: The MicroKernel cannot read the International Sorting Rule.

The MicroKernel returns this status code for one of the following reasons:

- The ISR is not found in the COLLATE.CFG file.
- The COLLATE.CFG file is missing or corrupt.
- The MicroKernel cannot read the ISR from the COLLATE.CFG file.

135: The specified ISR table is corrupt or otherwise invalid.

The MicroKernel returns this status code when it has found a readable COLLATE.CFG file, but the specific International Sorting Rule table is invalid.

136: The MicroKernel cannot find the specified Alternate Collating Sequence in the file.

The MicroKernel returns this status code in the following situations:

- You tried to create an index that uses an ACS, but the MicroKernel cannot locate an ACS with the specified name in the file.
- You called a Step Next Extended, Get Next Extended, Step Previous Extended, or Get Previous Extended operation and specified an ACS name, but the MicroKernel cannot locate an ACS with the specified name in the file.

139: The MicroKernel has detected an unacceptable value in the key number.

Certain Btrieve operations either use, or reserve the use of, the key number parameter as a subfunction number, rather than as a means to specify a file's index to be used with the operation (as is done, for example, in the Get

Equal operation). This status code is returned if an application does not specify a valid subfunction number (via the key number parameter) to one of these operations. For information about acceptable values for the operation's key number parameter, consult the section in the *Btrieve Programmer's Guide* that documents the operation that returned this status code.

140: The savepoint ID is invalid.

The application supplied an invalid savepoint ID in a nested transaction operation. Ensure that you specify valid savepoint IDs.

143: The MicroKernel cannot allow unauthorized access to files in a secure Scalable SQL database.

You attempted to open a data file bound to a Scalable SQL database that has security enabled. The MicroKernel does not allow access to such files, except through Scalable SQL.

The MicroKernel also returns this status code if you are not using Scalable SQL and all of the following are true:

- You attempt to create a file with the Replace option.
- A bound Scalable SQL data file with the same name and location already exists.
- The database to which the existing file is bound has security enabled.

147: The log segment is missing.

The MicroKernel cannot find a log segment that is necessary for rolling at least one file forward.

148: A roll forward error occurred.

The MicroKernel encountered an error while rolling a file forward. Depending on the operating system, the MicroKernel reports an error message as follows:

- The NetWare MicroKernel displays the message on the server's system console and writes the message to the Pervasive Event Log (PVSW.LOG), which resides in SYS:\SYSTEM.
- The Windows 95 and Windows NT workstation MicroKernel displays the message in the console message window and writes the message to the Pervasive Event Log (PVSW.LOG), which is located in the WINDOWS\SYSTEM or WINNT\SYSTEM32 directory.
- The Windows NT server MicroKernel does not display a message, but writes the message to the Pervasive Event Log (PVSW.LOG) in the WINNT\SYSTEM32 directory.

149: You must use Scalable SQL to make changes to this file.

You attempted to write to a data file for which triggers are defined, and you are not using Scalable SQL. You can only make changes to files that contain triggers if you are using Scalable SQL.

151: The nesting level is too deep.

You have exceeded the number of nesting levels that the MicroKernel allows. The MicroKernel's nesting level limit is 4,095. Specify fewer nesting levels. If you are using Scalable SQL and you specified fewer than 4,095 nesting levels, you may have specified more nesting levels than your system memory can hold.

161: The maximum number of user count licenses has been reached.

You attempted to open another session when you were at the limit of your user count license. Either close a session or upgrade your user count. For more information about purchasing and installing additive user counts, refer to *Getting Started With Pervasive.SQL*. You can use the Monitor utility to determine which users currently have connections to Scalable SQL. For information about the Monitor utility, refer to the *Pervasive.SQL User's Guide*.

162: The client table is full. Try increasing the number of Active Clients in the MicroKernel Database Engine System Resource setting.

You receive this status code when your setting for Active Clients in the Configuration utility is set too low, causing Pervasive.SQL to exhaust its client table space.



To increase your Active Clients setting and solve Status 162:

1. Click **Start**, point to **Programs**, then to **Pervasive SQL 7**.
2. Select **Configuration** (Setup) from the list of programs.
3. If you are running Configuration from a client, click **Connect** and type the name of the server on which Pervasive.SQL is installed.
4. Select **MicroKernel Database Engine** from the list of components.
5. Select **System Resources/Directories** from the list of categories for the MicroKernel Database Engine.
6. Select the **Active Clients** setting.
7. Increase the value to the number of clients that access the server engine at any one time.

200 to 999

Scalable SQL Engine Status Codes

This section describes status codes that the Scalable SQL engine returns.



Note: Many of the status codes in this section refer to APIs that pertain to Scalable SQL v3.x, and will soon become obsolete. Refer to the *What's New* for backward compatibility information.

200: The Scalable SQL local engine or Requester is not loaded.

You issued a Scalable SQL engine request before you loaded Scalable SQL. For information about loading the local engine or the Requester, refer to the *Pervasive.SQL User's Guide*.

201: The relational primitive function code is invalid.

Check the function code you are using. Valid function codes for the relational primitives are 0 to 41, inclusive; valid function codes for the SQL-level functions are 0 to 16, inclusive.

If you have written your own interface to the relational primitives or Scalable SQL APIs, be sure to use the function number defined in the *Scalable SQL Programmer's Guide*. Reload the application, and try the operation again.

202: The cursor ID parameter is invalid.

The iCursorID parameter does not contain a valid cursor value. Use either xNew or xRecall to allocate a cursor.

This error can also indicate an attempt to perform an operation on the incorrect type of cursor ID. The following operations are invalid on a cursor ID you obtained from xNew or xRecall:

XQLCompile	XQLDescribe	XQLExec
XQLFree	XQLSPUtility	XQLStatus
XQLSubst		

Similarly, you cannot call xInsert, xUpdate, or xRemove unless you have compiled a valid SELECT statement using the same cursor ID.

203: The master password is invalid.

The password you passed in the call is not the correct master password.

204: The table is not defined in the dictionary.

A table name in your SQL statement does not refer to a table defined in the data dictionary. Ensure that you spelled the table name correctly. You can also receive this status code when you invoke a stored procedure that contains a reference to a table that is not defined or has been dropped since the procedure was created.

For pre-v4.0 Scalable SQL engines, or when using v3.01 Compatibility mode with the Scalable SQL v4.0 engine, ensure that you have specified the name correctly and that you use the appropriate blank replacement character if the name contains embedded blanks.

205: The password is invalid.

The specified password is not valid for the current user.

206: The maximum number of cursors has been exceeded.

When you load the Scalable SQL Requester, you can specify the maximum number of cursor IDs that can be active at one time. If you are using an application and receive this status code, reconfigure the Scalable SQL Requester at your client machine with a larger number of cursors using the */V* option.

If you are developing an application and want to allocate another cursor ID, first use *XQLFree* (or *xFree*) to release one of the other cursor IDs. If you need all the active cursor IDs for the application, reconfigure the Scalable SQL Requester at your client machine with a larger number for the */V* option.

For information about loading the Scalable SQL Requester, refer to the *Pervasive.SQL User's Guide*.

207: The column does not exist in the dictionary.

Scalable SQL can return this status code for one of the following reasons:

- One of the columns in the list is not defined in the currently active dictionary. Ensure that you specified the names correctly. If you are using v3.01 Compatibility mode, ensure that you use the appropriate blank replacement character for any blanks in the name.
- The specified column does not belong to any of the tables in the view.
- You used a named index in place of a column name. Scalable SQL v4.0 or later expects a column name. Scalable SQL interprets this index name as a column name and can find no such column.
- If you can access the column by using the SQL statement `SELECT * FROM table` but cannot access the column by name, the `FIELD.DDF` dictionary file could be corrupt. Restore a backup copy of your data dictionary files.

208: The position parameter value is out of range for this function.

The position you specified for a column in a column list is out of range. Valid positions are between 1 and the number of columns in the view, inclusive. Scalable SQL can also return this status code if the value for the position in `XQLDescribe` or `XQLFormat` is out of range. Valid column list positions are between 1 and the number of columns in the list, inclusive.

209: The specified username or group name is invalid.

The username or group name does not exist in the currently active dictionary. Either the name was entered incorrectly, or the user or group does not have access to this dictionary.

210: The buffer is not large enough to hold all the data.

Scalable SQL cannot copy all the requested data into the data buffer. For most APIs, you can either work with smaller amounts of data or increase the buffer size. If you are using the `SQLGetRemoteDatabaseNames` or `SQLGetDatabaseNames` functions, check `iBufLen` to determine the number of names returned. To retrieve more names, set the first 20 characters of `bDataBuf` to the last name returned on the previous call and call the function again. Make sure the buffer is large enough to contain 20 bytes for each database name.

211: The maximum record size has been exceeded.

Either the largest record size or the total length of all the columns in the view exceeds the Communications Buffer Size specified when the MicroKernel was configured. Either remove some columns from the view definition or reconfigure the MicroKernel. For information about the MicroKernel and the Setup utility, refer to the *Pervasive.SQL User's Guide*.

212: The maximum index size has been exceeded.

The MicroKernel restricts the total length of all the segments of an index to 255 bytes. Reduce the number of columns included in the index.

213: You do not have Select (Read) rights for the active dictionary.

Security has been installed for the currently active dictionary, and the username that is currently logged in does not have access rights to the table specified in the current operation. Either log in as a user who has rights to the table, or contact your database administrator about obtaining Select (Read) rights.

214: You must revoke Alter rights before attempting to change column-level privileges.

A user who has Alter rights to a table has rights to every column in the table. You cannot restrict the user's access to a specific column in the table until you revoke Alter rights to the table for that user.

215: The lock flag is invalid.

You passed an invalid lock flag to Scalable SQL on an xFetch or XQLFetch call. For further discussion on xFetch or XQLFetch calls, refer to the *Scalable SQL Programmer's Guide*.

218: You have exceeded the maximum number of tables allowed in a join.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The maximum number of tables allowed in a join is determined by the Maximum Tables in a Join configuration option that you set using the Setup utility (as described in the *Pervasive.SQL User's Guide*). You can specify up to 20 tables for this option; the default is 8.

If you are joining a previously defined view to other tables, the number of tables in the defined view is added to the maximum. For example, if Maximum Tables in a Join is set to 8 and your join list includes a view that contains 5 tables, you can list only three other tables in the join list.

219: You can specify a maximum of 24 index segments for a join operation.

This status code is obsolete in Scalable SQL versions 4.0 and later.

You can use up to 24 columns in a join condition.

220: The secondary columns you defined have no matching indexes.

The columns you specified as secondary columns do not match any of the index paths defined for the secondary table. Change the column names for the xJoin call, or use xDDCreate to create a named index using the columns in that table.

221: The syntax for an expression or restriction clause is invalid.

The expression or restriction clause you specified contains invalid operators or the statement is incomplete. To determine the approximate byte offset of the syntax error, you can use the value that Scalable SQL returns in `iExpLen` for the `xCompute` or `xRestrict` primitive or `iStatementLen` for the `XQLCompile` function. You can also use Option 1 of `XQLStatus` to obtain information about the error. For more information about the WHERE clause, refer to the *SQL Language Reference*.

222: The memory allocation is insufficient for Scalable SQL to process the current request.

Scalable SQL is unable to allocate sufficient buffer space to process the current request. If other views or cursors are active, release some of them to free memory for the current request. If you are running a server-based engine, you may want to consider adding more memory to your system.

223: The data type is invalid for the expression.

The data types for the operands in the expression are incompatible. For example, you may have attempted to mix string and numeric operands (-6.02E23) in an operation that accepts only string ("XYZ") or numeric operands. The same error occurs for similar cases involving other incompatible types. Refer to the *SQL Language Reference* for information about type compatibility.

224: A character in the numeric data is invalid.

A numeric constant contains nonnumeric data. Verify that the data you passed in the expression or restriction clause contains valid data for the particular data type according to its mask. Numeric data can consist only of the ASCII characters 0 to 9, a period (.), plus sign (+), and minus sign (-). If the data type is MONEY, the dollar sign (\$) is also allowed.

225: Scalable SQL allows a maximum of 119 columns in an ORDER BY or WITH INDEX list.

You can specify up to 119 columns in the column list for an index when you use `xDDCreate`, `xDDModify`, the `CREATE INDEX` statement, or the `WITH INDEX` clause of the `CREATE TABLE` statement. Also, you can specify up to 119 columns in the order list of an `ORDER BY` clause or the `xOrder` primitive.

226: The record count is invalid.

When updating or removing records using `xUpdate` or `xRemove`, you cannot specify more records than Scalable SQL returned on the last `xFetch` call. Also, `xInsert`, `xFetch`, `xRemove`, and `xUpdate` can return this status code if you set `IRecordCount` to 0.

227: The expressions must be column names or constants.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The expressions in a computed column or restriction clause must be either valid columns that are defined in the dictionary or constant values. Scalable SQL can also return this status code when you specify a restriction clause that contains a computed column and you qualify the computed column name with its table name. If the SQL statement involves more than one table or view, substitute the expression for the computed column name. For example, refer to the computed column as `Price * 1.08`, not `TotalPrice`.

228: The data length in the buffer is invalid.

The length of the word preceding the data in the buffer is larger than the view width.

229: The column already exists in this table definition.

Each of the columns within a table definition must have a unique name. The column was not added; specify a new column name.

230: The restriction clause is incomplete.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The restriction clause is missing a condition following the && or || connector.

231: No columns are defined in the view.

You must include at least one column in the view definition.

232: The column must be in the current view.

The columns you reference must be included in the current view definition.

234: Scalable SQL cannot create the external index file.

Scalable SQL cannot create a temporary file to build the external index because one of the following conditions occurred:

- You specified an invalid path for the External Sort File Directory option in the Setup utility.
- The MicroKernel has reached the maximum number of open files or handles that it allows. Increase the values for the Open Files and Handles configuration options using the Setup utility, reinitialize the MicroKernel and the Scalable SQL engine, and retry the operation. For more information about how to do this, refer to the *Pervasive.SQL User's Guide*.
- In pre-v4.0 Scalable SQL, you specified a SQL statement with an ORDER BY or GROUP BY clause that contains columns totalling more than 255 bytes in length.

235: The option value is invalid for the Fetch, Update, or Delete operation.

Verify that the option value you selected for the relational primitive call is valid.

After you perform an operation that changes the query optimization, you must perform a Fetch First or Fetch Last operation to re-establish positioning. The following operations change the query optimization:

- An xRestrict call that modifies a restriction clause
- An XQLSubst call that substitutes new values into a restriction clause

For example, Scalable SQL returns this status code if you issue an xRestrict call followed by a Fetch Next or Fetch Previous operation.

236: You cannot update a view that is ordered on an external index.

This status code is obsolete in Scalable SQL versions 4.0 and later.

You cannot perform an Update or Delete operation if the view is ordered on an external index. Change the sort order of the view to a permanent index before attempting to update the view.

237: When you define a named index, you must specify the index descriptor.

When you define a named index in the data dictionary, specify the index descriptor block in the format that is described in the documentation for the xDDCreate or xDDModify primitives.

238: The named index does not exist.

The name you specified for a named index does not exist in the data dictionary or is not a valid name for a named index.

239: The table in the view definition is not in the active dictionary.

The table in the stored view definition does not exist in the dictionary. You can no longer recall this view. Use the DROP VIEW statement to remove this view.

240: The column in the view definition is not in the active dictionary.

A column referenced in the stored view definition does not exist in the dictionary. You can no longer recall this view. Use the DROP VIEW statement to remove this view.

241: The application cannot explicitly update a system table.

The system tables are automatically updated when you execute one of the data definition statements or data definition primitives.

242: You do not have Write access to this table.

Security has been installed for the dictionary, and you do not have Write access for the table.

243: You do not have Write access to all the columns you are attempting to update.

You do not have Write access to all the columns that are being updated.

244: You do not have Write access to all columns in this table.

You do not have Write access to all the columns in the table, so you cannot insert or delete records.

245: The application cannot delete a system table.

The application cannot delete a system table using xDDDrop or a DROP TABLE statement. To remove a system table, you must delete the entire database by dropping the dictionary.

246: You do not have rights to specify or alter table definitions.

You do not have access rights for an operation that creates, alters, or drops a table definition.

247: You do not have Read rights to all the columns you specified.

You do not have at least Read access to one of the columns in the list.

248: The specified size is invalid for the defined data type.

Scalable SQL returns this status code for one of the following reasons:

- The size you specified for the data type is not a valid size. For information about data type sizes, refer to the *SQL Language Reference*.
- For pre-v4.0 engines, the size you specified for the computed column or dictionary column is not valid for the defined data type.

249: The data types of the columns are incompatible.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The data types of the replacement columns must be the same as the data types of the columns you are updating.

250: Scalable SQL cannot recall the view.

The view definition does not match the current dictionary definition. Ensure that the VIEW.DDF dictionary file exists in the same directory as the other dictionary files. If the VIEW.DDF file is accessible, the dictionary definition may have been modified since the view was stored. In this case, you must store the view again.

251: Scalable SQL cannot read the X\$View system table.

The MicroKernel returned a file error while reading the VIEW.DDF file. Retry the operation. If you continue to receive this status code, the VIEW.DDF dictionary file may be corrupt. You can try the following remedies:

- Use the file recovery procedure described in the *Pervasive.SQL User's Guide*.
- Restore the file from a backup.
- Rebuild the VIEW.DDF file.

252: VIEW.DDF must be open before you can perform this operation.

Scalable SQL returns this status code if the X\$View definition exists in the data dictionary, but the VIEW.DDF file could not be opened. Ensure that the VIEW.DDF file exists in the same directory as the other dictionary files.

253: The view definition was not stored with the correct version of Scalable SQL.

Scalable SQL returns this status code for one of the following reasons:

- You may be trying to recall a v3.01 view with Scalable SQL v4.0. Convert the view to v4.0 using the View Conversion utility. The View Conversion utility is described in the *Pervasive.SQL User's Guide*.
- You stored a view with one version of Scalable SQL and then tried to recall the view with an earlier version. Make sure you are not using an earlier version than the one you used to store the view.

254: Scalable SQL cannot create the X\$View system table.

An error occurred while Scalable SQL was trying to create the VIEW.DDF file. Make sure you have operating system rights to create files in the directory in which the data dictionary resides.

255: The column is not defined as an index.

The column you specified using an xDDIndex operation is not defined as an index.

256: You must specify at least one table for an Insert, Update, or Delete operation.

Make sure you specify at least one table name for an Insert, Update, or Delete operation.

257: The specified table is already defined.

The name you specified for a table already exists in the dictionary. Select another name, or remove the current table definition first.

258: Scalable SQL cannot update the table definition in the dictionary.

A MicroKernel error occurred when an existing table definition was being updated. This usually indicates that an I/O error has occurred. Try the operation again.

259: Scalable SQL cannot insert the new table definition in the dictionary.

A MicroKernel error occurred when a new table definition was inserted into the system tables. This usually indicates that an I/O error has occurred. Try the operation again.

260: Scalable SQL cannot update the X\$Rights system table.

A MicroKernel error occurred when the X\$Rights system table was being updated. This usually indicates that an I/O error has occurred. Try the operation again.

261: A variable-length column must be defined last in a view or table definition.

When you define a NOTE or LVAR column in a view or table definition, that column must be the last column in the definition. Scalable SQL also returns this status code if you try to add a variable-length column to a table definition that already has a NOTE or LVAR column.

262: You cannot define a variable-length column as an index.

You cannot use a NOTE or LVAR column as an index in a table.

263: You cannot sort a view by a variable-length column.

This status code is obsolete in Scalable SQL versions 4.0 and later.

You cannot specify a NOTE or LVAR column as the sort column for a view.

264: You cannot create, modify, or remove the Master username.

You cannot use xUser to create, modify, or remove the master user rights. The master user always has access to the entire dictionary.

265: The session identifier is invalid.

Scalable SQL can return this status code for one of the following reasons:

- You attempted to call a Scalable SQL API function before calling xLogin/XQLLogin.

- You attempted to call a Scalable SQL API function after your session was asynchronously reset at the server.
- You attempted to call a Scalable SQL API function that requires a valid sessionID before calling xLogin or XQLLogin.
- You attempted to call the xPutSessionID function using an invalid session identifier. Before calling xPutSessionID, you must successfully call XQLLogin (or xLogin) first in order to get a valid session ID.
- You attempted to execute SQL statements with the background mode feature while connected to a Pervasive.SQL workstation engine. SQLScope does not support this mode when you are connected to a Pervasive.SQL workstation engine.

266: The maximum number of user count licenses has been reached.

You cannot establish another login session because you have already reached the maximum number of login licenses that the user count limit of your version of Pervasive.SQL allows. You can use the Monitor utility to determine which users currently have connections to Scalable SQL. For information about purchasing and installing additive user counts, refer to *Getting Started With Pervasive.SQL*.

269: Scalable SQL cannot create the X\$User system table.

Scalable SQL cannot define the X\$User system table for the data dictionary. Query the system tables to determine if the X\$User table already exists, and make sure you have not used X\$User or any of its column names (Xu\$ID, Xu\$Name, Xu\$Password, or Xu\$Flags) as data element names elsewhere in the database. If these names have been used elsewhere (for table or index names, for example), Scalable SQL cannot use them to create the X\$User table.

Scalable SQL also returns this status code if more than one user attempts to install security on the database at the same time.

270: Scalable SQL cannot open the X\$User system table.

A MicroKernel error occurred when an attempt was made to open the X\$User system table. Retry the operation. If you continue to receive this status code, do either of the following:

- Use the Setup utility to increase the values for the MicroKernel Open Files and Handles options. Reinitialize the MicroKernel and the Scalable SQL engine, and retry the operation. For more information about how to do this, refer to the *Pervasive.SQL User's Guide*.
- Restore all system tables from a backup.

271: Scalable SQL cannot create the X\$Rights system table.

Scalable SQL is not able to define the X\$Rights system table for the data dictionary. Query the system tables to find out if the X\$Rights table already exists, and make sure you have not used X\$Rights or any of its column names (Xr\$User, Xr\$Table, Xr\$Column, or Xr\$Rights) as data element names elsewhere in the database. If these names have been used elsewhere (for table or index names, for example), Scalable SQL cannot use them to create the X\$Rights table.

Scalable SQL also returns this status code if more than one user attempts to install security on the database at the same time.

272: Scalable SQL cannot open the X\$Rights system table.

A MicroKernel error occurred when an attempt was made to open the X\$Rights system table. Retry the operation. If you continue to receive this status code, do one of the following:

- Use the Setup utility to increase the values for the MicroKernel Open Files and Handles options. Reinitialize the MicroKernel and the Scalable SQL engine, and retry the operation. For more information about how to do this, refer to the *Pervasive.SQL User's Guide*.
- Restore all system tables from a backup.

273: Scalable SQL cannot read or update the X\$User system table.

A MicroKernel error occurred when an attempt was made to access the X\$User system table. Retry the operation. If you continue to receive this status code, do one of the following:

- Reinitialize the MicroKernel and the Scalable SQL engine and retry the operation. For more information about how to do this, refer to the *Pervasive.SQL User's Guide*.
- Restore all system tables from a backup.

274: Scalable SQL cannot assign the owner name to the data file.

A MicroKernel error occurred when the owner name was being assigned during a CREATE TABLE, xDDCreate, or xDDModify operation. Ensure that you specified a valid owner name.

275: The user or group already exists in the dictionary.

The user or group you are defining with xUser or CREATE GROUP or GRANT LOGIN statements has already been defined in the data dictionary. Choose another name or remove the current user or group if it is no longer needed.

276: Only the master user can perform this operation.

Only the master user is allowed to disable security on an xSecurity call.

277: The access rights code is invalid.

Valid values for the access rights code are 0, 1, and 2.

278: The option is invalid.

The value for the option on the last API call is not valid. The function does not support the specified value for the sub-function. You must specify a valid value or a valid combination of values for the sub-function code. Refer to the *Scalable SQL Programmer's Guide* for valid option values.

279: The mask is invalid for the data type specified.

The mask you specified does not contain valid mask characters for the column's data type. Check the mask to ensure that the characters in it are valid. For more information about edit mask characters, refer to the *Pervasive.SQL Programmer's Guide*. Scalable SQL also returns this status code if you pass in a mask length that is longer than the maximum mask length (30 bytes).

280: The USING clause contains an invalid pathname.

Check the pathname for valid characters. Verify that the directory path exists and is valid. If you are logged into a bound database or a named database, you must use either a simple filename or a relative path. If you specify a relative path, Scalable SQL interprets it relative to the first data file path associated with the database name.

281: The attribute is not defined in the current dictionary.

The attribute you requested with an xDDAttr Fetch operation or SET MASK statement is not defined in the current data dictionary.

282: The data you specified includes invalid characters.

The data for a column in the data buffer contains characters that are not in the valid set of characters, as specified in the X\$Attrib system table. Make sure the data validation rules for the columns allow the values you specify for the columns. Contact your database administrator to receive a list of valid values.

283: A column value is not within the range of valid values.

A column value in the data buffer is not in the range specified for that column in the X\$Attrib system table. Make sure the data validation rules for the columns allow the values you specify for the columns. Contact your database administrator to receive a list of valid values.

284: A column value is not in the list of valid values.

A column value in the data buffer is not in the value list specified for that column in the X\$Attrib system table. Make sure the data validation rules for the columns allow the values you specify for the columns. Contact your database administrator to receive a list of valid values.

285: Scalable SQL cannot open the system tables.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The data dictionary in the specified directory cannot be opened. Be sure you specified the directory path correctly, the directory contains the correct system tables, and you have operating system rights to read and write files in the directory.

286: You cannot turn security on when it is already on (or off when it is already off).

You cannot enable security when it is already enabled, nor can you disable it when it is already disabled.

287: No data dictionary is currently active.

This status code is obsolete in Scalable SQL versions 4.0 and later.

You attempted to call a primitive before logging in with xLogin or XQLLogin, or before setting a valid session identifier with xPutSessionID. If you have successfully logged in or set a valid session identifier, a Scalable SQL internal error occurred.

288: The data dictionary is in use.

You cannot remove the data dictionary because a user is still logged in. Before you can remove a data dictionary, all activity for that dictionary must be completed.

289: The specified dictionary name is invalid.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The data dictionary name you specified is not in the database configuration file. Ensure that the name is valid, or add the name to the configuration file using the Setup utility.

290: The specified data dictionary already exists.

The data dictionary you specified already exists, and you did not specify the Replace option in the xDD call or use the WITH REPLACE clause in the CREATE DICTIONARY statement. If you want to replace the existing dictionary, issue the xDD call with the Replace option or use the WITH REPLACE clause in the CREATE DICTIONARY statement.

291: The specified directory does not contain a data dictionary.

The directory specified for a dictionary to be dropped does not contain a valid Scalable SQL data dictionary. Ensure that you specified the path correctly.

292: The numeric value for the month is invalid.

The numeric value for the month must be between 1 and 12, inclusive.

293: The numeric value for the day of the month is invalid.

The numeric value for the day must be between 1 and the number of days in the month.

294: The specified numeric time value is not within the range of valid values.

If the time mask does not contain the *ap* (for am/pm) characters, the numeric value for hours must be between 0 and 23, inclusive. If the time mask does contain the *ap* characters, the numeric value for hours must be between 1 and 12, inclusive. The numeric value for minutes and seconds must be between 0 and 59, inclusive. The numeric value for hundredths of a second must be between 0 and 99, inclusive.

295: The specified open mode value is invalid.

You issued an invalid open mode in a SET OPENMODE statement. Valid open modes for data files are Accelerated, Exclusive, Normal, Read-only, and Verify. For pre-v3.1 Scalable SQL, ensure that you specify a valid open mode. The valid values for the open mode are 0, -1, -2, -3, and -4. If you are using an older version of Scalable SQL, refer to the discussion of xNew in the *What's New*.

296: Scalable SQL cannot open the data file to retrieve the file statistics.

Scalable SQL is not able to open the data file to retrieve the file statistics. One of the following may be the case:

- The file is locked.
- The file does not exist.
- You do not have the necessary permissions to access the file. For example, an owner name could be set on the file.

Check the file's path. If the path is correct and you continue to receive this status code, then another user owns the file, has the file locked in a transaction, or has the file in Exclusive mode.

297: You cannot delete the column from the view.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

You cannot delete a column from a view if it is referenced in a restriction clause or order list.

298: A string constant contains mismatched quotes.

This status code is obsolete in Scalable SQL versions 4.0 and later.

A string constant does not contain beginning and ending single quote characters. Scalable SQL also returns this status code when you use AND or OR to connect an EXISTS clause with one or more expressions in a WHERE clause. To avoid receiving the status code, place parentheses around the EXISTS clause.

299: The join index for tableName:n.columnName is invalid.

The n must be a number between 1 and the number of occurrences of the same table in the view.

300: A conditional computed column must contain a :: operator.

This status code is obsolete in Scalable SQL versions 4.0 and later.

A conditional computed column must contain a :: operator to indicate the second part of the condition.

301: No restriction clause has been defined for the view.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

A restriction clause has not been defined for this view, so you cannot add another condition. Do not begin the restriction clause with the && or || connector.

302: The specified column is not a computed column.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

You requested a computed column description, but the column whose position you specified is a dictionary column, not a computed column. Specify the position of a computed column instead.

303: Cannot update current sort index if more than one file is in the view.

This status code is obsolete in Scalable SQL versions 3.0 and later.

You cannot update the current sort index in a joined view because Scalable SQL uses that index to read by when updating the file.

304: You cannot change another user's password.

Only the master user is allowed to change the password for another user.

305: The specified view is not in the current dictionary.

The current dictionary does not contain a view definition with the name you specified. Ensure that the VIEW.DDF dictionary file exists in the same directory as the other dictionary files.

306: Scalable SQL encountered an error while attempting to store the view definition.

An operating system or MicroKernel I/O error occurred when Scalable SQL attempted to store the view definition in the data dictionary. Scalable SQL cannot store the definition. Retry the operation. If you continue to receive this status code, the VIEW.DDF dictionary file may be corrupt. Try the following remedies:

- Recover the file as described in the *Pervasive.SQL User's Guide*.

- Restore the file from a backup.
- Rebuild the VIEW.DDF file.

307: The computed column you specified for the join does not exist in the view.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

Add a computed column to the view before you specify it in a join condition.

308: Scalable SQL cannot read the master password from the system tables.

Only DOS engines return this status code. Ensure that you have the proper operating system rights to access the dictionary files and try the operation again.

309: The valid owner access codes are 0, 1, 2, and 3.

You specified an invalid owner access code in the xDDCreate or xDDModify operation or in a CREATE TABLE statement. Valid owner access codes are between 0 and 3, inclusive. For more information about the CREATE TABLE statement, refer to the *SQL Language Reference*. If you are using an older version of Scalable SQL, refer to the discussion of the xDDCreate or xDDModify primitive in the *What's New*.

310: You do not have Read access to all the specified tables.

You do not have at least Read access to one of the tables in the list. Either remove that table from the list, obtain access rights to the table, or log in as a user with access to the table.

311: You have not activated transaction processing.

You attempted an operation that requires the session to be in a transaction. You must be in a transaction in order to use SAVEPOINT, ROLLBACK, or COMMIT statements.

312: Scalable SQL aborted the transaction.

You issued one of the following calls before completing a transaction: xLogout, XQLLogout, xReset, xStop, or XQLStop. Scalable SQL has rolled back the operations that were performed within the transaction.

313: You have already started transaction processing.

If a transaction is still active, you cannot issue a START TRANSACTION statement or call one of these xTrans options: 0 (Start Transaction), 3 (Start Transaction Using Cursor Stability Isolation Level), or 4 (Start Transaction Using Exclusive Isolation Level). If you want to nest transactions, use the SAVEPOINT statement.

314: Scalable SQL cannot initiate the transaction; you must recover a damaged file.

A file that you tried to open within a transaction is damaged and must be recovered. Complete or abort the current transaction. You can either use the file recovery procedure described in the *Pervasive.SQL User's Guide* or restore the file from a backup.

315: An application cannot change the dictionary or data path within a transaction.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

The application issued an xDDPath operation within an active transaction. This status code can also occur if the application attempts to execute a SET DDPATH or SET DATAPATH statement. You must commit or abort the transaction before you can change the current dictionary or data file path.



Note: The xDDPath primitive and the SET DATAPATH statement will be obsolete in future versions. The SET DDPATH statement is obsolete.

316: The Scalable SQL engine or Requester cannot unload.

In the DOS environment, Scalable SQL returns this status code for one of the following reasons:

- You attempted to unload Scalable SQL when you have loaded another terminate and stay resident (TSR) program after you loaded Scalable SQL. Unload the other TSR before unloading Scalable SQL.
- You attempted to unload Scalable SQL from a Rational Extended DOS application.
- You attempted to unload the Btrieve Requester before unloading the Scalable SQL Requester. You must unload the Scalable SQL Requester before you can unload the Btrieve Requester.

317: Another user has locked the dictionary within a transaction.

You cannot access the dictionary at this time because another user currently has it locked in a transaction. Retry the operation.

318: You cannot update a read-only view.

You cannot perform an Insert, Update, or Remove operation on a view that is defined as read-only. Refer to the *Pervasive.SQL Programmer's Guide* for information about the criteria for a view to be read-only.

319: A value in the list is invalid for the defined column attributes.

Any value in a list must agree with other attributes you have previously defined for the associated column. Query the X\$Attrib table to determine the attributes that have been defined for the column.

320: A character in the list is invalid for the defined column attributes.

Any character in a list must agree with other attributes you have previously defined for the associated column. Query the X\$Attrib table to determine the attributes that have been defined for the column.

321: The specified default value is invalid for the defined column attributes.

Any default value you define for a column must agree with other attributes you have previously defined for that column. Query the X\$Attrib table to see the attributes that have been defined for the column.

322: The specified range value is invalid for the defined column attributes.

Any range value you define for a column must agree with other attributes you have previously defined for that column. Query the X\$Attrib table to determine the attributes that have been defined for the column.

323: You cannot perform the specified operation on a table defined for an

active view.

You cannot perform the specified operation on a table that is currently in use in an active view. You can receive this status code if you attempt to change or remove the definition of a table that is currently in use by another application or another cursor ID in your application.

Pre-v4.0 Scalable SQL returns this status code if you request an operation that requires Scalable SQL to open a data file in Exclusive mode. The operation you requested cannot be completed because another task has opened the data file. In the server environment, you can use the Monitor utility to determine which task has opened the file.

324: The restriction clause or computed column contains an unmatched parenthesis.

This status code is obsolete in Scalable SQL versions 4.0 and later.

An unmatched parenthesis was found in a computed column or restriction clause. Make sure that all parentheses are matched, that the length of the restriction clause or computed column is correct, and that the restriction clause or computed column contains only ASCII text.

325: The total row length must be greater than three bytes and less than 32,765 bytes.

The total row length of a table definition (including any variable-length column) must be at least 4 bytes and less than 32,765 bytes. If you receive this status code after attempting to create a table with an LVAR column, you may have specified a column length greater than 32,761 bytes, which is the maximum length allowed. Refer to the topic of Record Lengths in *Pervasive.SQL Programmer's Guide*.

326: Scalable SQL has exhausted its internal stack space.

If you are defining a computed column or restriction clause, try reducing the size of the computed column or restriction clause.

327: At least one index must be defined without the NULL attribute.

This status code is obsolete in Scalable SQL versions 4.0 and later.

You must define at least one index in a table definition without the NULL attribute.

328: The syntax in a restriction clause or computed column using the DATE, TIME, or TIMESTAMP data type is invalid.

A restriction clause or computed column contains invalid syntax for DATE, TIME, or TIMESTAMP arithmetic.

329: The specified number of decimal places exceeds the column size.

When specifying the number of decimal places for the NUMERIC, NUMERICSA, or NUMERICSTS data type using the L,D (length, decimal places) format, the value for D (number of decimal places) must be less than or equal to L (column length). For example, NUMERIC (5,3) is valid, whereas NUMERIC (5,6) is invalid.

Because the DECIMAL data type is packed, the values you can use for the L,D format are different. Each digit takes one half byte of storage, and the sign takes one half byte. The internal size of a packed column is $M/2 + 1$, where M is the number of significant displayable digits. Thus, the following example allows nine significant digits with three decimal places (xxxxx.yyy):

DECIMAL (5,3)

The next example allows nine significant digits with six decimal places (xxx.yyyyyy):

DECIMAL (5,6)

The next example is invalid because the number of decimal places (12) exceeds the number of significant displayable digits (9):

DECIMAL (5,12)

This status code has the same meaning as Status Code [“806: The specified number of decimal places exceeds the column size.”](#).

330: The data is not formatted according to the appropriate mask.

If you are using XQLConvert (or xConvert), you must send the data to XQLConvert (or xConvert) formatted according to the default mask for that data type or the mask specified in the XQLConvert (or xConvert) sMask parameter.

If you are using a computed column (such as entering expressions in a column list), you must use the default mask for all data types.

If you are inserting rows, Scalable SQL does not allow you to insert rows in views that meet all of the following conditions:

- The view contains a table that has a LOGICAL column.
- The LOGICAL column has a default value assigned and a mask that uses character strings (such as 'yes–no').
- The LOGICAL column is not included in the view.

To avoid receiving this status code, either include the LOGICAL column in the view or reset the mask to 0–1 rather than character strings.

331: The operation performed on a group is invalid.

The following conditions cause this error:

- You tried to add a user to a group that does not exist.
- You attempted to alter the access rights for a user who is a member of a group. You must either modify the group's rights or first drop the user from the group and then redefine the user.
- You attempted to drop a group that still has members defined. You must remove each user from the group before you drop the group.

332: You must convert the dictionary for use with Scalable SQL v3.x.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The dictionary either is not valid or was created with pre-v3.x Scalable SQL.

333: Scalable SQL cannot store the index definitions.

Scalable SQL cannot store the index definitions for the current view. Retry the operation. If you continue to receive this status code, the INDEX.DDF dictionary file may be corrupt. You can either use the file recovery procedure described in the *Pervasive.SQL User's Guide* or restore the file from a backup.

334: Index names must be unique in the dictionary and must not match any column names.

The name you specified for a new index already exists as an index name or column name in the dictionary. Specify a unique name for the index.

335: The specified join option is invalid.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

Valid xJoin options are from 0 through 6, inclusive.

336: You cannot use xConvert or XQLConvert to convert variable-length data values.

You cannot use xConvert or XQLConvert to format or unformat data values of type LVAR and NOTE.

337: You have specified a pathname that exceeds the maximum length.

The maximum length of a pathname cannot exceed 64 characters.

338: You cannot enable or disable security when a view is active.

You can enable or disable security only when the database is inactive. Be sure all users have logged out of the dictionary and all views have been released before proceeding. You can use the Monitor utility to obtain this information.

339: You cannot perform an insert, update, or delete operation on a view containing a self join.

This status code is obsolete in Scalable SQL versions 4.0 and later.

You cannot perform an Insert, Update, or Delete operation on a view that joins a table to itself.

340: You cannot define a column of data type BIT as an index.

You cannot include a BIT data type column in an index.

341: An index of data type AUTOINC must be unique.

The AUTOINC data type is a special type that does not support duplicate values. It can be defined as part of a segmented index only if the key number of the AUTOINC key is less than the key number of the segmented key.

342: A parameter for the scalar function in an expression is invalid.

One of the parameters for the scalar function contains an invalid value or is of the wrong data type. Parameters that specify the length or offset values must be integer data types.

343: The operation is not valid on the group PUBLIC.

The PUBLIC group is reserved; you cannot create it or delete it from the dictionary. You cannot add users to the PUBLIC group.

344: The DOS rename function failed.

Scalable SQL returns this status code when you issue an ALTER TABLE statement and the existing data file cannot be replaced. The most common reason for receiving this status code is that the data file is flagged transactional.

345: The computed column name is invalid.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

When using the xCompute function, you cannot specify a computed name that is invalid (for example, all blanks). Also, for each computed column in a view definition, you cannot specify the same name more than once.

346: Scalable SQL has detected a math overflow.

If you receive this status code, retry the operation. If the error persists, attempt to isolate the problem to a particular SQL statement or relational primitive call.

347: A column descriptor is missing.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

For xDDCreate or xDDModify, you must define at least one column descriptor when using Option 0.

348: Scalable SQL cannot locate the named database you specified.

You specified a database name that Scalable SQL cannot find. First, ensure you entered the database name correctly. If the name is correct, try the following remedies:

- If you are using Windows, Windows 95, or Windows NT and you are looking for a database name defined in the local engine, ensure that the DBNAMES.CFG file is in the correct directory. Use the Setup utility to verify the directory and that the name is defined.
- If you are looking for a database name defined on the network, use SQLScope to determine whether or not the database name is advertised on the network. (In DOS, the DBNAMES.CFG file is in the directory specified under the DBNamesDirectory parameter in the BTI.CFG file.)
 - If the name is advertised, try logging in from SQLScope with that database name. If you can log in, the problem may be with your application. If you cannot log in when using the local engine, ensure that your machine environment is configured to allow use of remote engines. For more information about configuring your environment, refer to the *Pervasive.SQL User's Guide*.
 - If the name is not advertised, ensure that the server in which the name is defined is running Scalable SQL. Also ensure that the database name is configured on the server using the Setup utility. For more information about defining database names, refer to the *Pervasive.SQL User's Guide*.

349: The specified value exceeds the maximum length defined for the column.

A data value you specified in an INSERT or UPDATE statement or on an xInsert, xUpdall, or xUpdate call is larger than the maximum length defined for the column. This status code may also occur if Scalable SQL encounters an existing value in the database that is larger than the maximum value defined for the column. This may occur during

an ALTER TABLE MODIFY column, SELECT, INSERT, UPDATE, or DELETE statement.

When using the CAST function, Scalable SQL assumes the default data type unless you specify the size of the CAST data type equal to the data type on which the CAST is being performed. If you do not specify the size, you may receive Status Code 349. In the following example, the first statement returns Status Code 349. The second statement specifies the data type size.

Incorrect Example (results in Status Code 349):

```
SELECT CAST(SUBSTR(TEST, 1, 255) AS CHAR)
```

Correct Example:

```
SELECT CAST(SUBSTR(TEST, 1, 255) AS CHAR(255))
```

350: Security has not been enabled on this dictionary.

You attempted to perform an operation that you can complete only when security is enabled. Database security is not enabled at this time. To perform the operation successfully, you must first enable security on the relevant database.

351: A transaction has not yet been started.

If you are using SQL statements, you specified a COMMIT or ROLLBACK statement without being inside a transaction. You must start a transaction with START TRANSACTION before you can specify a COMMIT or ROLLBACK statement.

If you are using relational primitives, you attempted to perform Option 1 (End Transaction) or Option 2 (Abort Transaction) of xTrans without being inside a transaction. You must first start a transaction with one of these xTrans options: 0 (Start Transaction), 3 (Start Transaction Using Cursor Stability Isolation Level), or 4 (Start Transaction Using Exclusive Isolation Level).

352: The record count must be greater than zero.

The record count for an xFetch call, Option 0 through 5, must be at least 1.

353: The data definition includes an inconsistency.

The definition of the data file does not match the definition defined in the dictionary.



Note: If you received the dictionary from a third-party vendor, you must check with the vendor about resolving this discrepancy. Do not change the dictionary or file definitions unless you created them.

You receive this status code under the following conditions:

- For the column being updated, the offset into the record is past the end of the record length defined in the dictionary.
- When preparing to insert a record, Scalable SQL detects that the record length defined in the dictionary exceeds the record length of the data file.

For either condition, verify that the total size of the fixed-length columns equals the record length:

1. Use the SQLScope Table Statistics command to display the sizes of the fixed-length columns.
2. Add the sizes of the fixed-length columns. Adjust the size for any BIT columns. For BIT columns, up to eight contiguous bits are contained in one byte.
3. Obtain the record length of the file using either the Maintenance utility STAT command, SQLScope, or a Btrieve Status operation, depending on your operating environment.
4. Compare the total of the dictionary columns to the record length.
5. Determine whether the file or the dictionary definition is correct:
 - If the dictionary definition is correct, recreate the file to match the definition.
 - If the file is correct, use ALTER TABLE to update the dictionary definition to match the file.

354: The iCount parameter is not set correctly for the specified API.

The iCount parameter must be greater than or equal to zero and less than the maximum for the specified API. Check the description of the API to determine the correct value for the iCount parameter.

357: The DDF is out of sync or corrupt.

Scalable SQL returns this status code for one of the following reasons:

- The DDF file is corrupt. Replace the DDF file with a backup copy, or try to retrace by resorting back to the set of DDFs that came with your distribution.
- You attempted to log in to the dictionary (in v4.0 mode) before you converted the v3.01 VIEW.DDF file to v4.x format. Use the View Conversion utilities to convert the VIEW.DDF file or resubmit the text of the view. For more information about these utilities, refer to the *Pervasive.SQL User's Guide*.

358: Scalable SQL was unable to assign a value for the column X\$File.Xf\$Id.

Scalable SQL was unable to assign a unique, unused ID for the table being created. You cannot create more than 32,767 tables. You must drop another table before you can create a new one.

359: Scalable SQL was unable to assign a value for the column X\$Field.Xe\$Id.

Scalable SQL was unable to assign a unique, unused ID for the column being created for the table. You cannot create more than 32,767 columns for a table. You must drop another column before you can create a new one.

360: Scalable SQL was unable to assign a value for the column X\$Index.Xi\$Number.

Scalable SQL was unable to assign a unique, unused number for the index to be created for the table. You cannot create more than 32,767 indexes for a table. You must drop another index before you can create a new one.

361: The specified collation sequence is incompatible with the dictionary definition.

The collation sequence defined for the column is different from the collation sequence defined in the dictionary. You must use the collation sequence defined in the dictionary.

362: The attempted operation is not available in the current compatibility mode.

The operation you attempted is not permitted in the session's current compatibility mode. Following are possible causes of this status code:

- For applications in v3.01 mode, one of the following:
 - Attempt to create or invoke a stored procedure
 - Attempt to declare a trigger
 - Attempt to declare a session variable or session cursor
- For applications in v4.0 mode, one of the following:
 - Attempt to create or recall a stored statement
 - Attempt to execute a view primitive: xNew, xJoin, xField, xRestrict, xOrder, xCompute, xStore, xRecall, or xMovefld
 - Attempt to set the blank replacement character using the SET BLANK statement or xChar

363: This trigger does not exist in the dictionary.

The specified trigger does not exist. Check to be sure that you specified the name correctly. You can use the CREATE TRIGGER statement to create a trigger.

364: This trigger already exists in the dictionary.

While trying to create a trigger, you specified a trigger name that already exists. Trigger names must be unique in the dictionary. Specify a different, unique name for the trigger.

365: A trigger with the same order of execution already exists for the specified table and operation.

While trying to create a trigger, you specified an ORDER value for the trigger that is identical to the ORDER value for a trigger that is already defined. The ORDER value must be unique within each set of trigger definitions that specify the same values for the following:

- The same table
- The same trigger event (an INSERT, UPDATE, or DELETE operation)
- The same triggered action time (BEFORE or AFTER)

For example, you can have only one BEFORE INSERT ORDER 1 trigger for the Student table, but you can have both an AFTER INSERT ORDER 1 trigger and a BEFORE INSERT ORDER 1 trigger for the Student table.

366: The stored procedure already exists.

While trying to create a stored procedure, you specified a procedure name that already exists. Specify a different, unique name for the stored procedure.

367: You must drop the trigger before completing this operation.

A defined trigger contains a dependency that prohibits the current operation from succeeding. The current operation may be a DROP TABLE, DROP PROCEDURE, DROP VIEW, or an ALTER TABLE statement that adds, modifies, or drops a column. In order to complete the current operation, you must first drop the trigger that contains the dependency.

368: A table cannot contain a foreign key with DELETE CASCADE and also contain a DELETE trigger.

You cannot alter a table so that it contains both a DELETE trigger and a DELETE CASCADE foreign key.

369: No more users or groups are allowed.

You have reached the maximum number of allowed users or groups defined by a CREATE GROUP or GRANT LOGIN statement. The limit for the number of defined users or groups is the maximum for a two byte signed integer, which is approximately 32K. Scalable SQL also returns this status code when you use the primitive API xUser with options 0 and 4.

370: The user or group is in use.

The following operations are not allowed:

- The master user cannot change a user's privileges while that user is logged in.
- The master user cannot change a group's privileges while any user in that group is logged in.
- The master user cannot drop a user if that user is logged in.
- The master user cannot drop a group if any user in that group is logged in.
- A user cannot log in if the master user is changing that user's privileges.
- A user cannot log in if they are in a group and the master user is changing that group's privileges.
- The master user cannot change a user's password if that user is logged in.
- A user cannot change their own password if that user is logged in more than once.
- A user cannot read the privileges for the PUBLIC user while the master user is changing the privileges for the PUBLIC user.
- The Master cannot change the privileges for the PUBLIC user while some user is reading the privileges for the PUBLIC user.
- The master user cannot change any user, group, or PUBLIC privileges, if another login for the master user is changing that same user, group, or PUBLIC privileges.

The master user can only drop, change the password of, or change privileges for a user or group when the specified user or user group is not logged in. Use the Scalable SQL Monitor utility to determine if the specified user is logged in.

If you are trying to log in and you get this status code, retry the operation after ensuring that the master user has completed changing your privileges or password. If you are trying to change your password, ensure that no other application has logged in with your username and then retry the operation.

If you are not the master user, and you are attempting any other operation and you get this status code, retry the operation after ensuring that the master user has completed changing both your privileges and privileges for the PUBLIC user. If you are the master user, and you get this status code, first ensure that no other Master login is

trying to change the same user's privileges. Once no other Master login is trying to change the same user's privileges, retry the operation.

371: The stored procedure is in use.

You cannot drop a stored procedure while that procedure is in use by you or another user. You also cannot execute a procedure if it is being dropped by another user. Retry the DROP PROCEDURE, CALL, or EXECUTE statement at a later time. Ensure that no other users are accessing the procedure if you are trying to drop it, and ensure that no user is attempting to drop the procedure if you are trying to execute it.

373: The numeric value for the year must be between 1 and 9999, inclusive.

The numeric value you specify for the year in a TIMESTAMP column must be between 1 and 9999, inclusive.

380: The IN DICTIONARY clause is not allowed.

You cannot use the IN DICTIONARY clause when the table you are accessing is in a bound, named database. Scalable SQL only allows you to modify the table and the data file in a bound database.

381: Data file exists.

The CREATE TABLE statement found the data file already present, so the table was not added to the dictionary. This error only occurs in a bound database.

501: The SQL statement contains invalid syntax.

This is a general syntax error that Scalable SQL returns when it is unable to determine the intended statement type. Refer to the *SQL Language Reference* for more information about SQL syntax and SQL keywords.

502: The SELECT statement is invalid.

The SELECT statement contains a syntax error. Check the syntax of the statement and try again.

503: The function code is invalid.

Check the function code you are using. Valid function codes for the relational primitives are 0 to 41, inclusive; valid function codes for the SQL-level functions are 0 to 16, inclusive. If you have written your own interface to the relational primitives or Scalable SQL APIs, be sure to use the function number defined in the *Scalable SQL Programmer's Guide*. Reload the application, and try the operation again.

This status code has the same meaning as Status Code ["201: The relational primitive function code is invalid."](#)

504: The specified open mode is invalid.

You issued an invalid open mode in a SET OPENMODE statement. Valid open modes for data files are Accelerated, Exclusive, Normal, Read-Only, and Verify. For pre-v3.1 Scalable SQL, the valid values for the open mode are 0, -1, -2, -3, and -4. If you are using an older version of Scalable SQL, refer to the discussion of xNew in the *What's New*.

This status code has the same meaning as Status Code ["295: The specified open mode value is invalid."](#)

505: Each item in the attribute list must be enclosed in quotation marks.

This status code is obsolete in Scalable SQL versions 4.0 and later.

In a SET statement, you must use single quotation marks to enclose all the values for string data types in an attribute list. A range list can include one or more range specifications. A range specification consists of two values separated by a dash (–) with a space on either side of the dash. Multiple range specifications are separated with a comma.

506: The INSERT statement is invalid.

Scalable SQL detected a syntax error when it was parsing the INSERT statement. Verify that the syntax you are using is correct. If you are using the XQLCompile function, check the iStatementLen parameter for the approximate byte offset of the error.

507: The UPDATE statement is invalid.

Scalable SQL detected a syntax error when it was parsing the UPDATE statement. Verify that the syntax you are using is correct. If you are using the XQLCompile function, check the iStatementLen parameter for the approximate byte offset of the error.

508: The DELETE statement is invalid.

Scalable SQL detected a syntax error when it was parsing the DELETE statement. Verify that the syntax you are using is correct. If you are using the XQLCompile function, check the iStatementLen parameter for the approximate byte offset of the error.

509: The SET statement is invalid.

Check the SET statement you are using to verify that the syntax is correct.

510: The keyword following CREATE must be TABLE, VIEW, INDEX, DICTIONARY, GROUP, PROCEDURE, or TRIGGER.

The only keywords you can use following a CREATE keyword are TABLE, VIEW, INDEX, DICTIONARY, GROUP, PROCEDURE, or TRIGGER. Ensure that you have entered the text correctly. This status code can also indicate a general syntax error within one of these statement types. For more information about SQL statement syntax, refer to the *SQL Language Reference*.

511: The index name is missing or invalid.

You did not specify an index name in a CREATE INDEX statement, or the index name contains invalid characters. Refer to the *Pervasive.SQL Programmer's Guide* for more information about valid characters for index names.

512: The GRANT statement is invalid.

The GRANT statement contains invalid keywords or names. Make sure the syntax is correct and that you have spelled the usernames or group names correctly.

513: You must specify a username in a GRANT or REVOKE statement.

You must specify the names of the users or groups to which you are assigning access rights in a GRANT statement or from which you are revoking access rights in a REVOKE statement.

514: The table name is missing or invalid.

You did not specify a table name in a CREATE TABLE, GRANT (access rights), or REVOKE (access rights) statement, or the table name contains invalid characters. Refer to the *Pervasive.SQL Programmer's Guide* for more information about valid characters for table names.

515: The syntax of the ALTER TABLE statement is invalid.

An ALTER TABLE statement requires that you include the ALTER and TABLE keywords followed by a table name. You must include a column name after the ADD, DROP, or MODIFY keyword.

516: The ALTER TABLE statement requires either an ADD, DROP, MODIFY, or USING clause.

The ALTER TABLE statement allows you to change a table definition. You must include one of the following clauses in the statement: ADD, DROP, MODIFY, or USING.

517: You must include a valid FROM clause in the SELECT statement.

You either omitted the FROM clause from the SELECT statement, or the statement contains invalid syntax. A SELECT statement must always contain a FROM clause that specifies the table or view that you are querying. For more information about syntax, refer to the *SQL Language Reference*.

518: The sort keyword must follow a column name in an ORDER BY clause.

Either the column name is missing from an ORDER BY clause, or you positioned the sort keyword (ASC or DESC) incorrectly in the clause.

519: The AS keyword is missing, or a column heading is invalid.

You tried to execute a CREATE VIEW statement without an AS keyword to specify the SELECT clause, or a column heading contains invalid characters. Refer to the *Pervasive.SQL Programmer's Guide* for details about naming conventions.

520: The keyword BY must follow ORDER and GROUP.

You must follow ORDER and GROUP with the BY keyword.

521: The view name is missing or invalid.

You did not specify a view name in a CREATE VIEW statement, or the view name contains invalid characters. Refer to the *Pervasive.SQL Programmer's Guide* for details about naming conventions.

522: The statement contains an unknown keyword.

The syntax for the SQL statement is incorrect. Review the statement and verify that the keywords are in the proper order. If you are using the XQLCompile function, check the iStatementLen parameter for the approximate byte offset of the error.

523: The access privilege keyword following GRANT or REVOKE is missing or invalid.

Scalable SQL does not recognize the access privilege keyword in a GRANT or REVOKE statement. You must specify one of the following keywords after the GRANT or REVOKE keyword:

ALL	DELETE	REFERENCES
ALTER	INSERT	SELECT

524: You can grant column-level privileges for Select and Update rights only.

Column-level privileges are allowed only for Select and Update rights. Do not specify a column list when granting All, Alter, Insert, Delete, or References rights. These rights affect every column in the table.

525: You must specify a master password in a SET SECURITY statement.

You did not specify a master password in a SET SECURITY statement, or you specified an illegal string. The password must be the master password for the dictionary.

526: The column name cannot be a Scalable SQL keyword.

A column name cannot be a Scalable SQL keyword. For a list of Scalable SQL keywords, refer to the *SQL Language Reference*.

527: You must specify the new user's password in a GRANT LOGIN statement.

When you create a new user using the GRANT LOGIN statement, you must assign a password to the user.

528: You must provide a valid value enclosed in single quotation marks after the equal sign.

The SET statement is incomplete if it does not have a valid value enclosed in single quotation marks or the keyword NULL after the equal sign. You can also receive this status code if you log into a database (using xLogin or XQLLogin) in v3.01 mode and specify an invalid blank character. Valid blank replacement characters in v3.01 mode are underscore (`_`), tilde (`~`), and caret (`^`).

529: You must supply one or more owner names after the equal sign.

An owner name, or a list of owner names, must follow the equal sign in the SET OWNER statement.

530: You must supply a dictionary path after the equal sign.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The pathname to the dictionary in a SET DDPATH statement is missing. The SET DDPATH statement is obsolete in Scalable SQL v4.0 and later.

531: You must supply a data file path after the equal sign.

The pathname to the data files in a SET DATAPATH statement is missing.



Note: The SET DATAPATH statement will become obsolete in future versions.

532: The SET statement must contain an equal sign.

Ensure that the SET statement contains an equal sign.

533: The specified data type is invalid.

Scalable SQL does not support the data type you specified. For a list of valid data types, refer to the *SQL Language Reference*.

534: You must include SELECT in a CREATE VIEW statement.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The SELECT keyword is missing from a CREATE VIEW statement. The CREATE VIEW statement allows you to store a SQL statement that begins with a SELECT keyword.

535: When defining an index, you must include index attributes.

When defining indexes for a table definition, the index attributes must follow the column name and precede the ACS (if present), comma (if present), or closing parenthesis.

536: The specified name is invalid.

When specifying a table, column, index, trigger, procedure, foreign key, user, or user group name, ensure that the name does not conflict with a reserved word or symbol used by Scalable SQL. If the name you specify conflicts with a reserved word, either specify a different name or enclose the name in double quotes.

537: The specified column name already exists in this definition.

All column names in a table definition must be unique. Verify that you have not entered the same column name more than once in the current table definition.

538: The column name you specified for the index definition is invalid.

You have attempted to define an index in the WITH INDEX clause of CREATE TABLE using a column name that is not in the table definition. Ensure that you have entered the column name correctly.

539: The specified index is already defined in this definition.

You have attempted to define the same index more than once in this definition. Rewrite the statement so the index is defined only once.

540: A column list must follow SELECT, ORDER BY, and GROUP BY.

The SELECT keyword must be followed by a selection list to specify the columns to retrieve. This selection list can be one or more column names, expressions, or aggregate functions. The ORDER BY and GROUP BY keywords must be followed by one or more column names or column positions.

541: A restriction clause must follow WHERE and HAVING.

You must include a restriction clause after the WHERE and HAVING keywords.

542: The keyword following DROP must be DICTIONARY, GROUP, INDEX,

PROCEDURE, TABLE, TRIGGER, or VIEW.

You can only drop a dictionary, group, index, procedure, table, trigger, or view.

543: You must enclose a string in single quotation marks.

This status code is obsolete in Scalable SQL versions 4.0 and later.

Make sure you enclose a string in single quotation marks.

544: You must specify a directory path in a DROP DICTIONARY statement.

When dropping a dictionary using DROP DICTIONARY, you must specify the directory path to the dictionary you are dropping. You cannot drop the dictionary that you are currently using.

545: You must include INTO in an INSERT statement.

The INTO keyword, which must be followed by a table or view name, is missing from an INSERT statement.

546: You exceeded the maximum token size allowed in a SQL statement.

Your statement included a token (probably a name or string) that exceeded the maximum allowable size of 255 bytes. Scalable SQL also returns this status code if you submit a large, complex statement that exhausts the parser's internal table space.

547: You must include a column attribute in a SET statement.

Scalable SQL expects a column attribute in a SET DEFAULT or SET MASK statement. Scalable SQL expects a list of column attributes in a SET CHAR, SET RANGE, or SET VALUES statement. This status code is specific to those types of SET statements.

548: TRANSACTION must follow START; COMMIT or ROLLBACK must precede WORK.

Always follow the START keyword with the TRANSACTION keyword, and precede the WORK keyword with the COMMIT or ROLLBACK keyword. You cannot use the AND CHAIN TO SAVEPOINT clause in a ROLLBACK WORK statement.

549: The ORDER BY clause must be specified last in a SELECT statement.

The ORDER BY clause must be the last clause in a SELECT statement.

550: An alias name has already been defined for the table.

This status code is obsolete in Scalable SQL versions 4.0 and later.

You have attempted to assign multiple aliases to the same occurrence of a table name in the SQL statement. If the FROM clause contains more than one table name, make sure a comma separates each alias from the following table name.

551: The specified dictionary name contains invalid characters.

The dictionary name you specified for a data element contains invalid characters. Refer to the *Pervasive.SQL*

Programmer's Guide for details about naming conventions.

552: A SELECT clause must follow UNION.

You must include a valid SELECT clause after the UNION keyword.

553: The database element name exceeds its character limit.

The name of the database element (such as a dictionary column or stored procedure) is too long. Ensure that the name does not exceed its maximum number of characters. For more information about the maximum lengths of database elements, refer to the *Scalable Pervasive.SQL Programmer's Guide*.

554: You must specify a group name with CREATE GROUP, DROP GROUP, or IN GROUP.

Scalable SQL expects a group name after a CREATE GROUP or DROP GROUP statement or after the IN GROUP clause of GRANT LOGIN. Ensure that the name you specify is a valid group name.

555: The keyword DICTIONARY must follow IN in an ALTER TABLE statement.

In an ALTER TABLE statement, the keyword DICTIONARY must follow IN.

556: A parameter in the CREATE TABLE statement is invalid.

The page size, preallocation, or threshold value in the CREATE TABLE statement may be invalid. Check the values you specified to be sure they are within the valid ranges. If you are using the data types DECIMAL, MONEY, NUMERIC, NUMERICSA, or NUMERICSTS, check for an invalid decimal count specification. Also, ensure that there are no double quotes surrounding pathnames.

557: Scalable SQL cannot access the alternate collating sequence file.

Scalable SQL cannot open the ACS file. Ensure that you entered a valid pathname for the file.

558: You cannot use a group aggregate function within an expression.

This status code is obsolete in Scalable SQL versions 4.0 and later.

You cannot use an aggregate function (AVG, COUNT, MIN, MAX, and SUM) within an expression.

559: The syntax in the stored procedure or trigger is invalid.

Scalable SQL usually returns a more specific syntax error when compiling an invalid statement within a stored procedure or trigger. This is a general error that Scalable SQL returns when it is unable to determine the intended statement type. In a trigger, the syntax error may also refer to an invalid time or event or other invalid syntax in the trigger definition. For more information about SQL syntax in stored procedures or triggers, refer to the *SQL Language Reference*.

560: Scalable SQL encountered an error while attempting to create a stored SQL statement.

If you receive this status code, try executing the SQL statement(s) before attempting to store the statement(s) in a stored SQL statement. If the statement(s) execute without error, try recovering the PROC.DDF file or restoring it from

backup.

561: A CREATE PROCEDURE statement must include a name for the stored SQL statement.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

You must follow the CREATE PROCEDURE keyword with a name; you use this name to call the stored statement when you execute it.

562: The parameter name must precede the parameter value.

The format for specifying parameters in a stored SQL statement is as follows:

```
@parameter_name [data_type] [= value]
```

The at symbol (@) must precede the parameter name.

563: You specified a default for a parameter that cannot have a default value.

This status code is obsolete in Scalable SQL versions 4.0 and later.

You tried to specify a default for a parameter that cannot have a default value. Remove the default value and try the statement again.

564: You must declare all parameters when you create the stored SQL statement.

When executing a stored SQL statement, you can provide values only for parameters that were declared when the stored SQL statement was created. This status code has the same meaning as Status Code [“864: The specified parameter name cannot be found.”](#).

565: The specified isolation level is invalid.

The valid isolation levels for a SET ISOLATION statement are EX for exclusive isolation level and CS for cursor stability isolation level.

566: You must specify an isolation level in a SET ISOLATION statement.

On a SET ISOLATION statement, you must specify an isolation level. Use EX for exclusive or CS for cursor stability.

567: The syntax for the primary or foreign key definition is invalid.

The CREATE TABLE or ALTER TABLE statement contains invalid syntax for the primary or foreign key definition. Check the syntax of the statement as defined in the *SQL Language Reference*.

568: A CREATE TABLE statement cannot reference the same table more than once.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The CREATE TABLE statement you issued contains multiple foreign key references to the same table. If you must create more than one foreign key that refers to the same table, use the ADD FOREIGN KEY clause of the ALTER

TABLE statement to add the additional keys.

569: A CREATE VIEW statement contains invalid syntax.

The syntax for the CREATE VIEW statement is not specified correctly. One possible cause of this status code is that the CREATE VIEW statement contains a UNION clause, which is not supported. Check the syntax of the statement as defined in the *SQL Language Reference*.

570: The CREATE TABLE or ALTER TABLE statement contains a duplicate, redundant, or conflicting option.

You specified a duplicate, redundant, or conflicting option in a CREATE TABLE or ALTER TABLE statement. For example, in the following statement, DCOMPRESS is duplicated:

```
CREATE TABLE Table1 DCOMPRESS DCOMPRESS (c1 INT(8))
```

For more information about the syntax of the CREATE TABLE and ALTER TABLE statements, refer to the *SQL Language Reference*.

571: Duplicate name not allowed in SELECT list.

You attempted to assign a name to a column in the SELECT list using the AS clause when the name is already in use. The name must be unique within the SELECT list.

800: The memory allocation is insufficient for Scalable SQL to process the current request.

Scalable SQL is unable to allocate sufficient buffer space to process the current request. If you have other cursors active, attempt to release some of them to free memory for the current request. This status code has the same meaning as Status Code

802: The Scalable SQL local engine or Requester is not loaded.

You must load the Scalable SQL local engine or the Requester on your machine. For more information about loading these components, refer to the *Pervasive.SQL User's Guide*.

When using a 16-bit application such as SQLScope to log in to a local Windows NT engine, thinking must be turned on.



To Enable Thinking in Scalable SQL:

1. Click **Start**, point to **Programs**, then to the **Pervasive SQL 7** folder.
2. Select **Setup (Win16)**.
3. Select the **Scalable SQL Requester** category.
4. Select the **Use Think** setting.
5. Ensure that **Use Think** is set to Yes.

6. Click **Save** and then **Exit**.
7. Stop and restart the Scalable SQL engine as described in *Getting Started With Pervasive.SQL*.

803: You must obtain a valid session ID before attempting to call the function.

You attempted to call a Scalable SQL function before calling XQLLogin. XQLLogin must be the first Scalable SQL function that you call from the application. This status code has the same meaning as Status Code [“265: The session identifier is invalid.”](#).

804: The specified column position is invalid.

There is no column at the position that you specified. Ensure that you specified the correct position. Scalable SQL also returns this status code from an XQLDescribe call to inform an application that the last column has been described.

805: The specified size for the data type is invalid.

The size for the data type is not a valid size. For information about data type sizes, refer to the *SQL Language Reference*. This status code has the same meaning as Status Code [“248: The specified size is invalid for the defined data type.”](#).

806: The specified number of decimal places exceeds the column size.

When specifying the number of decimal places for the NUMERIC, NUMERICSA, or NUMERICSTS data type using the L,D (length, decimal places) format, the value for D (number of decimal places) must be less than or equal to L (column length). For example, NUMERIC (5,3) is valid, whereas NUMERIC (5,6) is invalid.

Because the DECIMAL data type is packed, the values that you can use for the L,D format are different. Each digit takes one half byte of storage, and the sign takes one half byte. The internal size of a packed column is $M/2 + 1$, where M is the number of significant displayable digits. Thus, the following example allows nine significant digits with three decimal places (xxxxxx.yyy):

DECIMAL (5,3)

The next example allows nine significant digits with six decimal places (xxx.yyyyyy):

DECIMAL (5,6)

The next example is invalid because the number of decimal places (12) exceeds the number of significant displayable digits (9):

DECIMAL (5,12)

807: The variable reference is invalid.

The operation referred to a variable name that could not be resolved within the current scope. This status code can refer to two different kinds of variables:

- A SQL variable reference in an assignment statement, target list, or other context that Scalable SQL could not resolve, such as:
 - Inside a procedure—Using the procedure’s variables or parameters.
 - Inside or outside a procedure—Using the session’s variables.

- A substitution variable name in an XQLSubst call.

808: All variables were not replaced with values.

This status code can refer to two different kinds of variables:

- If your statement has substitution variables, you cannot issue an XQLExec call until you have supplied values for all the substitution variables in the SQL statement using XQLSubst.
- If your statement includes session or procedure variables, you cannot attempt to use the values of the variables unless they are assigned values or given default values when declared.

809: The specified column is not in the recalled view.

This status code is obsolete in Scalable SQL versions 4.0 and later.

One or more of the columns you specified in a SELECT statement is not defined in the recalled view.

810: The specified numeric data contains invalid characters.

Numeric data can consist only of the ASCII characters 0 to 9, the period (.), plus sign (+), and minus sign (-). If the data type is MONEY, the dollar sign (\$) is also allowed. This status code has the same meaning as Status Code [“224: A character in the numeric data is invalid.”](#).

811: The specified numeric value for the month is invalid.

The numeric value for the month must be between 1 and 12, inclusive. This status code has the same meaning as Status Code [“292: The numeric value for the month is invalid.”](#).

812: The specified numeric value for the day of the month is invalid.

The numeric value for the day must be between 1 and the number of days in the month, inclusive. This status code has the same meaning as Status Code [“293: The numeric value for the day of the month is invalid.”](#).

813: The specified numeric time value is not within the range of valid values.

The numeric value for hours must be between 0 and 23, inclusive. The numeric value for minutes and seconds must be between 0 and 59, inclusive. The numeric value for hundredths of a second must be between 0 and 99, inclusive.

815: The specified join column is invalid.

This status code is obsolete in Scalable SQL versions 4.0 and later.

A column you specified in the join condition of the WHERE clause is not defined in any of the tables specified in the SQL statement.

816: The specified column is not defined in this table.

You have attempted to modify, drop, or define a column constraint for a column that does not exist in the table you specified.

818: The select terms must be either group aggregate functions or columns specified in a GROUP BY clause.

When a SQL statement contains group aggregates, the selection list of that statement can contain only group aggregates or columns that are specified in the GROUP BY clause. If there is no GROUP BY clause, the selection list can contain only group aggregates.

819: The specified table name is invalid.

The table name you specified is not in the currently active dictionary. Be sure that you have specified the name correctly and that you use the appropriate blank replacement character if the name contains embedded blanks. This status code has the same meaning as Status Code [“204: The table is not defined in the dictionary.”](#).

820: You must include a comparison operator in the HAVING clause.

The restriction clause for the HAVING clause is incomplete. Ensure that you have included at least two expressions separated by a comparison operator.

821: The specified cursor ID is invalid.

Scalable SQL returns this status code for one of the following reasons:

- The iCursorID parameter does not contain a valid ID. You must receive the ID from XQLCursor and free it using XQLFree.
- Although the cursor ID is valid, it is not currently valid for the specified operation. For example, you cannot call XQLFetch without first successfully compiling a SELECT statement using that cursor ID.

822: The data buffer is too small to hold the requested data.

The data buffer for the XQLFetch call is not large enough to hold all the requested data. This status code has the same meaning as Status Code [“210: The buffer is not large enough to hold all the data.”](#).

823: The position you specified in the ORDER BY clause is invalid.

When you specified ordering by position in the selection list, you specified a position that is out of range. The valid positions are 1 through the number of columns in the list, inclusive.

824: The statement contains an unmatched parenthesis.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The number of left and right parentheses must be equal.

825: You cannot use the AVG or SUM functions on STRING, DATE, TIME, or TIMESTAMP data types.

Because the AVG and SUM functions perform arithmetic operations, you can use these aggregate functions only on the numeric data types (AUTOINC, BFLOAT, CURRENCY, DECIMAL, FLOAT, INTEGER, MONEY, NUMERIC, NUMERICSA, NUMERICSTS, and UNSIGNED). You can use the COUNT, MIN, and MAX aggregate functions on any data type.

826: The statement contains an unmatched parenthesis.

The statement contains an unequal number of left and right parentheses. The number of left and right parentheses must be the same. This status code can occur in any statement that requires parentheses around a portion of the statement (for example, the column list in an INSERT statement, an expression in a WHERE clause, or the column list in a CREATE TABLE statement). Make sure that the number of left and right parentheses are the same and re-

submit your statement.

827: The columns in the GROUP BY control column list must be defined in the select list.

Check the GROUP BY columns to ensure that they match the columns in the selection list.

828: The WHERE clause must be the last clause in the statement.

The WHERE clause must be the last clause in a DELETE or UPDATE statement.

829: The syntax for the computed column definition is invalid.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

A computed column can be a constant value or the result of arithmetic operations performed on one or more columns in the selection list. You can also build a computed column by concatenating several string columns. Verify the syntax of the computed column as defined in the *SQL Language Reference*, and make sure you put a comma after each item except the last one in the selection list.

830: The HAVING clause requires another condition or contains an incomplete condition.

A HAVING clause must specify another condition after AND or OR. Also, a HAVING clause cannot contain an incomplete condition, such as HAVING MIN(a).

831: An expression in a HAVING condition is invalid.

The HAVING clause restricts the view for group aggregates. The HAVING clause can contain multiple conditions. The first expression in a condition must be an aggregate function value. The second expression can be a substitution variable, a string or numeric constant, or a subquery. You must separate the expressions with a range or relational operator; you must separate multiple conditions with either AND or OR. Use the WHERE clause to restrict the view for columns that are not aggregates.

832: The statement cannot be recalled because it is not of the expected size.

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

The contents of the view buffer do not match the expected size. This usually indicates that a disk error occurred while the view was being saved or recalled. Retry the operation. If you continue to receive this status code, the VIEW.DDF dictionary file may be corrupt. Try the following remedies:

- Use the file recovery procedure described in the *Pervasive.SQL User's Guide*.
- Restore the file from a backup.
- Rebuild the VIEW.DDF file (use CREATE VIEW statements).

833: The columns in the ORDER BY clause must be defined in the select list.

The columns in the ORDER BY clause must be defined in the selection list if they are not defined as indexes in the database.

835: The specified option for XQLStatus is invalid.

Valid options for XQLStatus are 0 and 1.

836: The specified mask is invalid for the column's data type.

The mask you specified does not contain valid mask characters for the field's data type. Check the mask to make sure the characters in it are valid. This status code has the same meaning as Status Code [“279: The mask is invalid for the data type specified.”](#).

837: The delimiter separating the column size and decimal place values is invalid.

In a CREATE TABLE or ALTER TABLE statement, use a comma to separate the column size and the number of decimal places for the DECIMAL, MONEY, NUMERIC, NUMERICSA, and NUMERICSTS data types. In addition, you must separate the column definitions with commas.

838: In a union, the number of columns in each SELECT statement must be the same.

When you combine multiple queries with UNION, ensure that each SELECT statement contains the same number of columns.

839: In a union, the corresponding columns must be compatible data types.

When you combine multiple queries with UNION, the data types of corresponding columns in the SELECT statements must be compatible. The following matches are compatible:

- Columns of the same data type and size
- String columns (CHARACTER, LSTRING, LVAR, NOTE, and ZSTRING) with other string columns
- Numeric columns (AUTOINC, BFLOAT, CURRENCY, DECIMAL, FLOAT, INTEGER, MONEY, NUMERIC, NUMERICSA, NUMERICSTS, and UNSIGNED) with other numeric columns
- Boolean columns (BIT and LOGICAL) with other Boolean columns
- Date columns (DATE) with other date columns
- Time columns (TIME) with other time columns
- Timestamp columns (TIMESTAMP) with other timestamp columns

840: Scalable SQL has exhausted its internal stack space.

If you are defining a computed column or restriction clause, try reducing the size of the computed column expression or the restriction clause.

841: Scalable SQL cannot create a temporary sort file.

Ensure that you have specified a valid external sort file directory using the Setup utility. For more information about how to do this, refer to the *Pervasive.SQL User's Guide*.

842: The HAVING clause cannot contain a correlated subquery.

A HAVING clause cannot contain a correlated subquery. Rewrite the SELECT statement without a correlated

subquery in the HAVING clause.

843: An OS path must be specified if Add, Modify, or Drop is not entered.

This status code is obsolete in Scalable SQL versions 4.0 and later.

Use the ALTER TABLE statement to modify a dictionary definition or file path. If you do not specify the ADD, MODIFY, or DROP keywords, Scalable SQL assumes that you only want to modify the file path, and expects an operating system pathname.

844: Invalid number of rows selected.

A SELECT statement returned an invalid number of records for the context. A SELECT statement on the left hand or right hand side of a relational operator must select a single record. A SELECT statement in an assignment statement must select a single record. The following examples cause this status code:

- `SELECT * FROM table1 WHERE
column1 > (SELECT column2 FROM table2)`

If (SELECT column2 FROM table2) selects 0 or more than one record then this status code is returned. In this case, the SELECT statement is a subquery. The SELECT may also appear to the left or right of a relational operator (=, <>, >, >=, <, <=) in the condition of an IF in a stored procedure.

- `SET variable = SELECT column1 FROM table1`

If (SELECT column1 FROM table1) selects 0 or more than one record then this status code is returned.

845: The number of column headings must match the number of columns in the view.

In a CREATE VIEW statement, the number of column headings in the view must match the number of columns in the view. You must specify a heading for all computed columns in a view; they are required so that you can explicitly reference the computed column. If you do not specify headings for computed columns, Scalable SQL returns this status code, and you cannot create the view.

846: You must specify a column name for the group aggregate function.

You must specify a column name for the aggregate functions AVG, MIN, MAX, or SUM.

847: The syntax in the WHERE clause is invalid.

A WHERE clause contains invalid syntax. Several conditions can cause Scalable SQL to return this status code:

- The NOT keyword is not followed by BETWEEN, CONTAINS, EXISTS, IN, or LIKE.
- The BEGINS keyword is not followed by WITH.
- No condition operator appears after a column name or constant.
- A column name is included in the list of values for an IN condition.
- The values for a BETWEEN condition are not connected by an AND keyword.

Scalable SQL also returns this status code when a statement contains something in place of an expected WHERE clause.

848: A statement may contain only one DISTINCT keyword.

This status code is obsolete in Scalable SQL versions 4.0 and later.

In a single SELECT statement, you cannot specify DISTINCT at both the view level and the group aggregate level.

849: The buffer is not large enough for Scalable SQL to process the WHERE clause.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The internal buffer that Scalable SQL uses to reconstruct the restriction clause of the WHERE clause is not large enough. If possible, reduce the size of the restriction clause or reduce the size of the dictionary names.

850: Scalable SQL cannot convert the data values to the new data type.

The new data type for the column is not compatible with the previous data type. For example, Scalable SQL cannot change a string data type (CHARACTER, LSTRING, LVAR, NOTE, or ZSTRING) to a numeric data type (AUTOINC, BFLOAT, CURRENCY, DECIMAL, FLOAT, INTEGER, MONEY, NUMERIC, NUMERICSA, NUMERICSTS, or UNSIGNED). Scalable SQL also returns this status code if you reduce the size of a column and an existing column value is longer than the new size.

851: In an INSERT or UPDATE statement, a subquery cannot access the table being updated.

This status code is obsolete in Scalable SQL versions 4.0 and later.

The table referenced in the INSERT or UPDATE statement cannot also be referenced in the FROM clause of the SELECT clause.

852: The record count must be greater than zero.

The record count for an XQLFetch call, Option 0 through 4, must be greater than zero.

853: You do not have the necessary rights to access the specified table.

You do not have the necessary Scalable SQL security rights for the specified table. To determine which rights you need for the operation you are attempting, refer to the discussion of security in the *Pervasive.SQL Programmer's Guide*.

854: At least one index must be defined for this operation.

This status code is obsolete in Scalable SQL versions 4.0 and later.

A CREATE INDEX statement must specify at least one index name.

856: The number of values cannot exceed the number of columns in the view.

The number of the values must match the number of columns explicitly placed in the view and must not exceed the number of columns implicitly placed in the view. The following examples cause this status code:

- ```
INSERT INTO table1 (column1, column2)
VALUES (value1, value2, value3)
```

- INSERT INTO table1 (column1, column2)

VALUES (value1)

- Assume table1 has only three columns. If the columns are not explicitly specified in the INSERT statement, then all columns are implicitly specified.

INSERT INTO table1

VALUES (value1, value2, value3, value4)

Note that INSERT INTO table1 VALUES (value1) is allowed.

### **857: You cannot use the equal to (=) operator with ALL in a WHERE clause.**

This status code is obsolete in Scalable SQL versions 4.0 and later.

Make sure that you do not use the equal (=) operator with the ALL keyword in a WHERE clause.

### **858: You can use CASE only with string data types.**

Only string columns can be case-sensitive.

### **859: The ORDER BY clause must specify the column position.**

If the ORDER BY clause is part of a SELECT clause in a union, or if the selection list contains a computed column, you must use the column position (not the column name) to specify the columns by which to order the rows.

### **860: The buffer is not large enough for Scalable SQL to store the view definition.**

This status code is obsolete in Scalable SQL versions 4.0 and later.

The internal buffer that Scalable SQL uses to store a view is not large enough to hold all the data. You can reduce the size of the definition by reducing the number of columns in the view or by making the size of the column headings smaller.

### **861: A value must follow the equal sign (=).**

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

When you execute a stored SQL statement, you must specify a value after the equal sign if you use the syntax `@parameter_name [= value]`. For more information about the EXECUTE command, refer to the *What's New*.

### **862: The specified stored SQL statement name is not in the dictionary.**

Either the name you selected is an invalid name for the stored SQL statement, or the stored SQL statement has been removed from the dictionary.

### **863: Scalable SQL cannot recall the stored SQL statement.**

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

This status code usually indicates the dictionary definitions for a table or view referenced within a stored SQL statement have changed since the stored statement was created. Recreate the stored statement and retry the

operation.

### **864: The specified parameter name cannot be found.**

When invoking a stored procedure, you attempted to match the calling parameters using a name that was not defined in the called procedure.

### **865: You cannot update columns that exceed 255 characters by using an UPDATE statement.**

This status code is obsolete in Scalable SQL versions 4.0 and later.

To update a variablelength column that is longer than 255 characters, use the xUpdate relational primitive.

### **866: The XQLSubst function call is invalid.**

This status code is obsolete in Scalable SQL versions 4.0 and later, except when using v3.01 Compatibility mode.

You can only use XQLSubst to substitute values for variable names specified in a SELECT, UPDATE, INSERT, or DELETE statement. You cannot use the XQLSubst function with a stored SQL statement.

### **867: The stored SQL statement was not created with the version of Scalable SQL you are currently using.**

You either stored a statement with Scalable SQL v3.01 and then tried to execute it with an earlier version, or you tried to execute a v3.01 statement with the v4.0 engine. Ensure you are using the same version of the Scalable SQL engine when you execute the stored statement.

### **868: Scalable SQL cannot read or update the X\$Relate system table.**

Scalable SQL cannot access the RELATE.DDF file. Try the operation again. If you continue to receive this status code, the RELATE.DDF file may be corrupt. You can either use the file recovery procedure described in the *Pervasive.SQL User's Guide* or restore the file from a backup.

### **869: No primary key is defined on the parent table.**

The table name specified after the REFERENCES keyword of the FOREIGN KEY definition does not have a primary key defined. Make sure you have specified the correct table name after the REFERENCES keyword.

### **870: The index must exist before you can use ALTER TABLE to define a primary or foreign key.**

Before you can add a primary key or foreign key with the ALTER TABLE statement, you must create an index for the columns in the key. The index for a primary key must be unique and non-null; the index for a foreign key must be non-null. Use the CREATE INDEX statement to create the index, then add the key again.

### **871: The primary key must be a unique, non-null index.**

The group of columns you have defined as the primary key do not provide a unique, non-null index for the table. Redefine the primary key so that it is unique and non-null.

### **872: The specified foreign key name already exists in the dictionary.**

If you omitted the foreign key name, Scalable SQL tried to use the name of the first column in the key as the foreign key name, and there is already a key named for that column. Provide a unique name for the foreign key and try to create it again.

### **873: The parent table must exist before you define a foreign key.**

The table named after the REFERENCES keyword in the FOREIGN KEY definition does not yet exist. Before you can add the foreign key, you must create the table with a primary key that you can reference with the foreign key.

### **874: The referential integrity constraint does not exist.**

Either the foreign key or the primary key you are attempting to drop is not defined. If you are attempting to drop a foreign key, ensure that you spelled the key name correctly. If you are attempting to drop a primary key, ensure that you spelled the specified table name correctly, and that it has a primary key.

### **875: You must place the plus sign (+) for an outer join after the rightmost column in the join condition.**

The syntax for an outer join is table1.column = table2.column(+). The syntax table1.column(+) = table2.column is incorrect.

### **876: You cannot drop a primary key that is referenced.**

You cannot drop the primary key until you drop all foreign keys that reference it.

### **877: Before performing an RI operation, you must log in to a database using the database name.**

Before you try to perform an RI operation, make sure that you log in to the database using a database name.

### **878: A referenced table does not exist.**

Scalable SQL cannot find a table that is referenced by another table. Check the RI constraints on the database. For information about how to do this, refer to the *Pervasive.SQL User's Guide*.

### **879: The foreign and primary keys must match exactly.**

The columns in the foreign key must be the same types and sizes and in the same order as the columns in the primary key. The only exception is that the foreign key can contain a column of type INTEGER to refer to an AUTOINC column in the primary key. However, the INTEGER column in the foreign key must be the same size as the AUTOINC column.

### **880: You cannot define a key as both a foreign key and a primary key.**

You cannot define the same key as both the primary key and a foreign key in the CREATE TABLE statement. Create the table with the primary key only, then add the foreign key using an ALTER TABLE statement.

### **881: You must remove all referential constraints before performing the operation.**

You must remove all the referential constraints from the table before you can drop the index or add, drop, or modify any columns.

## **882: The delete rule for self-referencing tables must be cascade.**

Ensure that the delete rule on a foreign key that references the same table's primary key is defined as cascade instead of restrict.

## **883: The delete rule for at least two of the foreign key connections in the cycle must be restrict.**

Ensure that the delete rule for at least two of the foreign key connections in the cycle is restrict instead of cascade.

## **884: The delete rules for multiple paths to the same parent table must match.**

When multiple dependent tables have delete paths to the same parent table, ensure that the delete rules for each path are the same. They must be either all restrict or all cascade.

## **885: Scalable SQL allows a maximum of 16 adjacent delete cascade rules in a reference path.**

You cannot have more than 16 adjacent delete cascade rules in a reference path. If your reference path contains more than 16 tables, use restrict instead of cascade for at least one of the delete rules.

## **886: The MicroKernel data file version does not support referential integrity.**

The MicroKernel encountered an error when it attempted to add RI information to a file. This status code can occur for one of two reasons:

- The file is in pre-v6.0 format. Pre-v6.0 file formats do not support RI. For information about converting v5.x files to v6.x, or either format to v7.x format, refer to the *What's New*.
- The version of the MicroKernel accessing the file is earlier than v6.0. Scalable SQL v3.0 and later cannot run with pre-v6.0 MicroKernels.

## **887: A CREATE TABLE statement cannot reference the same table more than once.**

The CREATE TABLE statement you issued contains multiple foreign key references to the same table. Check that you have specified the parameters for the CALL statement correctly. Refer to the CALL statement syntax in the *SQL Language Reference* for more information. If you must create more than one foreign key that refers to the same table, use the ADD FOREIGN KEY clause of the ALTER TABLE statement to add the additional keys.

## **888: A primary key is already defined on this table.**

You attempted to define a primary key on a table that already has a primary key defined. A table can contain only one primary key.

## **889: The column is not in the referenced table.**

A column comprising the primary key is not part of the table on which the primary key is being defined. This status code typically occurs within a CREATE TABLE or ALTER TABLE statement.

## **890: The label reference is invalid.**

A LEAVE, END LOOP, or END statement specified a label that was not previously defined. This status code can only

occur within a CREATE PROCEDURE or CREATE TRIGGER statement.

### **891: The condition reference or declaration is invalid.**

The referenced condition label is invalid either because it was not previously declared or because the declaration statement itself was invalid.

### **892: The savepoint reference is invalid.**

A statement referred to a savepoint that was not active for the current session. Scalable SQL also returns this status code if a statement attempts to define a savepoint that is already active.

### **893: The cursor reference is invalid.**

A statement referred to a cursor name that was not defined in the current scope. Scalable SQL also returns this status code if a statement attempts to declare a cursor that is already defined in the current scope.

### **894: The named cursor is already open.**

An OPEN CURSOR statement was issued for a cursor that is already open.

### **895: The named cursor is not open.**

A statement requiring an open cursor was not preceded by an OPEN CURSOR statement for the named cursor. Statements requiring open cursors are FETCH, DELETE, UPDATE, and CLOSE.

### **896: The parameter syntax is invalid.**

The parameter syntax is invalid for the stored procedure or statement type. This status code can also indicate generally invalid parameter syntax. Refer to the *SQL Language Reference* for information about SQL statement syntax.

### **897: The calling syntax is invalid.**

A CALL or EXEC statement was encountered where the other was expected. This status code can occur when creating or recalling a procedure. Ensure that you have specified the parameters for the CALL statement correctly. Refer to the CALL statement syntax in the *SQL Language Reference* for more information. This status code can also indicate a generally invalid procedure call syntax.

### **898: You attempted to drop a dependent element.**

The object of a DROP statement is currently defined as a dependency for a trigger.

### **899: The operation requires a scrollable cursor.**

While using the cursor name of a non-scrollable cursor, the application attempted an operation that is permitted only on a scrollable cursor. Specify the SCROLL keyword in the cursor declaration.

### **900: The record number is invalid.**

The specified record number must belong to the current set of records that the chunk operation obtained.

### **901: The correlation name is invalid.**

The correlation name is invalid; this is due to any of the following causes:

- Using an old correlation name with an INSERT trigger.
- Using a new correlation name with a DELETE trigger.
- A correlated column name in the trigger body that did not resolve to one of the trigger's correlation names.

### **902: The stored procedure or trigger body contains a disallowed statement.**

You included a statement in a stored procedure or trigger body that is not allowed in that type of statement. For a complete listing of statements that you cannot include in stored procedures or triggers, refer to the CREATE PROCEDURE and CREATE TRIGGER syntax in the *SQL Language Reference*.

This status code can occur when you use a SELECT statement in a stored procedure and it is not the only statement contained in the procedure. See the CREATE PROCEDURE syntax in the *SQL Language Reference* for more information.

### **903: The trigger body contains an invalid subject table reference.**

The trigger body either included a statement that modified the subject table or called a procedure that modified the subject table. Check the trigger body and all invoked procedures to ensure that no statement attempts to modify the subject table.

### **904: The output parameter must be a variable or parameter name.**

If you define a procedure parameter as OUT or INOUT, then calls made to that procedure must provide a variable or parameter name in that position.

### **905: You attempted to assign a value to a constant parameter.**

A procedure parameter that was passed in as a constant was used later in the procedure as the target of an assignment or FETCH operation.

### **906: You cannot re-execute the statement.**

Scalable SQL returns this status code both inside and outside of procedures, as follows:

- Inside a procedure, it occurs any time a data definition (DDL) statement (such as CREATE TABLE or ALTER TABLE) is encountered within a control loop (such as WHILE). In this case, Scalable SQL does not execute the statement.
- Outside a procedure, it occurs if you use XQLExec to re-execute a DDL statement.

### **907: The IF statement contains invalid syntax.**

The CREATE PROCEDURE statement contains an IF statement with invalid syntax. For information about the syntax of IF statements, refer to the *SQL Language Reference*.

### **908: The WHILE or LOOP statement contains invalid syntax.**

The CREATE PROCEDURE statement contains a WHILE or LOOP statement with invalid syntax. For information about the syntax of WHILE and LOOP statements, refer to the *SQL Language Reference*.

### **909: The FETCH statement contains invalid syntax.**

A FETCH statement contains invalid syntax. For information about the syntax of FETCH statements, refer to the *SQL Language Reference*.

### **910: The exception condition is undefined.**

A stored procedure or trigger raised an exception condition that does not correspond to a specific Scalable SQL status code. Ensure that you specified a valid status code.

### **911: The triggered action failed to complete.**

The trigger failed. Scalable SQL returns this status code when a triggered action signals a SQLSTATE value of 09000, which is the generic value for trigger failure.

### **913: The statement contains an invalid declaration.**

Scalable SQL returns this general syntax error when a DECLARE statement contains invalid syntax. For more information about the syntax of the DECLARE statement, refer to the *SQL Language Reference*.

### **914: The same column may appear only once in a primary or foreign key definition.**

Each column can appear at most once in either a primary or foreign key definition, or at most once in both. Remove the duplicate reference to the column from the definition and re-submit your statement to Scalable SQL.

### **915: The column is not in the view or is not a variable length data type.**

The column from which chunks are to be fetched must be in the view and must be of a variable length data type.

### **916: Stored statement contains parameter without assigned data type or default value.**

This status is returned for option 4 of the xDDProc API when the stored statement (which are only valid for Scalable SQL 3.x or Scalable SQL 4.x when logged in as 3.x Compatibility mode) was generated without either a data type or default value assigned to a parameter. Option 4 of xDDProc returns stored statement or stored procedure parameter information, which is indeterminate if the parameter is assigned neither a data type nor a default value.

### **917: Cannot create system file in a pre-Btrieve 6.0 format.**

You receive this status code when the MicroKernel is configured to generate Btrieve 5.x format files and an operation is performed that causes Scalable SQL 4 to generate a system file such as a DDF file or DBNAMES.CFG.

To solve this condition, reconfigure the MicroKernel to generate files in a Btrieve 6.x or later format as described in the following procedure.



#### **To solve Status 917:**

1. Click **Start**, Point to **Programs** and select **Pervasive SQL 7**.
2. Select **Setup** (Win32 or Win16).
3. If you are performing this operation from a client, click **Connect** and enter the name of the server on which Pervasive.SQL is installed.

4. Select **MicroKernel Database Engine** from the list of components.
5. Select **File Settings** from the list of categories.
6. Select **Create File Version** from the list of settings.
7. Change the setting to 6.x or greater.
8. Click **Save** or **Exit** to retain the new value.
9. Shut down and restart the MicroKernel (refer to *Getting Started With Pervasive.SQL*)

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## 1000 to 1999

# MicroKernel Status Codes for Windows, OS/2, and DOS Workstations

The workstation MicroKernel engine returns the following status codes in Windows, OS/2, and DOS environments.

### **1001: The lock parameter is out of range.**

This status code is obsolete in MicroKernel versions 7.0 and later.

Version 6.x MicroKernels return this status code when the value specified for the Number of Locks configuration option is out of range. Pre-v6.0 MicroKernels return this status code when the value specified for the Multiple Locks configuration option is out of range.

### **1002: A memory allocation error occurred.**

Make sure that the workstation has enough memory to load all the programs it requires. For workstation versions of the MicroKernel, or client requesters, the insufficient memory may apply to conventional memory, expanded memory, or both.

### **1003: An invalid memory size parameter was specified.**

Ensure that the value for the Cache Allocation configuration option is large enough to accommodate the required cache size. See the *Pervasive.SQL User's Guide* for more information.

### **1004: A page size error occurred.**

This status code is obsolete in MicroKernel versions 7.0 and later.

The value of the Page Size configuration option must be an even multiple of 512, and it must be between 512 and 4096, inclusive.

### **1005: The pre-image or lock file drive parameter is invalid.**

You must specify a valid drive letter for the Pre-Image/Lock File Drive configuration option.



**Note:** Pre-image files are used only for pre-v6.0 files. Lock files are used only for v6.0 and later files opened in MEFS mode.

### **1006: The pre-image buffer parameter is out of range.**

The Extended Operation Buffer Size configuration option must be between 0 and 64,000, inclusive. Both the v6.15 and v7.0 MicroKernel return Status 1006 for this reason.

The Pre-Image Buffer Size configuration option must be between 1 and 64, inclusive. The pre-image file is used in pre-v6.0 files and by v6.0 and later MicroKernels that are loaded with the Create Files in Pre-v6.x Format configuration option set to Yes.

### **1007: The files parameter is out of range.**

For pre-v6.0 engines, the Open Files configuration option must be between 1 and 250, inclusive. For v6.0 and later engines, refer to the *Pervasive.SQL User's Guide* for valid range information.

### **1008: The initialization parameter is invalid.**

The specified configuration options contain invalid or unidentifiable values.

### **1009: The transaction filename parameter is invalid.**

This status code is obsolete in MicroKernel versions 7.0 and later.

The filename specified for the Transaction filename configuration option is not valid. Ensure that the transaction filename is correct.

### **1010: An error occurred during the access of the transaction control file.**

This status code is obsolete in MicroKernel versions 7.0 and later.

The MicroKernel is unable to create, open, read, or write to BTRIEVE.TRN or MKDE.TRN. Set the TRNFILE setting in the BTI.INI file to C:\.

### **1011: The compression buffer parameter is out of range.**

The Compression Buffer Size configuration option is out of range. Refer to the *Pervasive.SQL User's Guide* for valid range information.

### **1012: Invalid /n: option.**

This status code is obsolete in MicroKernel versions 6.0 and later.

The Maximum Number Of Files In A Transaction configuration option is invalid. Valid values are 0 to 18; the default is 12.

### **1013: The task list is full.**

In the Windows environment, this status code is returned if the task entry table is full. You can change the value for the Number of Tasks option using the Setup utility. For more information, refer to the *Pervasive.SQL User's Guide*.

In the DOS environment, this status code is returned if the BREQUEST or BREQNT requesters are used without the /t parameter when BTRVID calls are present in the application. Reload the requester with a non-zero value for the /t parameter. For more information on the DOS Requester parameters, refer to *Getting Started With Pervasive.SQL*.

### **1014: Stop warning. The application still has files or transactions active.**

In the OS/2 environment, the MicroKernel returns this status code when an application calls the BTRVSTOP function while files are still open or while a transaction is still active. The application must close all files and end all transactions before calling BTRVSTOP.

### **1015: One of the pointer parameters passed to the MicroKernel is invalid.**

One of the pointer parameters passed into the MicroKernel is invalid. The MicroKernel checks for invalid pointers (and therefore only returns this status code) if you put the following line under the [BTRIEVE] heading in your

initialization file: CHKPARMS=YES. By default, the MicroKernel performs no pointer checking.

### **1016: The MicroKernel is already initialized.**

This status code is obsolete in MicroKernel versions 6.0 and later.

The MicroKernel returns this status code if you attempt to initialize the MicroKernel when it is already initialized. To reinitialize the MicroKernel, close all files, end/abort all transactions, and issue Btrieve operation 25, using the BTRV API before calling the initialization function.

### **1017: The Btrieve Requester is unable to find the resource file WBTRVRES.DLL.**

WBTRCALL.DLL returns this status code when it cannot find the resource file WBTRVRES.DLL. Place the WBTRVRES.DLL file in the same directory as the WBTRCALL.DLL file.

### **1018: The application attempted to call the MicroKernel from a Btrieve callback function.**

The Windows MicroKernel does not allow a task to call the MicroKernel from a Btrieve callback function. You can only use the callback function with Btrieve for Windows.

### **1019: The MicroKernel cancelled the current Btrieve operation at the request of the application's Btrieve callback function.**

The application's MicroKernel callback function returned a nonzero value, indicating that the application wants to terminate the current operation immediately. When the MicroKernel receives such a cancellation request, it attempts to terminate the currently executing operation and ceases to call the callback function for the duration of that operation. The MicroKernel may be unable to cancel the operation. However, if successful in doing so, the MicroKernel returns this status code.

### **1020: Btrieve Requester Interface communications error.**

The MicroKernel loader and requester Interface returns this status code when it cannot send a message to the MicroKernel. This occurs when Windows is shutting down or when you terminate the MicroKernel using Ctrl+Alt+Delete. This is an informational status code only. No action is required. Your application continues the shutdown process. You also receive this status code when running an application that prevents Windows from processing messages.

### **1021: The MicroKernel failed to initialize.**

The MicroKernel could not complete its initialization tasks. Check the MicroKernel's console or error log for a message that specifies the problem that prevented the MicroKernel from initializing. The NetWare MicroKernel displays the message on the server's system console, and it also writes the message to the Pervasive Event Log (PVSU.LOG), which is located at SYS:SYSTEM. The Win32 workstation MicroKernel displays the message in the console message window and writes the message to the Pervasive Event Log (PVSU.LOG), which is located in the C:\WINDOWS directory. The Windows NT server MicroKernel does not display a message, but writes the message in the Pervasive Event Log (PVSU.LOG), which is located in the C:\WINNT directory.

Correct the problem, using the Setup utility if necessary, then retry the operation. If you are using a NetWare MicroKernel, you must first correct the problem, then unload and reload the MicroKernel before you can retry the operation.

### **1022: The MicroKernel is shutting down.**

The operation cannot be completed because the MicroKernel is shutting down. To correct this problem, allow the engine to completely shut down, and then restart.

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## 2000 to 2099

### Btrieve Requester Status Codes

This section lists the status codes that the Btrieve Requesters generate.

#### **2000: Internal error.**

The Btrieve requester encountered an internal error. Check the Pervasive Event Log (PVSW.LOG) for more information.

#### **2001: The memory allocation is insufficient.**

In an OS/2 environment, the Requester cannot allocate enough memory for the parameters specified with the BRQPARMS environment variable. In a DOS environment, reduce the value specified for the /D configuration option.

#### **2002: The option is invalid or out of range.**

In an OS/2 environment, either one of the options specified with the BRQPARMS environment variable is invalid (such as /P instead of /D) or the value specified for a parameter is out of range. Check the SET BRQPARMS statement to make sure it is correct.

#### **2003: The Requester does not allow local access to the specified file.**

The application attempted to access a file stored on a local drive. The configuration of the MicroKernel installed at the client machine does not allow access to local files.

This status code has been supplemented in Btrieve 7.0. For a list of the new status codes, see [“Status Code 2003 replacement”](#).

#### **2004: SPX is not installed.**

Install the NetWare SPX v2.0 or later communications software for OS/2.

#### **2005: An incorrect version of SPX is installed.**

Install the NetWare SPX v2.0 or later communications software for OS/2.

#### **2006: There is no available SPX connection.**

SPX has already established the maximum number of sessions it can handle. To increase the maximum, edit the NET.CFG file. Refer to your NetWare documentation for more information about NET.CFG.

#### **2007: A pointer parameter is invalid.**

One of the pointer parameters passed to the MicroKernel is invalid. Check the program to ensure that the pointer parameters are correct.

#### **2008: Router cannot find engine.**

The MicroKernel Router cannot communicate with the 6.15 engine. This status code is only used with Scalable SQL v4.0.100.

## **2009: Cannot load MicroKernel Router component.**

The Btrieve requester cannot load the MicroKernel Router. This can occur if the DLL fails to load or fails to obtain the necessary DLL entry point.

## **2010: Thunk not possible - the Win32 MicroKernel router could not be found or is incompatible.**

You receive this status code when the system is configured for thunking and the path is not accessible. This can occur if a necessary system file cannot be found (for example, kernel32.dll). It can also occur if the required system calls (for example, Win32s library) are not available.

## **2011: Btrieve requester resource DLL not loaded.**

The resource DLL controls options that you can set with the Setup utility. You receive this status code when the resource DLL is either missing or incompatible with the current version of the requester. If this happens, the MicroKernel reverts to its default settings and continues to run.

## **2012: The Btrieve requester encountered an operating system error.**

Check the Pervasive Event Log (PVSU.LOG) for more information.

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](http://www.pervasive.com) All rights reserved.

## 2100 to 2199

### Scalable SQL Interface Component Status Codes

The Scalable SQL Interface component returns the following status codes to an application if the Interface component encounters an error while the application is running at a DOS, Windows, Windows NT, Windows 95, or OS/2 machine. The application that is using Scalable SQL should intercept these errors and instruct the user to take the appropriate action.

#### **2101: The data message buffer at the workstation is too small.**

The data message buffer at the client machine is too small to contain the data to be transferred from the client machine to the server. If you are working in the DOS environment, reload the Scalable SQL Requester for DOS at your client machine, specifying a larger value for the Data Message Length (/D) configuration option. If you are working in the Windows or OS/2 environment, increase the memory at your client machine; in these environments, the communications buffer size is dynamically allocated as needed, if enough memory is available.

#### **2102: The redirection buffer is too small.**

The number of redirected devices for the client machine has increased since the Requester was loaded. If you are working in the DOS environment, reload the Scalable SQL Requester at your client machine, specifying a larger value for the Mapped Drives (/R) configuration option. If you are working in the Windows or OS/2 environment, increase the memory at your client machine; in these environments, the communications buffer size is dynamically allocated as needed if enough memory is available.

#### **2103: Scalable SQL is not active on the requested server.**

Scalable SQL returns this status code for one of the following reasons:

- Scalable SQL is not loaded on the server to which the application sent the Scalable SQL request.
- Scalable SQL at the server was unloaded and then reloaded, but the Scalable SQL Requester at a client machine was not reloaded.
- The Requester cannot find all the required components at the client machine.
- In a Windows, Windows 95, or Windows NT application, you are attempting to log in to a database, and you have not specified a full pathname.
- You do not have the NT Server name configured correctly when connecting to an NT server via database names, and the MS Client is installed on the workstation.

To resolve this, use the Configuration utility to add the NT Server name to the following setting: DBNames Interface/Access Control/Scalable SQL Servers list. For more information on the Scalable SQL Servers setting, refer to the *Pervasive.SQL User's Guide*.

#### **2104: The maximum number of logins has been exceeded.**

A Scalable SQL application has exceeded the number of concurrent logins. Either log out of a database or increase the Requester's Number of Concurrent Sessions configuration option. For more information about changing this setting, refer to the *Pervasive.SQL User's Guide*.

#### **2105: The server buffer is too small.**

The number of Scalable SQL servers to which you are attached exceeds the number you set with the Number of Servers configuration option. If you are working in the DOS environment, reload the Scalable SQL Requester at your

client machine, specifying a larger value for the Number of Servers option. If you are working in the Windows, Windows 95, Windows NT, or OS/2 environment, increase the memory at your client machine; in these environments, the communications buffer size is dynamically allocated as needed, if enough memory is available.

### **2106: The Scalable SQL data message buffer is too small.**

The application attempted to send data to a Scalable SQL server engine, but the data exceeded the size of the Scalable SQL data buffer length. Perform one of the following:

- Use the Setup utility to increase the Communications Buffer Size in the Scalable SQL Requester component to 32767, which is the default maximum.
- Use the Setup utility to increase the size of the Communications Buffer Size option on the client. The Communications Buffer Size option is contained in the Scalable SQL Requester component section of the Setup utility.

For more information about the setup utility, see the *Pervasive.SQL User's Guide*.

The value you specify for these options should be equal to or greater than the largest value set for the Communications Buffer Size option for a client machine.

### **2107: The application has activated the maximum number of views.**

This status code is obsolete in Scalable SQL versions 4.0 and later.

The application attempted to activate more views than were specified by the Number of Active Views configuration option. If you are working in the DOS environment, reload the Scalable SQL Requester at your client machine, specifying a larger value for the /V option. If you are working in the Windows or OS/2 environment, increase the memory at your client machine; in these environments, the communications buffer length is dynamically allocated as needed if enough memory is available.

### **2108: The device is not assigned to a server.**

You have attempted to access a server that is not available. Be sure that the client machine has drive letter mappings that correspond to every drive letter specified in the database definition.

### **2109: You have attempted to attach to too many servers.**

This status code is obsolete in Scalable SQL versions 4.0 and later.

The maximum number of file servers to which a client machine can attach is 8.

### **2111: The session was reset from the server console.**

A session was reset from the Scalable SQL server console while the server was actively processing a request for that session. As soon as the server detects the reset command, all resources assigned to that session are released. All active processing of requests is halted, and any requests that are within a transaction and waiting to be processed are aborted. You must log in to the server-based database and establish a new Scalable SQL session.

### **2112: The Scalable SQL Requester is not loaded.**

This status code is obsolete in Scalable SQL versions 4.0 and later.

A Windows application attempted to access server-based Scalable SQL, but the Scalable SQL Communications Requester for Windows did not load. Windows could not locate the component. You receive this status code only when the application is using the Windows Requester interface.

## **2113: The configuration file is missing or contains an invalid drive mapping specification.**

Either the Scalable SQL Requester for DOS cannot find the configuration file, or there is an invalid drive letter mapping in the configuration file. Refer to the *Pervasive.SQL User's Guide*.

## **2114: The specified server and volume for the drive alias is invalid.**

In the Scalable SQL Requester for DOS configuration file, you specified a drive letter mapping incorrectly. Refer to the *Pervasive.SQL User's Guide*.

## **2115: The Interface component could not allocate or access memory used to communicate with the local Scalable SQL engine.**

An internal error occurred. Try unloading and reloading the Interface component and ensure that you have enough available memory to run the application.

## **2116: There is insufficient server memory to process the request.**

The server no longer has enough dynamic system memory to process any Scalable SQL requests, possibly including a logout request. This situation arises only if the demands on the server are too great for the available amount of memory on the server. The overload could be incurred by too many sessions, concurrent requests, or active cursors, or by other memory demands.

To solve this problem, either use the Monitor utility to determine the source of the memory demands or increase the server's total RAM.

## **2117: The share count is invalid.**

When you use xShareSessionID to change the current share count, the resulting share count must be between 0 and 32,767 (both inclusive). You receive this status code only when the application is using the Windows Requester Interface. Refer to the *Scalable SQL Programmer's Guide* for information about sharing sessions and the xShareSessionID function.

## **2118: This session is not shareable.**

A task attempted to switch to an active login session that it did not own and the session is not specified as shareable. Tasks can specify a session as shareable by using the xShareSessionID function. You receive this status code only when the application is using the Windows Requester Interface. Refer to the *Scalable SQL Programmer's Guide* for information about sharing sessions and the xShareSessionID function.

## **2119: The API is invalid for a sharer.**

A task attempted to do one of the following:

- Use the xShareSessionID function on an invalid login session or a login session that the task did not own.
- Use XQLLogout, xLogout, or xReset to log out of or reset a login session that it shared but did not own.

You receive this status code only when the application is using the Windows Requester interface. Refer to the *Scalable SQL Programmer's Guide* for information about sharing sessions and the xShareSessionID function.

## **2120: The task table is full.**

A Scalable SQL application has exceeded the maximum number of concurrent tasks. Either terminate a task with the XQLLogout function or increase the setting for the Number of Sessions configuration option. To change the setting for the Number of Sessions option, use the Setup utility. For more information, refer to the *Pervasive.SQL User's Guide*.

### **2121: The connection between the Interface DLL and the engine has been lost.**

The Interface DLL cannot communicate with the Scalable SQL engine. Ensure that the engine was not unloaded while an application was using it.

### **2122: No thread is available to process the request.**

The Interface DLL could not gain a thread in the Scalable SQL engine to process a request. All the Worker Threads in the Scalable SQL engine were busy processing requests, and no additional threads could be allocated because the maximum number of threads has been reached. Retry the Scalable SQL operation that received this status code. If the condition persists, terminate one or more Scalable SQL applications that is running on the server, and then retry the operation.

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## 2200 to 2299

### Scalable SQL Miscellaneous Status Codes

This section describes the status codes that Scalable SQL applications can receive in a platform-specific environment, such as the server-based environment or the Windows environment.

#### **2200: The maximum number of sessions has been exceeded.**

Use the Setup utility to increase the maximum number of sessions allowed on the machine. For information about how to do this, refer to the *Pervasive.SQL User's Guide*. The number of sessions you can specify is limited, in part, by the amount of memory available.

#### **2201: The Scalable SQL for Windows DLL cannot yield to allow the other task to complete.**

Your task made a Scalable SQL API call while the Scalable SQL for Windows DLL was processing an API request for another task. The Scalable SQL local engine does not process more than one API request per task at a time. This status code occurs only if your task has not registered a callback function with Scalable SQL.

Scalable SQL cannot call your task's callback to yield back to the task for which it is processing the API request. Either register a callback for your task or disable callbacks for all tasks by setting the Allow Callbacks option to NO in the [Scalable SQL] section of the initialization file. However, if you disable callbacks, none of the tasks can yield time to allow other tasks to run.

#### **2202: The task removing a callback did not specify a correct previous callback address.**

You must unregister callbacks in the reverse order in which they were registered. Ensure that you pass the correct address of the previous callback. When you register a callback, XQLCallback returns the address of the previous callback in the fPrevCallback parameter. By passing the address of the previous callback when you remove a callback, you reregister the previous callback.

#### **2203: At the user's request, Scalable SQL canceled the operation.**

The user of the application has aborted the Scalable SQL API request.

#### **2204: A re-entrant call was detected.**

Either a Scalable SQL for Windows application attempted to make a Scalable SQL API call while in a callback, or an application made a request using a session ID that is already processing a request.

The Scalable SQL for Windows local engine does not currently process more than one API request at a time. Because the engine is already processing an API request for the task, the callback cannot make an additional Scalable SQL API request.

The Scalable SQL engines do not process more than one API request for a given session at a time. Because the engine is already processing an API request for a given session, the task cannot make an additional Scalable SQL API request using that session ID.

#### **2205: Invalid data was passed to Scalable SQL.**

The API call contained incorrect or incomplete data. Ensure that all parameters and options are set correctly before calling the API.

## **2206: A semaphore error occurred.**

An internal semaphore error occurred in the Scalable SQL engine or Interface DLL. The Scalable SQL engine returns this status code when it tries to perform an operation using an incompatible version of the DLLs. Shut down the engine and make sure that you are using the most recent version of the DLLs.

## **2208: The Scalable SQL engine is shutting down.**

The operation cannot be completed because the Scalable SQL engine is in the process of shutting down. To correct this problem, allow the engine to completely shut down, and then restart.

## 2300 to 2399

### Database Names Functions Status Codes

This section lists the status codes you can receive when using the database names functions.

#### **2300: No more database names are defined.**

No more database names are defined. If the buffer is large enough to hold multiple database names, you might receive this status code and still have one or more database names returned. The application should check `iBufLen` to determine the number of names returned.

#### **2303: The database name must be unique; the specified name already exists.**

While creating a database, you specified a database name that already exists in the `DBNAMES.CFG` file. Specify a different, unique name for the database, or remove the existing database name first.

#### **2305: The specified path for the dictionary or data file locations is invalid.**

The path you specified for the bound, named database is invalid or the paths you specified for the data file locations are invalid. Ensure that the entered paths are either UNC paths or local paths (a local path for a NetWare server would be *volume:path*) and then retry the operation. Paths that contain mapped drive letters are not allowed.

#### **2306: The specified path for the dictionary or data file locations is invalid.**

The path you specified for the bound, named database is invalid or the paths you specified for the data file locations are invalid. Ensure that the specified paths are either UNC or local paths (a local path for a NetWare server would be *volume:path*) and then retry the operation. Paths that contain mapped drive letters are not allowed.

#### **2307: Scalable SQL cannot open the DBNAMES.CFG file, or a bindery error occurred.**

If you are trying to access the workstation `DBNAMES.CFG` file to obtain a list of database names defined for the workstation engine, ensure that the `DBNAMES.CFG` file is in your Windows directory or in the directory specified by the Database Names Directory configuration option.

If you are trying to obtain a list of database names available on Scalable SQL for NetWare servers (or Scalable SQL for Windows NT servers that have SAP Agent enabled), the `DBNames` DLL received an error while scanning the bindery, which may mean that the bindery is corrupt. Use the appropriate NetWare utility for your NetWare environment to repair the bindery.

#### **2309: The database is in use.**

You cannot modify the definition of a named database if another user is modifying the definition, or if a user is connected to it. You also cannot connect to a named database if someone is modifying its definition.

#### **2310: The database is not bound.**

You cannot unbind a database that is not bound. You must bind the database before you can unbind it. For more information about bound databases, refer to the *Pervasive.SQL Programmer's Guide*.

#### **2311: The database is already bound.**

You cannot bind a database that is already bound. You must unbind the database before you can bind it. For more

information about bound databases, refer to the *Pervasive.SQL Programmer's Guide*.

### **2312: The bound database cannot share table data files.**

You cannot bind a data file referenced by a table in a bound, named database to another named database, or to another table in the same named database. For more information about bound databases, refer to the *Pervasive.SQL Programmer's Guide*.

### **2313: The bound database cannot share data dictionary files.**

The data dictionary files for a bound, named database cannot be referenced by another named database. For more information about bound databases, refer to the *Pervasive.SQL Programmer's Guide*.

### **2314: Scalable SQL cannot create the DBNAMES.CFG file.**

Scalable SQL cannot create the DBNAMES.CFG file. Verify that the file does not already exist, then retry the operation.

### **2315: Unable to bind or unbind the data dictionary files.**

Scalable SQL encountered a MicroKernel error while attempting to write binding information to one of the data dictionary files for the database. Retry the operation. If you continue to receive this status code, ensure that the disk is not full and that the data dictionary files are v6.x or higher.

### **2316: Scalable SQL cannot create DDF files for the bound database.**

Scalable SQL cannot create the database files for the bound database. Ensure that the data dictionary files do not already exist at the specified location. Remove any existing data dictionary files before creating the bound, named database. For more information about bound databases, refer to the *Pervasive.SQL Programmer's Guide*.

### **2317: Unable to open DDF files for the database to bind or unbind the database.**

Scalable SQL cannot open the data dictionary files to bind or unbind the database. This could be caused by a variety of unusual conditions, but *not* because the data dictionary files are not in the specified location. (If the data dictionary files are not in the specified location, you receive Status Code 2323.) Try to login to the same database through an application (such as SQLScope). If this fails, then note and correct the error indicated.

### **2318: Not allowed to change the name for the named database.**

You cannot change the name for a bound, named database. You also cannot change the name for a named database if any of the tables in the named database contain RI constraints or triggers. To change the name, you must make sure that the database is not bound and remove all triggers and RI constraints from the database. You can then change the name, add the triggers and RI constraints, and bind the database.

### **2319: Unable to bind or unbind the data file.**

Scalable SQL encountered a MicroKernel error while attempting to write binding information to a data file for a table in the database. Retry the operation. If you continue to receive the status code, ensure that the disk is not full and that the data file is v6.x or higher.

### **2320: Unable to bind or unbind a secured database.**

Binding and unbinding is not supported on a secured database. In order to bind or unbind a secured database, you

must first uninstall security from the database. For more information about bound databases, refer to the *Pervasive.SQL Programmer's Guide*.

### **2321: Data dictionary files do not support binding information.**

Scalable SQL attempted to bind the data dictionary files to a named database, but the data dictionary files were not the correct version. In order to support binding information, the files must be v6.x or higher. For more information about bound databases, refer to the *Pervasive.SQL Programmer's Guide*.

### **2322: Data file does not support binding information.**

Scalable SQL attempted to bind a data file for a table in the named database, but the data file was not the correct version. In order to support binding information, the data file must be v6.x or higher. For more information about bound databases, refer to the *Pervasive.SQL Programmer's Guide*.

### **2323: Data dictionary files cannot be found.**

Scalable SQL was unable to locate the data dictionary files in the specified directory. Ensure that the specified directory exists and contains data dictionary files, and then retry the operation.

### **2324: Data dictionary files are not bound.**

The data dictionary files for the specified named database are not bound, but should be. This normally indicates a situation in which the data dictionary files have been restored from a backup prior to the database being bound. Make sure that you restore your data dictionary from a backup which is consistent with the bound state of the databases. For more information about bound databases, refer to the *Pervasive.SQL Programmer's Guide*.

### **2325: Data dictionary files are already bound.**

The data dictionary files for the specified named database are already bound to another database. However, the other named database is not defined to the Scalable SQL engine that is trying to bind or unbind the database. Make sure that you only reference the data dictionary files for a bound database from one Scalable SQL engine. For more information about bound databases, refer to the *Pervasive.SQL Programmer's Guide*.

### **2326: Data dictionary files are bound but do not need to be bound.**

This status code is returned when an unbound named database is checked. This does not indicate an invalid, or corrupted, named database, but does indicate an inconsistency between the named database definition and the actual state of the database.

This status code often indicates a situation in which the data dictionary files have been restored from a backup that was made when the database was bound. Either restore the data dictionary from a backup that is consistent with the bound state of the database, or bind and then unbind the database. The unbinding of the database removes the binding information from the data dictionary files.

### **2327: Data dictionary files are bound to an unknown database.**

This status code typically occurs when securing, or unsecuring, a bound database. The data dictionary files contain incorrect binding information. The binding information indicates an unknown named database. This could indicate that an incorrect backup of the data dictionary files was restored for this named database.

### **2328: Data file for a table cannot be found.**

Scalable SQL was unable to locate a data file for a table in the database. Ensure that the data files for all of the tables in the data file exist in one of the data file paths for the named database and retry the operation.

### **2329: Data file for a table is not bound.**

The data file for a table in the database is not bound, but should be. This normally indicates a situation in which the data files for the database have been restored from a backup prior to the database being bound. Make sure that you restore your data from a backup that is consistent with the bound state of the database.

### **2330: Data file for a table is bound, but does not need to be bound.**

This status code is returned when an unbound, named database is checked. This does not indicate an invalid, or corrupted, named database but does indicate an inconsistency between the named database definition and the actual state of the database. This status code often indicates a situation in which the data files for the database was bound. Either restore your data files from a backup that is consistent with the bound state of the database, or bind and then unbind the database.

The unbinding of the database removes the binding information from the data files if the binding information is no longer needed.

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## 2900 to 2999

### Inscribe Status Codes

This section lists the status codes you can receive when using the Inscribe engine.

#### **2900: The Inscribe Interpreter could not be initialized correctly.**

The EPInit API made a call to initialize the Inscribe Interpreter, but the initialization call was unsuccessful. Verify that the proper Inscribe Interpreter is installed on your system (WINSIA6.DLL for Windows 3.1, WINSIA32.DLL for Windows 95 and Windows NT, NWINSIA.NLM for NetWare). Also verify that you have sufficient available memory for Inscribe to run.

#### **2901: The Inscribe Interpreter could not be terminated correctly.**

The EPDeInit API made a call to terminate the Inscribe Interpreter, but the call was unsuccessful. You can receive this status code when the Inscribe engine terminates the Inscribe Interpreter. This error should not affect active applications.

#### **2902: An attempt to create an internal thread within the Inscribe Interpreter failed.**

When EPCall executes procedures, internal resources (or threads) are allocated for those procedures. This status code indicates that an attempt to perform this allocation for a procedure failed. Verify that you have sufficient available memory for Inscribe to run.

#### **2903: An attempt to kill an internal thread within the Inscribe Interpreter failed.**

When EPCall executes procedures, internal resources (or threads) are allocated for those procedures. Once the procedures complete execution, those resources are freed, or killed, so other procedures can use them. This status code indicates that an attempt to perform this deallocation for a procedure failed. This error should not affect active applications; the effect of this status code is that a small amount of allocated memory is not freed.

#### **2904: An attempt to load a module containing an Inscribe procedure failed.**

You used the EPCall API to execute a procedure, but when EPCall attempted to load the module containing the procedure, the load failed. Because you use the Inscribe Developer Kit to install the modules containing the procedures into a database, verify that you installed the modules correctly and have sufficient available memory to allow Inscribe to load new modules. This status code can also be returned if the procedure being loaded is unable to find a DLL or NLM referenced by the procedure.

#### **2905: An attempt to unload a module containing an Inscribe procedure failed.**

You used the EPCall API to execute a procedure, but when EPCall attempted to unload a previously loaded module to make room for the module containing this procedure, the unload failed. Verify that you have enough memory for Inscribe to run; some memory is temporarily allocated to perform the unload operation.

#### **2906: A request was made to execute an Inscribe procedure, but no procedure with the specified name was found.**

You used the EPCall API to execute a procedure, but when EPCall attempted to locate the procedure in the database specified in the EPLogin call, the search failed.

If you receive this status code, check the following:

- Because different procedures are associated with different databases, verify that the database name specified in EPLogin is correct.
- Verify that the procedure name passed as a parameter to EPCall is spelled correctly.
- Because you use the Inscribe Developer Kit to install the modules containing the procedures into a database, verify that you installed the modules correctly.

### **2907: The execution of an Inscribe procedure failed.**

You used the EPCall API to execute a procedure, but when EPCall executed it, the procedure did not terminate normally. Use the Inscribe Developer Kit to locate and correct any problems in the Inscribe procedure that could result in errors during execution. The Visual Basic-compatible language Inscribe uses supports the "On Error" statement, which you can also use to handle run-time errors such as an attempt to open a non-existent file.

### **2908: The type or number of the arguments passed to an Inscribe procedure does not match the type or number of the parameters defined for that procedure.**

You used the EPCall API to execute a procedure, but when EPCall compared the number and types of the arguments passed to EPCall with the number and types of the parameters defined for the procedure, it found one or more conflicts. Correct any discrepancies between the number and types of the parameters defined for the procedure and the arguments that you passed to EPCall.

### **2909: An error was encountered building the argument vector to be passed to an Inscribe procedure.**

You used the EPCall API to execute a procedure, but when EPCall processed the arguments passed in to EPCall, Inscribe attempted to convert the input information into an argument vector usable by the Inscribe procedure and then detected an error. Verify that the input argument information is in the correct format. For example, you must specify the DATE, TIME, and TIMESTAMP data types in a particular format when you pass them to EPCall. Also verify that you have sufficient available memory for Inscribe to run; Inscribe uses additional memory for some parameter types.

### **2910: An error was encountered extracting values from the argument vector returned by an Inscribe procedure.**

You used the EPCall API to execute a procedure, but when EPCall processed the argument vector returned to EPCall after the procedure completed, Inscribe attempted to process an output argument and detected an error. When a string argument is returned from an Inscribe procedure, a buffer is provided by the function that called EPCall to store the string. Inscribe returns this status code when the buffer is too small to hold the string. Check all output string arguments and verify that sufficient space is provided to hold all output values.

### **2911: The Inscribe Interpreter did not load successfully.**

Under Windows 3.1, Windows 95, Windows NT, or NetWare, the Inscribe Interface loads the Inscribe Interpreter to execute procedures. The Inscribe Interface failed to load the Inscribe Interpreter. Verify that you installed the proper Inscribe Interpreter module on your system and that you have sufficient available memory to run Inscribe.

Scalable SQL also returns this status code if it is unable to load Inscribe. Verify that you installed the proper Inscribe components on your system and that Scalable SQL's Enable External Procedures configuration option is set to On. For more information about configuring Scalable SQL, refer to the *Pervasive.SQL User's Guide*.

### **2912: The Inscribe Interface encountered an error setting the internal**

## **instance flags that specify execution characteristics for the Inscribe Interpreter.**

The Inscribe Interface sets certain instance flags, such as flags for multi-threading, to allow the Inscribe Interpreter to run correctly in a particular environment. Inscribe returns this status code when an attempt to set these instance flags fails. If you receive this status code, the Inscribe Interpreter (WINSIA32.DLL, WINSIA16.DLL, or NWINSIA.NLM) may be corrupt. Try reinstalling it.

### **2913: The amount of available memory is less than that needed for Inscribe to run.**

The operating system has insufficient memory for Inscribe to run. Terminate other applications to make additional memory available.

### **2915: The database specified as an argument to the EPLogin API is not a valid database.**

The string passed to EPLogin representing the database may be invalid.

In Inscribe 1.0, if the string is valid, Inscribe was unable to access and read the database's EPMap.TXT file to get information about the Inscribe procedures defined for that database. Verify that the database name is correct. If it is correct, verify that the Inscribe Developer Kit was used to create an EPMap.DDF file for the database that contains information about the Inscribe procedures associated with that database.

### **2916: The Inscribe session ID is invalid.**

The Inscribe session ID passed as an input argument to EPLogout or EPCall is not a valid Inscribe session ID. Verify that the Inscribe session ID that was returned as an output value from EPLogin is the same one that is being used as the input session ID to EPLogout and EPCall.

### **2917: The credential specified by the caller when Inscribe was initialized does not match the credential specified when Inscribe was terminated.**

The Inscribe credential passed as an input argument to EPDeInit is not a valid Inscribe credential. Verify that the Inscribe credential that was specified as an input argument to EPInit is the same one that is used as the input argument to EPDeInit.

### **2918: An error was encountered deregistering a callback function in EPCallback, or an error was encountered in EPCall specifying a callback function for use by the Inscribe Interpreter.**

If EPCallback returns this status code, then the callback function specified in the fCallback parameter does not match the currently registered callback. If EPCall returns this status code, then there was an error establishing this callback function as the one for the Inscribe Interpreter to use. For EPCallback, verify that the fCallback parameter contains the correct value. (Callbacks must be deregistered in the reverse order in which they were registered.) For EPCall, verify that the callback function specified in EPCallback is a valid pointer to a callback function.

### **2919: While logged in with the same Inscribe login ID, an attempt was made to execute an Inscribe procedure before a previous Inscribe procedure completed.**

When logged in with the same Inscribe login ID, only one Inscribe procedure can be executed at a time so that

callback functions can be handled correctly. Typically, Inscribe only returns this status code in multi-threaded applications. If you need multiple threads to concurrently execute Inscribe procedures, then have each thread perform its own separate EPLogin call to get a unique login ID. Do not try to share the same login ID among multiple threads.

## **2920: An attempt was made to call an Inscribe Interface API function, but Inscribe has been disabled in the user's environment.**

A call was made to the EPVersion API to get version information for an Inscribe component, but an invalid value was passed in the flags argument. Verify that the value you set in the flags argument is correct.

## **2921: The Inscribe engine encountered an error reading an Inscribe module.**

Inscribe modules are .SBX files, which you produce by compiling an Inscribe script using the SBL Developer Tools and then copy to a database directory. The Inscribe engine returns this status code when the engine is unable to read a module or the contents of a module are invalid.

Verify that no other active applications (such as the 16- or 32-bit SBL Developer Tools) are accessing any Inscribe modules; such access may prevent the Inscribe engine from reading the modules. If no other applications are accessing any Inscribe modules, one or more of the modules may contain invalid data, caused by one of the following:

- Disk errors.
- Presence of modules compiled with Inscribe v1.0. (Any modules created with Inscribe v1.0 are not compatible with Inscribe v2.0 and must be recompiled with Inscribe v2.0.)
- Another application that overwrote some of the modules.

Recreate the Inscribe modules by recompiling them in a SBL Developer Tool and copying them to the database directory.

## **2922: The Inscribe engine encountered a duplicate procedure name when processing an Inscribe module.**

Inscribe modules are .SBX files, which you produce by compiling an Inscribe script using the SBL Developer Tools and then copy to a database directory. The procedure names within these modules must be unique so that Inscribe can uniquely determine which Inscribe procedure to invoke for a particular Scalable SQL external procedure name.

Examine the Inscribe modules to find the duplicate names, change the names to avoid the conflict, then recompile the modules in a SBL Developer Tool and copy them to the database directory.



**Note:** The *main* procedure can appear in multiple modules without causing this status code, because *main* procedures are intended for debugging with the SBL Developer Tools and are ignored by the Inscribe engine. Similarly, private procedures that appear in multiple modules can have the same name, because private procedures are also ignored by the Inscribe engine.

## 3000 to 3099

### MicroKernel Router Status Codes

This section lists the status codes you can receive from the MicroKernel router, which receives requests from the Btrieve requesters and routes them to the correct version of the MicroKernel.

#### **3000: The MicroKernel router encountered a memory allocation error.**

The MicroKernel router could not internally allocate memory. Check the Pervasive Event Log (PVSW.LOG) for more information.

#### **3001: Local access is unavailable to the MicroKernel router.**

A call to the MicroKernel failed. This is most likely the result of an incorrect configuration. For example, if the MicroKernel is unavailable because of configuration settings and the file to be opened is local, then the MicroKernel router returns this status code.

#### **3002: The MicroKernel router resource DLL is unavailable.**

The MicroKernel router returns this status code when its resource DLL is either missing or incompatible with the current version of the router. If this happens, the MicroKernel reverts to its default settings and continues to run.

Check the Pervasive Event Log (PVSW.LOG) for more information.

#### **3003: The MicroKernel router detected an incompatible network component.**

The networking services component is not compatible with this version of the MicroKernel router. The most likely cause is that the Networking services DLL has been replaced by an older version.

Check the Pervasive Event Log (PVSW.LOG) for more information.

#### **3004: The MicroKernel is not responding to the MicroKernel router.**

The MicroKernel is not responding to requests from the MicroKernel router. Verify that the MicroKernel is running.

#### **3005: The MicroKernel router encountered an operating system error.**

The MicroKernel router encountered an unexpected error from the operating system, such as a shared memory failure.

Check the Pervasive Event Log (PVSW.LOG) for more information.

#### **3006: The MicroKernel router detected an invalid session.**

The session information contained in a position block is invalid. This status code occurs for one of the following reasons:

- The application is trying to use a position block for a file that resides on a remote server and for which the connection has been previously terminated via a B\_STOP (25) or B\_RESET (28) operation.
- The application is trying to use a position block for a file that resides on a remote server and that encountered an abnormal network error on a previous operation.

### **3008: Invalid configuration for MicroKernel router.**

This status code results from an invalid configuration. For example, if the Local, Remote, and Thunk options have all been set to *No*, the router has no valid path to take. If this happens, set one of these options to *Yes*, depending on your setup.

### **3009: NETAPI.DLL is not loaded.**

The MicroKernel router could not find NETAPI.DLL (or NETAPI32.DLL in OS/2).

### **3010: NetWare API DLLs are not loaded.**

The MicroKernel router could not find the NetWare API DLLs.

### **3011: Thunk not possible - the Win32 MicroKernel router could not be found or is incompatible.**

This status code is only relevant to the Win16 MicroKernel router. You receive this status code when the Use Thunk option is set to *Yes* and the thunk path is not accessible. Possible causes include:

- The Win32 MicroKernel Router DLL does not exist or is not compatible.
- The Win32s routines failed to initialize.

Check the Pervasive Event Log (PVSU.LOG) for more information.

### **3012: Local engine is not accessible to the MicroKernel router.**

Access to the local engine is not possible because it is not loaded or could not be launched. You can receive this status code if you try to access a local file on a client and you do not have a MicroKernel workstation engine installed or if you try to access a local file on a server and the MicroKernel server engine is not running.

If you have only a server engine installed and your Pervasive Event Log (PVSU.LOG) contains Status Code 3012 warning entries, perform the following steps:



#### **To Turn Off Local Engine support:**

1. Click **Start**, point to **Programs**, then to the **Pervasive folder**.
2. Select **Setup** (Win32 or Win16).
3. Select the **MicroKernel router** component.
4. Select the **Local** setting and change it to Off.
5. Select the **Scalable SQL Requester** component.
6. Select the **Local Usage** setting and change it to Off.
7. Click **Save** and then **Exit**.

### **3013: The remote engine is inaccessible to the MicroKernel router because the**

## **networking component is not loaded.**

Access to the remote engine is not possible because the MicroKernel router could not initialize the networking component. Possible causes include:

- The Remote option is incorrectly set to *No*.
- The required networking component could not be found or is not compatible.

## **3014: The MicroKernel router cannot find an engine.**

The MicroKernel router could not complete the operation because it did not find an engine (local or remote) that could process the operation. Possible causes include:

- The target network operating system is not available.
- The target engine is not available.
- In a Windows 3.x environment, you are using a shared drive name that contains a space. To resolve the condition, remove the space from the shared drive name. For example, if you connect to a server named "D Drive", change the share name to "D\_Drive".

## **3015: The MicroKernel router encountered an initialization error.**

An unexpected error occurred during the initialization of the MicroKernel router. Check the Pervasive Event Log (PVS.W.LOG) for more information.

## **3016: The MicroKernel router encountered an internal error.**

The MicroKernel router encountered an internal error. Check the Pervasive Event Log (PVS.W.LOG) for more information.

## **3017: Data buffer of the local engine is too small.**

You receive this status code when the requester has more data to send the MicroKernel than the MicroKernel's buffers can handle. This status code is only relevant to an engine running locally; it does not apply to client/server environments. Some possible solutions include:

- Specify a higher value for the Communications buffer size parameter in the Setup utility.
- Use smaller buffers in Btrieve operations.

## **3018: The file is already closed.**

You receive this status code when the MicroKernel router is shut down and is not accepting any more requests.

## **3019: The MicroKernel router encountered a semaphore error.**

A semaphore error occurred while trying to establish contact with the local engine. Check the Pervasive Event Log (PVS.W.LOG) for more information.

## **3020: An error occurred while loading the MicroKernel.**

You receive this status code when an error occurs while loading the MicroKernel or when access to the server and client machine's shared memory objects are unavailable.

Check the Pervasive Event Log (PVSU.LOG) for more information.

### **3021: The MicroKernel router received a badly formatted data packet.**

The MicroKernel router rejected the response from the engine because it was badly formatted.

Check the Pervasive Event Log (PVSU.LOG) for more information.

### **3022: The MicroKernel router could not send the request to the remote engine because the specified data buffer length resulted in a data packet object that is too large.**

A request to the MicroKernel router specified a data buffer length that resulted in a data packet size greater than 64 KB, which is the maximum. This can occur even if the data buffer length for the request is smaller than 64 KB because the MicroKernel router adds some packet overhead as it forms the data packet.

Check the Pervasive Event Log (PVSU.LOG) for more information.

### **3023: The task table of the MicroKernel router is full.**

An internal table used by the MicroKernel router to manage tasks is full. This error should only occur in Win16 environments and indicates that a large number of applications are using the Win16 MicroKernel router at the same time.

### **3025 : Cannot load IDS client requester.**

The MicroKernel router could not load the IDS client requester.

The IDS client requester is W3ICRxxx.DLL, where xxx is a revision number (100, 101, etc.). The IDS client requester needs to be located in a directory listed in your PERSVasive\_PATH or PATH environment variable.

### **3026 : Incompatible IDS client requester.**

The MicroKernel router found an incompatible IDS client requester.

The IDS client requester was found, but is a version that cannot be used by the MicroKernel router, or the file is corrupt.

### **3027 : Connection failure to IDS target server.**

The MicroKernel router could not connect to the IDS target server.

Be sure that IDS is running on the target server. If you are connecting to the IDS target by a host name, ensure that the name is spelled correctly and can be resolved to a TCP/IP address. If you are connecting by a TCP/IP dotted notation address, be sure that the address is correct and accessible from the client. You can use SmartScout as a diagnostic tool.

### **3028 : Login failure to IDS target server.**

The MicroKernel router could not login to the IDS target server.

Be sure that the user name, password, and database set name used are correct. Some applications may prompt for this information; others may require MicroKernel Router configuration parameters to be set. See the Pervasive.SQL User's Guide or the Getting Started with I\*net Data Server for more information regarding MicroKernel Router

configuration parameters.

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## 3100 to 3199

# Network Services Layer Status Codes

The following status codes originate from the Pervasive Network Services Layer (PNLS).

### **3103: Server name not found by Pervasive Network Services Layer.**

The search for a target server name was unable to resolve an address after searching NDS, bindery, named pipes, and DNS. Some possible causes include:

- *No common communication protocol is available between the server and the client. To check this, use the Setup utility to check the Supported Protocols setting of the requester (under Communications Requester section) and the server (under Btrieve Communications Manager and Scalable SQL Communications Manager sections). Ensure that there is a common protocol (either SPX or TCP/IP) between the server and the requester.*
- *For Windows NT: The MicroKernel engine not running.*
- *For NetWare:*
  - *The Server Addressing Protocol (SAP) filtered out the server name. SAP is used in NetWare 3.x and can be used in version 4.x. This type of filtering is common in WAN environments.*
  - *NWCALLS.DLL and/or NWIPXSPX.DLL is missing when using the Win16 client. This is especially likely if the Win32 clients are functioning but the Win16 clients do not. These are NetWare system files that can be obtained from Novell's web site (<http://www.novell.com>).*
- *For NetWare using TCP/IP: The server name is not in DNS.*
- *For Windows NT or OS/2 Warp Server: Named pipes are used in these two operating systems to resolve server names. The networking at the operating system level was unable to find the target named pipe.*

Check the Pervasive Event Log (PVSU.LOG) for more information.

### **3104: The Pervasive Network Services Layer encountered a permission error.**

This status code is relevant to Windows NT and OS/2 Warp only.

Permission to access the target named pipe is denied. If you receive this status code:

- Verify that you have access to the inter process communication mechanism.
- Ensure that your passwords are identical if you have the same user name on two systems.

### **3105: No available transport protocol for the Pervasive Network Services Layer.**

No transport protocol that is common to both the target server engine and clients is available. For example, Status Code 3105 could be caused by a client using SPX when the server engine only has TCP/IP available. To check this, use the Setup utility to examine the Supported Protocols setting of the client (under the Communications Requester section) and the Supported Protocols setting of the server (under the Btrieve Communications Manager and Scalable SQL Communications Manager sections). For more information on the Setup utility, refer to the *Pervasive.SQL User's Guide*.

### **3106: The Pervasive Network Services Layer encountered a connection failure.**

The Pervasive Network Services Layer was able to establish a transport connection at the client side, but the connection attempt at the target side failed. Some possible causes are:

- The Btrieve or Scalable SQL engine is not running on the server.
- The network is overloaded.
- The connection path is invalid.
- You have more than one mapped drive to the same server.
- You are trying to access a Btrieve engine on a Windows NT server and the Accept Remote Requests setting of the Btrieve Communications Manager on that server is set to No. See the *Pervasive.SQL User's Guide* for more information.

Check the Pervasive Event Log (PVSW.LOG) for more information if you receive this status code.

### **3107: The Pervasive Network Services Layer is out of memory.**

The Pervasive Network Services Layer has insufficient memory to continue.

### **3108: The Pervasive Network Services Layer detected an invalid session.**

The application attempted to use a network session that was not recognized by the Pervasive Network Services Layer. If the error persists, contact Pervasive Software Customer Support.

### **3110: The network layer is not connected.**

The application attempted to use a network connection that was no longer active. This happens when the session with the Pervasive Network Services Layer is still valid, but was dropped by the network. Stop and restart the application.

### **3111: Failure during send to the target server.**

The Pervasive Network Services Layer attempted to send an application request to the target server and encountered a network specific error from the target. Check the Pervasive Event Log (PVSW.LOG) for more information.

You can receive this error when the server's Btrieve Communication Manager or Scalable SQL Communication Manager's Number of Sessions setting is too low. To resolve this, increase the Number of Sessions setting for these two components using the Setup utility. Refer to the *Pervasive.SQL User's Guide* for more information on the Number of Sessions setting and the Configuration utility.

### **3112: Failure during receive from the target server.**

The Pervasive Network Services Layer attempted to receive data from the target server and encountered a network specific error. Check the Pervasive Event Log (PVSW.LOG) for more information.

### **3114: The routing table of the Pervasive Network Services Layer is full.**

The Pervasive Network Services Layer's Server Routing Table is full. This table normally grows dynamically as needed. Verify that you are not in a low memory condition.

### **3115: Pervasive Network Services Layer encountered a transport failure.**

The Pervasive Network Services Layer encountered a WinSock (Win32, Win16) or Berkeley Sockets (OS/2) error while creating the local transport endpoint on the client side. Check the Pervasive Event Log (PVSW.LOG) for more information.

### **3125: The Pervasive Network Services Layer was unable to resolve the Novell Directory Services name into a NetWare server name.**

The Pervasive Network Services Layer was unable to resolve the Novell Directory Services name into a NetWare server name. This occurs on Windows 95 workstations with the Microsoft Service for Novell Directory Services in cases where the system is unable to load the Windows 95 Requester support modules. Ensure that the Requester Win95 support modules are installed and available, and that NWCALLS.DLL is available on the Windows 95 workstation.

### **3126: The Pervasive Network Services Layer was unable to resolve the given filename into a valid path.**

Check that the given filename is a valid file format.

### **3127 : Failure during send to the IDS target server.**

Failure during send to the IDS target server.

The IDS client requester attempted to send data to the IDS target server and encountered a TCP/IP specific error. Stop and restart the application. If the problem persists, consult your network administrator.

### **3128 : Failure during receive from the IDS target server.**

Failure during receive from the IDS target server.

The IDS client requester attempted to receive data from the IDS target server and encountered a TCP/IP specific error. Stop and restart the application. If the problem persists, consult your network administrator.

## 7000 to 7099

# User Count Manager Status Codes

This section lists the status codes you can receive when using the User Count Manager.

### **7002: The key number is invalid.**

For a description of this status code, refer to the message [“UCMGR–301: The key number used for increasing the user count is invalid.”](#).

### **7003: The User Count Manager cannot write to the key file.**

For a description of this status code, refer to the message [“UCMGR–405: Error writing to user count key file. \(errno=nn\).”](#).

### **7004: The User Count Manager cannot write to the specified file.**

For a description of this status code, refer to the message [“UCMGR–406: Error writing to file, file = filename. \(errno=nn\).”](#).

### **7005: The User Count Manager cannot read the key file.**

For a description of this status code, refer to the message [“UCMGR–408: Error reading from user count key file. \(errno=nn\).”](#).

### **7006: The User Count Manager cannot read the specified file.**

For a description of this status code, refer to the message [“UCMGR–409: Error reading from file, file = filename. \(errno=nn\).”](#).

### **7007: The User Count Manager encountered a premature end of file.**

For a description of this status code, refer to the message [“UCMGR–410: Unexpected end-of-file found, file = filename. \(errno=nn\).”](#).

### **7008: The User Count Manager cannot open the specified file.**

For a description of this status code, refer to the message [“UCMGR–404: File open error, file = filename. \(errno=nn\).”](#).

### **7009: The User Count Manager cannot modify the specified file.**

For a description of this status code, refer to the message [“UCMGR–403: File change mode error, file = filename. \(errno=nn\).”](#).

### **7010: The User Count Manager cannot find the specified file.**

For a description of this status code, refer to the message [“UCMGR–402: Invalid file pathname = path, file = filename. \(errno = nn\).”](#).

### **7011: You cannot increase the user count with the same key number twice.**

For a description of this status code, refer to the message [“UCMGR–305: Duplicate key number error. You cannot increase the user count with the same key twice.”](#).

### **7012: The key file’s integrity has been compromised.**

For a description of this status code, refer to the message [“UCMGR–302: Invalid key file check sequence. The integrity of the user count key file has been compromised.”](#).

### **7013: The key type is invalid.**

For a description of this status code, refer to the message [“UCMGR–303: Invalid key number \(key number is of the wrong type\).”](#).

### **7014: The User Count Manager is already in use. Try again later.**

For a description of this status code, refer to the message [“UCMGR–501: The User Count Manager is busy. Only one instance of the User Count Manager can be in use at one time.”](#).

### **7016: The User Count Manager cannot read the specified file; the file may be corrupt.**

For a description of this status code, refer to the message [“UCMGR–411: File seek error, file = filename, \(errno=nn\).”](#)

### **7017 – 7023: The component serial numbers do not match.**

For a description of this status code, refer to the UCMGR messages 310 through 316, beginning with [“UCMGR–310: Serial number synchronization error. Unable to find valid stamp area.”](#) on page 2-80.

7017: [“UCMGR–310: Serial number synchronization error. Unable to find valid stamp area.”](#) on page 2-80

7018: [“UCMGR–311: Serial number synchronization error. Invalid stamp area read.”](#) on page 2-80

7019: [“UCMGR–312: Serial number synchronization error. One or more stamp numbers read do not match.”](#) on page 2-80

7020: [“UCMGR–313: Serial number synchronization error. Could not write stamp number to one of the files. Make sure the database engine is not loaded and that you have access rights to the files.”](#) on page 2-80

7021: [“UCMGR–314: Serial number synchronization error. Files have not been stamped and the user count key file exists.”](#) on page 2-80

7022: [“UCMGR–315: Serial number synchronization error. The key file stamp number does not match the stamp number in the other files.”](#) on page 2-81

7023: [“UCMGR–316: Serial number synchronization error. Files have been stamped but the user count key file does not exist. If you have made a backup copy of your user count key file \(<plat>UCMGR.SYS\), please restore it to the database engine directory.”](#) on page 2-81

### **7026: The User Count Manager cannot find the diskette key file: UCMGR.KEY.**

For a description of this status code, refer to the message [“UCMGR–401: Cannot find diskette key file, file = filename, \(errno = nn\).”](#).

## **7028: The diskette key has already been applied.**

For a description of this status code, refer to the message ["UCMGR-306: The diskette key has already been applied. You cannot increase the user count with the same diskette key twice."](#).

## **7029: The User Count Manager cannot read the diskette key file.**

For a description of this status code, refer to the message ["UCMGR-307: The diskette key file is invalid. file = filename. \(errno=nn\)"](#).

## **7030: DOS partition services are not present.**

This status code applies to the NetWare platform only. For a description of this status code, refer to the message ["UCMGR-415: DOS Partition Services are not present."](#).

## **7031: The User Count Manager cannot copy the diskette key file.**

This status code applies to the NetWare platform only. For a description of this status code, refer to the message ["UCMGR-416: DOS Copy error. The User Count Manager cannot copy the diskette key file. \(errno=nn\)"](#).

## **7032: The User Count Manager cannot open the diskette key file.**

This status code applies to the NetWare platform only. For a description of this status code, refer to the message ["UCMGR-417: DOS Open error. The User Count Manager cannot open the diskette key file. \(errno=nn\)"](#).

## **7033: The User Count Manager cannot write to the diskette key file.**

This status code applies to the NetWare platform only. For a description of this status code, refer to the message ["UCMGR-418: DOS Write error. The User Count Manager cannot write to the diskette key file \(errno=nn\)"](#).

## **7035-7038: The serial numbers do not match.**

For a description of these status codes, refer to UCMGR messages 317 through 319, beginning on [page -81](#). Messages 7035-7038 share the same error messages as 7017-7023. The difference is that 7035-7038 occur only during a user count migration (such as a migration from a Btrieve v6.15 user count).

## **7039: The User Count Manager cannot delete the specified file.**

For a description of this status code, refer to the message ["UCMGR-412: File delete error. file = filename. \(errno=nn\)"](#).

## **7040: The User Count Manager cannot rename the specified file.**

For a description of this status code, refer to the message ["UCMGR-413: File rename error. file = filename. \(errno=nn\)"](#).

## **7041: The User Count Manager is not initialized.**

For a description of this status code, refer to the message ["UCMGR-500: The User Count Manager is not initialized. Unload or stop your database engine and run the User Count Initialization utility \(UCINIT.EXE, NTUCINIT.EXE, or NWUCINIT.NLM\)"](#).

## **7042: The User Count Manager cannot determine the platform.**

For a description of this status code, refer to the message [“UCMGR–502: The User Count Manager cannot determine the platform.”](#).

### **7046: The User Count Manager cannot migrate Btrieve user count.**

An invalid Btrieve database engine was specified for migrating the Btrieve user count. On NetWare, only Btrieve v6.15 and Btrieve v6.10 user count licenses can be migrated to Scalable SQL v4.0. On Windows NT, only Btrieve v6.15 user count licenses can be migrated to Scalable SQL v4.0.

### **7048: The operation is invalid for Btrieve v6.1x migration.**

This status code has the same meaning as the message [“UCMGR–321: You have previously migrated xx users and have tried to migrate xx users at this time. This is an invalid operation because the migrated user count is less than the previously migrated user count.”](#).

### **7049: User count license is already unlimited for this product.**

You receive this status code when you attempt to add a diskette-based license key on top of a product that already is licensed for unlimited users. Since you already have an unlimited license, you do not need to increase your user count.

### **7050: Incorrect user count resource DLL.**

You receive this status code when the user count manager loads the wrong message file. Update your resource DLL to the version specified by the message.

### **7051: Incorrect major version of user count resource DLL.**

You receive this status code when the user count manager loads an older major version of the message file. Update your resource DLL to the version specified by the message.

### **7052: Incorrect minor version of user count resource DLL.**

You receive this status code when the user count manager loads an older minor version of the message file. Update your resource DLL to the version specified by the message.

### **7053: Incorrect patch number for user count resource DLL.**

You receive this status code when the user count manager loads an older patch level of the message file. Update your resource DLL to the path number specified by the message.

### **7054: License key does not match product version.**

You receive this status code when you attempt to add a key for an older version of the product. Obtain a user license key for the current product level and try again.

### **7061: Your Pervasive Database engine expired on <date>. Please contact Pervasive Software for more information.**

You receive this status code when your temporary user license for the database engine is expired. Contact Pervasive Software to obtain a permanent user license key.

### **7062: Your Pervasive Database engine will expire on <date> because you have**

## **a temporary license key.**

You receive this message when you start the MicroKernel and you have a temporary license key. This is an informational message only that reminds you when your license will expire. Contact Pervasive Software to obtain a permanent license key.

## **7063: Your user count license limit has been reached.**

The maximum allowed users for the database engine has been reached. Please contact Pervasive Software if you need additional users for your license. To check the current setting for maximum users, run the User Count Administrator:

1. Click **Start**, Point to **Programs** and then to **Pervasive SQL 7**.
2. Select the **User Count Administrator**.

The User count Administrator displays the current maximum user licenses and allows you to add additional licenses. Refer to *Getting Started With Pervasive.SQL* for more information.

## **7064: The license you tried to remove does not exist.**

You receive this status code when the User Count Manager receives a request to remove a license for a product code, and the license does not exist.

## **7065: The product for which you requested the user count license does not exist.**

You receive this status code when the User Count Manager receives a request to display a user count license for a product code, and that product code does not exist. Consult *Getting Started with Pervasive Database* for a list of valid product codes.

## 8000 to 8499

### Component Management Status Codes

These status codes originate from the Smart Component Management and Event Logging API of the MicroKernel.

#### **8001: The Services DLL encountered an error while allocating memory.**

Memory could not be allocated. Verify that your system resources are not exhausted.

#### **8002: The component ID string was not found.**

The Services DLL attempted to retrieve an ID from a component, and the component did not have an ID. This status code results from an error in the application.

#### **8003: Thunk path is not accessible to the Services DLL.**

Smart Component Management cannot find the address of the Win32 APIs.

#### **8005: API not initialized.**

You receive this status code when an error causes an API to not initialize properly. Check the Pervasive Event Log (PVSW.LOG) for more information.

#### **8006: The specified component was not found.**

Smart Component Management cannot operate on the component because it was not found. Check the Pervasive Event Log (PVSW.LOG) for more information.

#### **8007: Error retrieving Win32 version information.**

An error occurred while retrieving the Win32 version information from the operating system. This error should not occur and it indicates a problem with your Windows environment. Contact Pervasive Software Customer Support.

#### **8008: Information mismatch in component ID.**

You receive this status code when an application has specified multiple platforms, types or versions in a component ID. The Services DLL therefore cannot bind to the correct component.

#### **8009: Required information missing in component ID.**

You receive this status code when a component ID string is missing required information such as:

- Platform
- Type
- Major/minor functional level

#### **8010: Error while opening a file to retrieve component ID.**

The Services DLL attempted to open a file in order to retrieve a component ID, and the file did not exist.

## **8012: Required information for binding rule is missing.**

You receive this status code when required information is missing, such as:

- Platform type
- Major/minor functional level

## **8013: Error locating a component.**

An error occurred while opening an operating system directory to search for a component. Check the Pervasive Event Log (PVSU.LOG) for more information.

## **8016: Invalid instance specified.**

This error occurs when an application passes an instance handle to the Services DLL that is invalid.

## **8017: Invalid version information structure.**

This error occurs when an application tries to free a version information structure and the Services DLL determines that the version structure is invalid.

## **8018: Item not found by Services DLL.**

This error occurs when an application requests that the Services DLL free an item, and the Services DLL determines that the item did not originate from the Services DLL.

## **8019: Library not accepting API requests.**

This error occurs when the Services DLL cannot accept requests. In a multithreaded application, this can occur when one thread attempts to make a request while other threads are overloading the Services DLL. In this situation, subsequent retries should clear the error.

## **8020: Error loading component.**

You receive this status code when the Services DLL receives an operating system error when trying to load a component. Check the Pervasive Event Log (PVSU.LOG) for more information. If there is no information in the Pervasive Event Log related to receiving a status code 8020, refer to "Diagnosing Load Errors" in the "Component Architecture" chapter of the *User's Guide*.

## **8022: Component not initialized.**

You should not receive this status code except when using a multithreaded application where one thread has shut down a component and the others continue trying to use it.

## **8023: Invalid filename passed to Services DLL.**

You receive this status code when the Services DLL receives a filename that does not exist.

## **8024: Invalid function name.**

You receive this status code when an application asks to import a function that does not exist.

### **8027: Required attribute missing in component ID.**

You receive this status code when an application program does not specify a required attribute in the component ID.

### **8030: Return buffer too small for information.**

This error occurs when an application passes the Services DLL a buffer in which to return information, and that buffer is not large enough to contain the information.

### **8032: Component ID is invalid.**

You receive this status code when a component ID string was found but it may be invalid because it is too short, too long, or contains invalid characters.

### **8035: Invalid API for NetWare platform.**

You receive this status code when an application makes a call to an API that is not supported on the NetWare platform.

### **8039: Invalid operating system handle specified.**

You receive this status code when an operating system handle to a component ID string is invalid. This can occur when an application erroneously modifies a handle obtained using the GetFileVersionByHandle or the GetFileSizeByHandle functions.

### **8040: CRC check on component ID failed.**

A routine CRC check on the component ID failed. You receive this status code when an application erroneously alters a component ID.

## 8500 to 8589

### ECAS API Status Codes

This section describes the status codes returned by the ECAS (Enhanced Common Address Space) API. Most of the errors are system errors and cause an entry in the Pervasive event log. In many cases when you receive these errors, there may be some instability in the environment or memory allocation and you will need to restart the machine.

#### **8500: An error occurred during the Smart Components initialization.**

In its attempt to auto-load the workstation engine, the application failed to initialize the smart component library. See the Pervasive event log for more information. You may also get more information about the problem that caused this error by setting the PVSW\_DISP\_LOAD\_ERRS=AIF environment variable and running the application. This enables the Services DLL to display on-screen module load errors. However, this environment variable should only be set to diagnose module load errors. In all other cases, it should not be set.

#### **8502: An error occurred when trying to locate W3UPIxyy.DLL.**

In its attempt to auto-load the workstation engine, the application failed to locate or load W3UPIxyy.DLL in the path specified by the PERVASIVE\_PATH environment variable. See the Pervasive event log for more information.

#### **8503: An invalid W3UPIxyy.DLL has been found.**

In its attempt to auto-load the workstation engine, the application discovered an incorrect version of W3UPIxyy.DLL. This problem may have been caused by a corrupted version of W3UPIxyy.DLL.

#### **8504: An error occurred when trying to create system semaphore.**

In its attempt to auto-load the workstation engine, the application failed to create the system semaphore. This problem may have been caused by the operating system running out of resources. Check the Pervasive event log for more information.

#### **8505: An initialization error occurred when trying to establish a session with the workstation engine.**

In its attempt to auto-load the workstation engine, the application failed to establish the session with the workstation engine. This is a system error. Check the Pervasive event log for more information.

#### **8506: A fatal error occurred when loading the MicroKernel.**

In its attempt to auto-load the workstation engine, the application failed to load the MicroKernel, W3MKDE.DLL. This may have been caused by a missing W3MKDE.DLL. Check the Pervasive event log for more information.

#### **8507: No valid database utility session was found.**

The application lost its session with the workstation engine. Check the Pervasive event log for more information.

#### **8508: An error occurred when attempting to enable Btrieve access.**

The application failed to enable the Btrieve access method in the workstation engine. Check the Pervasive event log for more information.

### **8509: A timeout occurred during the initialization of the MicroKernel.**

The application timed out during the initialization of the MicroKernel. This may have been caused by a bad configuration option or a malfunction of the workstation engine. Check the Pervasive event log for more information.

### **8510: A fatal error occurred when loading the Scalable SQL engine.**

In its attempt to auto-load the workstation engine, the application failed to load the Scalable SQL engine, W3SSQL.DLL. This may have been caused by a missing W3SSQL.DLL. Check the Pervasive event log for more information.

### **8511: An error occurred when attempting to enable SQL access.**

The application failed to enable the SQL access method in the workstation engine. Check the Pervasive event log for more information.

### **8512: A timeout occurred during the initialization of the Scalable SQL engine.**

The application timed out during the initialization of the Scalable SQL engine. This may have been caused by a bad Scalable SQL engine configuration option or a malfunction of the workstation engine. Check the Pervasive event log for more information.

### **8513: An error occurred when disabling Btrieve access.**

The application failed to disable the Btrieve access method in the workstation engine. Check the Pervasive event log for more information.

### **8514: An error occurred when unloading the MicroKernel.**

The application failed to unload the MicroKernel. Check the Pervasive event log for more information.

### **8515: An error occurred when disabling SQL access.**

The application failed to disable the SQL access method in the workstation engine. Check the Pervasive event log for more information.

### **8516: An error occurred when unloading the Scalable SQL engine.**

The application failed to unload the Scalable SQL engine. Check the Pervasive event log for more information.

### **8517: An error occurred when closing the session with the workstation engine.**

The application could not close the session with the workstation engine. Check the Pervasive event log for more information.

### **8518: An error occurred when attempting to allocate system memory.**

The application failed to allocate memory from the system. Possible ways to avoid this include closing all other applications and restarting the engine, decreasing the size of the cache, and changing the engine settings so that a smaller number of files and/or file handles are open.

## 8590 to 8599

### W3DBSMGR Status Codes

This section describes the status codes returned by W3DBSMGR logic. These errors are in the range 8590 to 8599 and cause an entry in the event log. In many cases when you receive these errors, there may be some instability in the environment or memory allocation and you will need to restart the machine.

#### **8590: An error occurred in the database manager while initializing Smart Components Manager.**

This error occurs when W3DBSMGR (Pervasive.SQL workstation engine) failed to initialize the smart component library. See the Pervasive event log for more information. You may also get more information about the problem that caused this error by setting the PVSW\_DISP\_LOAD\_ERRS=AIF environment variable and running the application. This enables the Services DLL to display on-screen module load errors. However, this environment variable should only be set to diagnose module load errors. In all other cases, it should not be set.

#### **8591: The database manager is already loaded.**

In its attempt to initialize, the workstation engine detected that another copy of the Pervasive.SQL workstation engine is already running in memory. To avoid this error, shut down the program that is already running as specified below:

- If the program is autoloaded (i.e., loaded through a Btrieve or a Scalable application), then terminate the application
- If the program is preloaded, then shut down the application using the tray icon.

#### **8592: Insufficient memory to load the database manager.**

This error indicates that the system is out of resources. Close some of the applications and try again.

#### **8593: An error occurred while the database manager was creating a system thread.**

This error indicates that the system is out of resources. Close some of the applications and try again.

#### **8594: Engine cannot be restarted.**

The workstation engine has been previously stopped while there were active applications and cannot be restarted. Close all current applications and try again. If you receive this code once more, you will need to restart the computer.



# Messages

This chapter describes the messages you can encounter while using the Scalable SQL or Btrieve data access systems or the MicroKernel. The messages are listed in alphabetic order by category. Some of the messages listed in this chapter may be generated by utilities that are not part of your Pervasive.SQL distribution.

Messages are not the same as status codes. (see [Chapter 1, "Status Codes"](#)); messages are returned to end users by utilities or specific components and generally begin with a prefix and a number. For example: "[MKDE-16: There is insufficient memory to load the MicroKernel Database Engine.](#)"

Some messages refer to status codes, as in the following example:

Error Accessing INDEX.DDF. Status = nn

The status code is tied to a particular component, and it varies depending on the nature and source of the error. The value for *nn* or *xx* refers to a software-supplied number (for example, a status code, operation code, or number of records) and *xxxx* refers to a software-supplied name (for example, a filename, a data type or key type, or a command).

The messages are arranged according to the originating component. [Table 2-1](#) lists the prefixes for each type of message.

**Table 2-1**  
**Message Groups**

| Prefix   | Type of Code                                                           |
|----------|------------------------------------------------------------------------|
| (none)   | <a href="#">Generic Messages</a>                                       |
| BDROUTER | <a href="#">Btrieve Remote Communications Manager Messages</a>         |
| BREQUEST | <a href="#">Btrieve Requester Messages</a>                             |
| BREQUTIL | <a href="#">Requester Utility Messages</a>                             |
| BROUTER  | <a href="#">Btrieve Message Router Messages</a>                        |
| BSPXCOM  | <a href="#">Btrieve SPX Communications Module Messages</a>             |
| BTCPCOM  | <a href="#">Btrieve TCP/IP Communications Module Messages</a>          |
| BUTIL    | <a href="#">Btrieve Maintenance Utility Messages</a>                   |
| MKDE     | <a href="#">MicroKernel Database Engine Messages</a>                   |
| NWBSRVCM | <a href="#">NetWare Btrieve Server Communications Manager Messages</a> |

|          |                                                                                    |
|----------|------------------------------------------------------------------------------------|
| NWSSRVCM | <a href="#"><u>NetWare Scalable SQL Server Communications Manager Messages</u></a> |
| SSPXCOM  | <a href="#"><u>Scalable SQL SPX Communications Module Messages</u></a>             |
| SSQL     | <a href="#"><u>Scalable SQL Engine Messages</u></a>                                |
| STPCOM   | <a href="#"><u>Scalable SQL TCP/IP Communications Module Messages</u></a>          |
| UCINIT   | <a href="#"><u>User Count Initialization Utility Messages</u></a>                  |
| UCMGR    | <a href="#"><u>User Count Manager Messages</u></a>                                 |
| UCUTIL   | <a href="#"><u>User Count Utility Messages</u></a>                                 |

## (No prefix) Generic Messages

The following messages may be returned by more than one utility or component. These messages are not preceded by a component prefix or number.

### **Bfloat and Float Fields Must Be 4 or 8 Bytes**

You have defined a FLOAT or BFLOAT column with a size other than 4 or 8 bytes. Change the column size to 4 or 8 bytes.

### **First and Second Password Entries Do Not Match, Please Retype**

The password you entered the second time is not the same as the first one you specified. Specify and verify your password again.

### **Incorrect Netware Version**

The NetWare version loaded is incompatible with the version of BREQUEST that is loaded. The NetWare shell for NetWare v3.12 or above must be installed at the client machine.

### **Incorrect Parameter**

The Scalable SQL or Btrieve Requester returns this message if it encounters an error while loading at a DOS client machine. You specified an invalid configuration option. For information about the Setup utility and specifying the configuration options, refer to the *Pervasive.SQL User's Guide*.

### **Insufficient Memory**

The Scalable SQL or Btrieve Requester returns this message if it encounters an error while loading at a DOS client machine. The Requester cannot allocate enough memory for the configuration options specified. Ensure that the client machine has enough memory to load all the programs you want to run.

### **Invalid Load File Format. End of Record Marker Not Found**

The Btrieve Maintenance utility (BUTIL) encountered an invalid record terminator in the sequential load file. The Btrieve Maintenance utility expects a carriage return/line feed at the end of each record in a load file. This error usually occurs because the length specified at the beginning of the sequential record is incorrect.

### **Load File Referenced Undefined Public Variable. Module xxxx Not Loaded**

In the NetWare v3.12 or later environment, a module you attempted to load (such as the Setup utility) requires that the correct version of another module be loaded. All Scalable SQL utilities in Scalable SQL 4.0 require that the following software is loaded:

- NetWare v3.12 or later, including the AFTER311 NLM. Load AFTER311 from AUTOEXEC.NCF before calling BSTART.
- Scalable SQL 4.0, including the Scalable SQL and SSPXCOM NLMs
- MicroKernel Database Engine v7.x and the BSPXCOM NLM

### **Local Engine Is Already Loaded**

The Scalable SQL Requester returns this message if it encounters an error while loading at a DOS client machine. A local engine (XQL v2.11 or earlier) is loaded at the machine. You cannot load the older local engine and the newer Requester at the same time.

### **Must Have DOS 3.1 or Greater**

The Scalable SQL or Btrieve Requester returns this message if it encounters an error while loading at a DOS client machine. The Requesters require DOS v3.1 or later.

### **NetWare Shell Is Not Installed**

The Scalable SQL Requester returns this message if the NetWare shell is not installed on the DOS client machine at which you want to load the Requester. You must install the NetWare shell on the client machine before an application can communicate with Scalable SQL.

### **Redirection List Not Large Enough**

The Scalable SQL Requester returns this message if it encounters an error while loading at a DOS client machine, and it cannot store all the redirected devices in its internal redirection list. Increase the value for the /R configuration option.

### **Scalable SQL Requester Is Already Loaded**

The Scalable SQL Requester returns this message if you try to load the Requester when it is already loaded at the DOS client machine.

### **SPX Is Not Loaded**

The Scalable SQL or Btrieve Requester returns this message if it encounters an error while loading at a DOS client machine. You must load the NetWare SPX communications software before an application can access the server-based version of Scalable SQL or the MicroKernel. For more information about installation, refer to *Getting Started With Pervasive.SQL*.

# BDROUTER

## Btrieve Remote Communications Manager Messages

The following messages are specific to the BDROUTER Remote Communications Manager. This component is obsolete in Pervasive.SQL 7.0, so these messages are only returned by v6.x engines.

### **BDROUTER–1: The server has insufficient memory to execute BDROUTER.**

The server has insufficient memory to load the BDROUTER NLM. Free some memory by unloading other NLMs or reconfiguring the NLMs to use less memory.

### **BDROUTER–2: The value specified for a configuration option is invalid.**

The value specified for a configuration option is invalid. Reload BDROUTER.NLM using valid options.

### **BDROUTER–3: An internal error has occurred; the SPXOpenSocket function failed.**

An internal diagnostic error has occurred. The SPXOpenSocket function failed. Another NLM may be using the socket number reserved for BDROUTER. If you receive this message, unload all other NLMs and then load BTRIEVE.NLM and BDROUTER.NLM. Finally, reload the other NLMs. This process reveals the NLM that is using BDROUTER's socket number.

### **BDROUTER–11: The value specified for Directory Services username and password is invalid.**

Before an application can access the NetWare Directory Services database, the user must be authenticated to NDS. For the client requesters, it is assumed that you have already been authenticated; therefore, the Requesters do not explicitly perform an NDS login. However, BDROUTER must explicitly log in to NDS and therefore must provide an NDS username and password. For this purpose, BDROUTER has two command line options, as follows:

- */U:<NDSusername>*
- */P:<NDSPassword>*

These options are required and have no default values.

# BREQUEST

## Btrieve Requester Messages

The following messages are specific to Btrieve Requesters.

### **BREQUEST-1: The message file xxxx is invalid; BREQUEST cannot be loaded.**

Install a valid message file so that the Requester can be loaded.

### **BREQUEST-4: The value specified for the Data Message Length (/d) option is invalid.**

Specify the /D option as /D:n, where n is a number between 532 and 57,000. For more information about DOS Requester options, refer to the *Getting Started With Pervasive.SQL*.

### **BREQUEST-5: The workstation has insufficient memory to load BREQUEST.**

Unload unnecessary programs or try a smaller value for the /D parameter. For more information about DOS Requester options, refer to *Getting Started With Pervasive.SQL*.

### **BREQUEST-7: The Btrieve Requester must be loaded before the Client Engine or Requester.**

The Scalable SQL workstation engine or Requester is already loaded; the Btrieve Requester for DOS must be loaded first. Unload the Scalable SQL workstation engine or Requester and then load the Btrieve Requester for DOS.

### **BREQUEST-8: DOS 3.00 or greater is not loaded; load DOS 3.00 or greater.**

Load DOS 3.x or later to proceed.

### **BREQUEST-9: SPX is not available on the workstation.**

Load SPX.COM on your client machine.

### **BREQUEST-10: The function SPXInitialize returned an error.**

Ensure that the file IPX.COM is loaded.

### **BREQUEST-11: The IPX socket table is full.**

The Btrieve Requester for DOS returns this message when the IPX socket table is full.

### **BREQUEST-12: The value specified for the NetWare Runtime Server support (/C) option is invalid.**

Specify this option in one of these forms:

**/C:0**    Disables NetWare Runtime server support.

**/C:1** Enables NetWare Runtime server support. To authenticate requests on the NetWare Runtime server, provide a username and password, separating them with commas, as follows:  
`/C:1,username,password`

For more information about DOS Requester options, refer to the *Getting Started With Pervasive.SQL*.

### **BREQUEST-13: The value specified for the Number of Servers (/S) is invalid.**

Specify the /S option as /S:n, where n is a number between 1 and 8. For more information about DOS Requester options, refer to *Getting Started With Pervasive.SQL*.

### **BREQUEST-14: Btrieve has been removed from memory.**

This is an informational message that displays when you use the BREQUEST or BREQNT requester programs with the /u (unload) parameter. See *Getting Started With Pervasive.SQL* for more information on the command-line parameters of BREQUEST and BREQNT.

### **BREQUEST-15: Btrieve is not loaded.**

You used the BREQUEST or BREQNT requester programs with the /u parameter that unloads Btrieve from memory, and Btrieve was not running at the time.

## BREQUTIL

### Requester Utility Messages

This section lists the messages that the Scalable SQL and Btrieve Requester utilities generate.

#### **BREQUTIL–8: MicroKernel or Btrieve Requester is not loaded.**

You must load the MicroKernel or a Btrieve Requester before you can perform a Btrieve operation.

#### **BREQUTIL–9: Btrieve operation *nn* was unsuccessful. The number of the applicable Btrieve status code is *nn*.**

Refer to [Chapter 1. “Status Codes”](#) for a description of Status Code *nn*.

#### **BREQUTIL–16: MicroKernel or Btrieve Requester cannot be removed from memory while Scalable SQL is loaded.**

If you want to unload the MicroKernel or Btrieve Requester, unload Scalable SQL first.

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

# BROUTER

## Btrieve Message Router Messages

The following messages are specific to the BROUTER Remote Communications Manager.

### **BROUTER-2: The value specified for a configuration option is invalid.**

Reload BROUTER.NLM using valid configuration options.

### **BROUTER-3: An internal error has occurred; the SPXOpenSocket function failed.**

An internal diagnostic error occurred. The SPXOpenSocket function failed. Another NLM may be using the socket number reserved for BROUTER. If you receive this message, unload all other NLMs and then load BTRIEVE.NLM and BROUTER.NLM. Finally, reload the other NLMs. This process reveals the NLM that is using BROUTER's socket number.

### **BROUTER-7: The server has insufficient memory to execute BROUTER.**

Free some memory by unloading NLMs or reconfiguring NLMs to use less memory.

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](http://www.pervasive.com) All rights reserved.

## BSPXCOM

### Btrieve SPX Communications Module Messages

The following messages are specific to the Btrieve Communications Manager.

#### **BSPXCOM-2: The server has insufficient memory to execute BSPXCOM.NLM**

An error occurred during a memory allocation operation. Free memory by unloading NLMs or reconfiguring NLMs to use less memory. Retry the operation. If the error persists, contact Pervasive Software Customer Support.

#### **BSPXCOM-3: An internal error has occurred. BSPXCOM detected a semaphore allocation failure.**

An internal error occurred during a semaphore allocation operation. Retry the operation. If the error persists, contact Pervasive Software Customer Support.

#### **BSPXCOM-4: The Service Request Block (SRB) function code nn contains invalid data. Check for an incompatible version of the file BSPXCOM.NLM.**

Ensure that BSPXCOM's version is compatible with the version number of the workstation's Btrieve Requester.

#### **BSPXCOM-6: Another NLM is using the socket number reserved for BSPXCOM.NLM.**

To solve this condition, unload all other NLMs. Load BTRIEVE.NLM and NWBSRVCM.NLM, and then reload the other NLMs. This process reveals the NLM that is using BSPXCOM.NLM's socket number.

#### **BSPXCOM-7: An SPX receive I/O error (nn) has occurred. The connection has been lost.**

An error occurred during an SPX-level receive operation. Refer to your operating system documentation for the definition of the *error number*. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **BSPXCOM-8: An SPX send I/O error (nn) has occurred. The connection has been lost.**

An error occurred during an SPX-level send operation. Refer to your operating system documentation for the definition of the *error number*. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **BSPXCOM-10: Bad connection ID detected on receive. The SPX connection was lost after the initial request began.**

If this message appears frequently, increase your requester's SPX Timeout parameter, which is controlled with the Configuration (Setup) utility. Also, check for NLMs that are monopolizing the CPU time.

#### **BSPXCOM-11: Bad connection ID detected on send. The SPX connection was**

## **lost after the initial request began.**

If this message appears frequently, increase your requester's SPX Timeout parameter, which is controlled with the Configuration (Setup) utility. Also, check for NLMs that are monopolizing the CPU time.

## **BSPXCOM-12: An error (nn) was detected while trying to establish an SPX session requested by a remote workstation.**

An error occurred during an SPX-level session establishment operation. Refer to your operating system documentation for the definition of the *error number* value. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

## **BSPXCOM-13: The session was rejected because the session limit was reached. Increase the value specified for the Number of Remote Sessions option.**

Increase the value specified for the Number of Remote Sessions configuration option, as described in the *Pervasive User's Guide*.

## **BSPXCOM-14: The request for statistics from the Monitor utility was not recognized. Check for an incompatible version of the utility or BSPXCOM.NLM.**

Ensure that your BSPXCOM.NLM version is compatible with the version number of the Monitor utility.

## **BSPXCOM-15: An internal error has occurred. BSPXCOM did not recognize the GET\_EIM\_STATS function.**

Ensure that the version of BSPXCOM.NLM is compatible with the version number of the Monitor utility.

## **BSPXCOM-18: A thread initialization error has occurred.**

An error occurred during BSPXCOM.NLM thread initialization. Refer to the Pervasive Software Event Log (PVSU.LOG) for error messages preceding this one. Correct the problem and retry the operation.

## **BSPXCOM-19: A BeginThread() error (nn) has occurred.**

An error occurred during a BeginThread operation. Refer to your operating system documentation for the definition of the *error number* value. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

## BTCPCOM

### Btrieve TCP/IP Communications Module Messages

#### **BTCPCOM-1: The option specified is not a valid option.**

Specify a valid option.

#### **BTCPCOM-2: The server has insufficient memory to execute BTCPCOM.NLM.**

An error occurred during a memory allocation operation. Free memory by unloading NLMs or reconfiguring NLMs to use less memory. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **BTCPCOM-3: An internal error has occurred. BTCPCOM.NLM detected a semaphore allocation failure.**

An error occurred during a semaphore allocation operation. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **BTCPCOM-4: The Service Request Block (SRB) function code nn contains invalid data. Check for an incompatible version of the file BTCPCOM.NLM.**

Ensure that BTCPCOM's version is compatible with the version number of the workstation's Btrieve Requester.

#### **BTCPCOM-5: A TCP/IP level receive I/O error nn has occurred. The connection has been lost.**

An error occurred during a TCP/IP level receive operation. Refer to your operating system documentation for the definition of the *error number* value. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **BTCPCOM-6: A TCP/IP level send I/O error nn has occurred. The connection has been lost.**

An error occurred during a TCP/IP level send operation. Refer to your operating system documentation for the definition of the *error number*. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **BTCPCOM-7: A TCP/IP accept error (nn) has occurred.**

An error occurred during a TCP/IP level accept session establishment operation. Refer to your operating system documentation for the definition of the *error number*. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **BTCPCOM-8: The session was rejected because the session limit was reached. Increase the value specified for the Number of Remote Sessions option.**

Increase the value specified for the Number of Remote Sessions option. This option is described in the *Pervasive User's Guide*.

**BTCPCOM-9: The request for statistics from the Monitor utility was not recognized. Check for an incompatible version of the utility or BTCPCOM.NLM.**

Ensure that the BTCPCOM.NLM version is compatible with the version number of the Monitor utility.

**BTCPCOM-10: An internal error has occurred. BTCPCOM did not recognize the GET\_EIM\_STATS function.**

Ensure that the BTCPCOM.NLM version is compatible with the version number of the Monitor utility.

**BTCPCOM-12: A TCP/IP select() error (nn) has occurred while trying to establish a session requested by a remote workstation.**

An error occurred during a TCP/IP-level select session establishment operation. Refer to your operating system documentation for the definition of the *error number*. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

**BTCPCOM-13: A thread initialization error has occurred.**

An error occurred during BTCPCOM.NLM thread initialization. Refer to the Pervasive Software event log (PVSW.LOG) for error messages preceding this one. Correct the problem and retry the operation.

**BTCPCOM-15: A TCP/IP socket() error (nn) has occurred.**

An error occurred during a TCP/IP-level socket operation. Refer to your operating system documentation for the definition of the *error* value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

**BTCPCOM-16: A TCP/IP setsockopt() error (nn) has occurred.**

An error occurred during a TCP/IP-level setsockopt operation. Refer to your operating system documentation for the definition of the *error* value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

**BTCPCOM-17: A TCP/IP bind() error (nn) has occurred.**

An error occurred during a TCP/IP-level bind operation. Refer to your operating system documentation for the definition of the *error* value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

**BTCPCOM-18: A TCP/IP listen() error (nn) has occurred.**

An error occurred during a TCP/IP-level listen operation. Refer to your operating system documentation for the definition of the *error* value. Retry the operation again. If the error persists, call Pervasive //Software technical support for more information.

**BTCPCOM-19: A TCP/IP gethostname() error (nn) has occurred.**

An error occurred during a TCP/IP-level gethostname operation. Refer to your operating system documentation for the definition of the *error* value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

### **BTCPCOM-20: A TCP/IP NWgethostbyname() error (nn) has occurred.**

An error occurred during a TCP/IP-level NWgethostbyname operation. Refer to your operating system documentation for the definition of the *error* value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

### **BTCPCOM-21: A TCP/IP NetDBgethostbyname() error (nn) has occurred.**

An error occurred during a TCP/IP-level NetDBgethostbyname operation. Refer to your operating system documentation for the definition of the *error* value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

### **BTCPCOM-22: A BeginThread() error (nn) has occurred.**

An error occurred during a BeginThread operation. Refer to your operating system documentation for the definition of the *error* value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

# BUTIL

## Btrieve Maintenance Utility Messages

The following messages are specific to the Btrieve Maintenance utility.

### **BUTIL-1: Message file initialization failed.**

Only BUTIL.EXE in DOS returns this message. The utility could not find the BUTILDOS.MSG file, which contains messages that the utility uses. Ensure that you installed NetWare Communication Services correctly.

### **BUTIL-2: Message file initialization failed. BUTIL was unable to access BTIUTILD.MSG.**

Only BUTIL.EXE in DOS returns this message. The utility could not find the BTIUTILD.MSG file, which contains messages that the utility uses. Ensure that you installed Btrieve for DOS correctly.

### **BUTIL-6: The BUTIL command is invalid.**

The syntax of the command you entered is incorrect. Verify the syntax before re-entering the command.

### **BUTIL-8: The command completed, but one or more errors occurred.**

An error occurred when you executed a command that performed a number of MicroKernel operations. These commands include COPY, LOAD, or CLONE. This message is accompanied by additional messages that can help you identify the problem.

### **BUTIL-9: The command did not complete due to an unrecoverable error.**

Verify that the syntax you entered is correct before re-entering the command. This message is accompanied by additional messages that can help you identify the problem.

### **BUTIL-10: The command line contains a syntax error.**

The syntax of the command you entered is incorrect. Verify the syntax before re-entering the command.

### **BUTIL-11: The command line requires the index file.**

If you specify the BUTIL -INDEX or -SAVE command (modified by the Y parameter) to the Btrieve Maintenance utility, you must specify the full pathname of an external index file.

### **BUTIL-12: The command line requires the key number.**

If you specify the DROP command or the SAVE command (modified by the N parameter) to the Btrieve Maintenance utility, you must specify the key number of the key you want to drop or by which you want to save the data file.

### **BUTIL-13: The key size for key of type xxxx is invalid.**

In a description file, the specified value of the Key Length element for a particular key is incorrect. Ensure that the value of each Key Length element is appropriate for the matching Key Type element.

### **BUTIL-15: Error occurred on key segment descriptor xxx of description file.**

An error occurred when the Maintenance utility attempted to read the key segment descriptor portion of the description file. Ensure that the elements that describe the specified key segment are valid.

### **BUTIL–16: BUTIL could not open the description file.**

Before attempting to re-enter the CREATE, INDEX, or SINDEXT commands, ensure that the file exists and that you specify the correct full pathname.

### **BUTIL–18: An error occurred during access of the sequential file.**

Ensure that the source file is valid.

### **BUTIL–19: BUTIL could not open the alternate collating sequence file.**

Ensure that you assigned a valid pathname to the Alternate Collating Sequence filename element in the description file.

### **BUTIL–20: An error occurred during access of the alternate collating sequence file.**

Ensure that the information in the alternate collating sequence file is formatted correctly.

### **BUTIL–21: The file version is earlier than 6.0.**

The RECOVER command cannot recover data from a Btrieve v5.x file.

### **BUTIL–23: The /D parameter specified to the Requester was too small for BUTIL to receive the entire record. BUTIL is writing only nn bytes.**

The Maintenance utility is writing only as many bytes as the value of the /D option allows. If you want the utility to write all the bytes in the record, specify a value for the /D option that is at least as large as the affected record.

### **BUTIL–25: The /D parameter specified to BUTIL was too small for BUTIL to receive any part of the record.**

You specified an invalid value for the Requester's /D option. Use the Setup utility to increase the value specified.

### **BUTIL–26: The data buffer is too small to hold any part of the record.**

The MicroKernel cannot return any data in the data buffer because the data buffer is too small to hold it. Use the Setup utility to increase the value specified for the Largest Record Size configuration option.

### **BUTIL–27: An error occurred during the access of the variable page. BUTIL is writing the obtainable portion of the variable page.**

An error occurred during the recovery of a file with variable-length records. The file is corrupt.

### **BUTIL–30: The key position cannot exceed the record length.**

The range of the specified key position is invalid. The key position you specify on a Btrieve call must be within the range of the record's length. For example, for a record that is 100 bytes long, a key position of 50 is within the correct range. However, a key position of 150 is not.

### **BUTIL–31: The key position plus key length cannot exceed the record length.**

The range of the key position you specified is invalid. The key position of a key plus its length cannot be larger than the record length. Verify that the key is defined so that its position plus its length does not exceed the record length.

### **BUTIL–32: The key length must be an even number for key type xxxx.**

You specified an invalid key length for the key type. Some key types must contain an even number of bytes. Specify a valid Key Length element.

### **BUTIL–36: The page size must be a multiple of 512, from 512 to 4,096.**

The page size you specified is not a multiple of 512, from 512 to 4096. Specify an appropriate page size.

### **BUTIL–37: The record length cannot exceed the page size.**

The record length you specified is invalid. In the description file, the record length you specified for the Record Length element is larger than the page size you specified for the Page Size element. Specify a record length that is smaller than the page size or increase the page size.

### **BUTIL–38: The record length must be at least 4 and no greater than 4,096.**

Specify a record length between 4 and 4096 (inclusive) for Btrieve v5.x, or between 4 and 4088 for Btrieve v6.x.

### **BUTIL–41: The alternate collating sequence cannot be found.**

The Maintenance utility cannot find the alternate collating sequence file you specified in the definition file. Verify that the alternate collating sequence file exists and that the name is correct in the definition file.

### **BUTIL–43: The file exists, but the Replace option was not specified.**

The Maintenance utility did not create a file when you specified the BUTIL –CREATE command because the file already exists. To recreate this file, specify the Replace Existing File element in the description file as Y.

### **BUTIL–44: The file access error nn occurred for file filename.**

The Maintenance utility returns the appropriate status code and filename for a file on which a file access error occurred during the beginning or end of continuous operation. The corrective measure depends on the status code received. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **BUTIL–45: The number of duplicate keys must be between 1 and 119.**

Check the value specified for the Duplicate Key element in the description file.

### **BUTIL–47: BUTIL cannot open the command file.**

Ensure that the command file exists and that you specified the command file location and filename correctly.

### **BUTIL–48: The command file is empty.**

Specify the desired commands in the command file before attempting to use the command file again. In addition, ensure you specified the correct command filename.

### **BUTIL–49: The command file exceeds 1,000 bytes.**

A command file cannot contain more than 1,000 bytes. Verify that the command file adheres to this requirement.

### **BUTIL–50: An internal error caused BUTIL to terminate.**

The Maintenance utility detected an internal diagnostic error that caused it to terminate.

### **BUTIL–52: Btrieve cannot be stopped when Scalable SQL is loaded.**

Unload Scalable SQL before attempting to unload the MicroKernel.

### **BUTIL–53: Btrieve error nn occurred for file or command xxxx.**

The Maintenance utility returns a status code related to a particular file or command. The corrective measure depends on the status code received. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **BUTIL–60: The end of the file occurred while BUTIL was expecting keyword xxxx on key segment descriptor nn.**

Check the syntax of the description file.

### **BUTIL–61: The end of the file occurred while BUTIL was expecting keyword xxxx.**

Check the syntax of the description file.

### **BUTIL–62: BUTIL was expecting keyword xxxx on key segment descriptor nn.**

Check the syntax of the description file.

### **BUTIL–63: BUTIL was expecting keyword xxxx.**

Check the syntax of the description file.

### **BUTIL–65: BUTIL has loaded no records.**

Verify that you specified the command correctly and that the input file is in the correct format.

### **BUTIL–70: The Btrieve error nn occurred on closing a file.**

The Maintenance utility returns this status code while closing a file. The corrective measure depends on the status code received. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **BUTIL–74: Btrieve error nn was returned for the Stop Command.**

This message applies only to the DOS environment. The Maintenance utility returns this status code after you issue the BUTIL –STOP command. The corrective measure depends on the status code received. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **BUTIL-76: When BUTIL wrote the Page Allocation Table at page #nn, an error occurred.**

The Maintenance utility returns this message while salvaging a file if the file is corrupt, or when a hardware error occurs.

### **BUTIL-77: When BUTIL wrote a mirror copy of the Page Allocation Table at page #nn, an error occurred.**

The Maintenance utility returns this message while salvaging a file if the file is corrupted, or when a hardware error occurs.

### **BUTIL-82: Error occurred on file specification of the description file.**

An error occurred when the Maintenance utility attempted to read the file specification portion of the description file. Ensure that the elements that describe the file specification are valid.

### **BUTIL-84: BUTIL internal error.**

The Maintenance utility encountered an internal error, such as an invalid pointer. Try to run the Maintenance utility again.

### **BUTIL-90: BUTIL could not allocate enough memory.**

Free some memory at the server by unloading unused applications.

### **BUTIL-91: BUTIL could not determine the size of the file.**

Try to recover the file using the BUTIL -RECOVER command.

### **BUTIL-101: The list of files that were not processed is: list.**

The Maintenance utility returns this error during the ROLLFWD operation if the utility encounters a file that cannot be rolled forward. The utility stops rolling forward and lists the remaining files not processed.

If you receive this message in conjunction with the Maintenance utility message number 103, restore the missing log segment and any log segments created after it to the log directory. If you cannot restore all missing log segments, restore as many as possible. Then, perform the ROLLFWD operation and specify the *segment\_number* parameter of the earliest log segment you could restore.

### **BUTIL-102: A Roll Forward error occurred.**

An internal error occurred while the MicroKernel was rolling forward files.

### **BUTIL-103: Log segment number nn is missing.**

The Maintenance utility returns this error during ROLLFWD operation if you do not specify a *segment\_number* parameter and one of the files to be rolled forward has a Restore Segment Number that indicates a log segment number that the MicroKernel cannot find.

### **BUTIL-131: BUTIL was unable to create or open the sequential file.**

The Maintenance utility returns this message when it is unable to create or open the specified file. Check the

sequential file to make sure it exists and has the read-only attribute set.

**BUTIL–132: The disk volume is full.**

You must have more disk space to create or enlarge any data files.

**BUTIL–134: BUTIL was unable to create or open the new file.**

Check the file specified for the BUTIL –SAVE or –RECOVER command. The file may already exist.

**BUTIL–136: BUTIL was unable to write the new backup file.**

Verify that you specified the correct path and filename for the backup file. Also, ensure you have enough disk space for the file to be written.

**BUTIL–141: File xx has an invalid extension. This extension is reserved for extended unformatted files.**

Specify a valid extension. For more information, see the *Pervasive.SQL User's Guide*.

**BUTIL-142: File xx exists.**

For more information about extended files, see the *Pervasive.SQL User's Guide*.

**BUTIL-143: File xx exists. This file is a potential extension of file yy. Remove this file and other potential extension files in this directory.**

For more information about extended files, see the *Pervasive.SQL User's Guide*.

**BUTIL-144: File xx has reached the size limit. The utility created one or more extension files. The last extension file is yy.**

For more information about extended files, see the *Pervasive.SQL User's Guide*.

**BUTIL-145: The utility found file xx. Remove this file before loading data from file yy.**

For more information about extended files, see the *Pervasive.SQL User's Guide*.

**BUTIL–152: There was an error opening file filename.**

Check the data file attributes, path, and filename.

**BUTIL–155: BUTIL cannot open the Btrieve file filename.**

Check the path, filename, and file attributes.

## MKDE

### MicroKernel Database Engine Messages

This section lists the messages that the MicroKernel Database Engine generates.

#### **MKDE–12: The value specified for the –option option is invalid.**

Enter the correct value in the BSTART.NCF, BTI.INI, or BTI.CFG file.

#### **MKDE–13: The option specified is not a valid option.**

Remove the invalid option from the BSTART.NCF, BTI.INI, or BTI.CFG file.

#### **MKDE–16: There is insufficient memory to load the MicroKernel Database Engine.**

The system has insufficient memory to allow the MicroKernel to load as it is configured. Use the Setup utility to reconfigure the MicroKernel to use less memory, or unload any unnecessary applications.

#### **MKDE–20: The log file filename cannot be written. Check disk space.**

If the disk is full, free some space by deleting any unnecessary files. When an error occurs while writing to the transaction log file, the MicroKernel attempts to resume transaction logging after every update to a data file containing system data. Therefore, when the condition causing the error is corrected (for example, when more disk space is made available), the MicroKernel automatically resumes transaction logging.

#### **MKDE–25: The file filename is rolling back.**

Rolling back is the procedure the MicroKernel uses to restore a file to a consistent state after a system crash. This procedure undoes the most recent changes to the file (changes in the most recent system transaction).

#### **MKDE–51: Program initialization failed; the program was not loaded.**

This is an informative message that the MicroKernel returns when an error prevents the MicroKernel from loading. Another message accompanies this one and provides more specific information about the nature of the error.

#### **MKDE–57: The file filename cannot be created in the log segment directory xxxx.**

While initializing, the MicroKernel could not find the specified transaction log segment directory, so it did not load. Ensure that the directory exists and that the MicroKernel has rights to it. (This is not a problem for the NetWare NLM version of the MicroKernel, because it has rights to all directories and files.) Use the Setup utility to specify the transaction log directory.

#### **MKDE–58: The file filename cannot be opened in the log segment directory xxxx.**

While initializing, the MicroKernel could not open the specified file in the transaction log segment directory, so it did not load. Ensure that the directory and file exist and that the MicroKernel has rights to it. (This is not a problem for the NetWare NLM version of the MicroKernel, because it has rights to all directories and files.) Another possibility is that someone has opened the file with an exclusive lock. Use the Monitor utility to see if another user has locked the file.

### **MKDE–59: The log segment file filename cannot be created.**

The MicroKernel could not create the specified transaction log segment file. Ensure that you have enough disk space and that the MicroKernel has rights to the file. (This is not a problem for the NetWare NLM version of the MicroKernel, because it has rights to all directories and files.) Another possibility is that someone has opened the file with an exclusive lock. Use the Monitor utility to see if another user has locked the file.

This error can occur while the MicroKernel is loading, or later, during normal operations. If it occurs while the MicroKernel is loading, the loading fails and no operations are possible until the problem is corrected. If the error occurs later, during normal operations, there are some additional issues to consider.

If the MicroKernel is already loaded when an error occurs creating the transaction log file, it ceases transaction logging; however, the MicroKernel attempts to resume transaction logging after every update to a data file containing system data. Therefore, when the condition causing the error is corrected (for example, when more disk space is made available), the MicroKernel automatically resumes transaction logging.

### **MKDE–75: The log file format is invalid.**

The MicroKernel returns this message when the log segment file format is incompatible. Ensure that the version of the MicroKernel doing the roll forward is the same as the version of the MicroKernel that created the log file.

### **MKDE–76: An Open operation failed on “xxxx” with Btrieve error nn.**

The MicroKernel returns this message during roll forward operations when a Btrieve error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–77: An unexpected Btrieve error nn occurred on INSERT to “xxxx”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Insert error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–78: An unexpected Btrieve error nn occurred on UPDATE in “xxxx”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Update error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–79: An unexpected Btrieve error nn occurred on DELETE from “xxxx”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Delete error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–80: The system is out of memory.**

The MicroKernel returns this message when the server has insufficient memory to allow the program to operate. Unload any unnecessary applications and retry the roll forward operation.

### **MKDE–81: An internal error nn occurred.**

The MicroKernel returns this message when an internal error occurs. Try to perform the operation again.

## **MKDE–82: An unexpected Btrieve error nn occurred on GET EQUAL in “xxxx”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Get Equal error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

## **MKDE–83: The log file is corrupt.**

The MicroKernel returns this message when the log segment file is damaged and unreadable. The Log Segment Manager cannot perform roll forward operations using this log segment.

## **MKDE–84: An error occurred while reading the log file.**

The MicroKernel returns this message when an error prevents the program from reading the log segment file. The MicroKernel cannot perform roll forward operations using this log segment.

## **MKDE–85: An unexpected Btrieve error nn occurred on CREATE INDEX to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a CREATE INDEX operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

## **MKDE–86: An unexpected Btrieve error nn occurred on DROP INDEX to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a DROP INDEX operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

## **MKDE–87: An unexpected Btrieve error nn occurred on BEGIN TRANSACTION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a BEGIN TRANSACTION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

## **MKDE–88: An unexpected Btrieve error nn occurred on END TRANSACTION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on an END TRANSACTION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

## **MKDE–89: An unexpected Btrieve error nn occurred on ABORT TRANSACTION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on an ABORT TRANSACTION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

### **MKDE–90: An unexpected Btrieve error nn occurred on CLOSE to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a CLOSE operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

### **MKDE–91: An unexpected Btrieve error nn occurred on GET POSITION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a GET POSITION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

### **MKDE–92: An unexpected Btrieve error nn occurred on BEGIN TRANSACTION.**

The MicroKernel returns this message during roll forward operations when an error occurs on a BEGIN TRANSACTION operation. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

### **MKDE–93: An unexpected Btrieve error nn occurred on END TRANSACTION.**

The MicroKernel returns this message during roll forward operations when an error occurs on an END TRANSACTION operation. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

### **MKDE–94: An unexpected Btrieve error nn occurred on ABORT TRANSACTION.**

The MicroKernel returns this message during roll forward operations when an error occurs on an ABORT TRANSACTION operation. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

### **MKDE–1001: The value specified for Cache Allocation is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 32 to the client machine’s available memory.

### **MKDE–1002: The value specified for Longest Compressed Record Size is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 0 to 64.

### **MKDE–1003: The value specified for Extended Operation Buffer Size is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 0 to 64.

### **MKDE–1004: The value specified for Open Files is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 1 to 64,000.

**MKDE–1005: The value specified for Operation Bundle Limit is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 0 to 64,000.

**MKDE–1007: The value specified for Number of Handles is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 1 to 64,000.

**MKDE–1008: The value specified for I/O Threads is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 0 to 64.

**MKDE–1009: The value specified for Maximum Sort Buffer is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 0 to one half of the client machine's available memory.

**MKDE–1010: The path specified for Temporary Sorting Files is invalid.**

Specify a valid path. In OS/2 the MicroKernel accepts any valid directory path.

**MKDE–1011: Info: The Number of Handles cannot be less than the Number of Open Files.**

Either specify a higher value for the Handles or specify a lower value for the Open Files option.

**MKDE–1015: The value specified for Maximum Record Size is invalid.**

Specify a valid value for this option. For more information about valid values, refer to the *Pervasive.SQL User's Guide*.

**MKDE–1017: The value specified for System Transaction Hold Limit is invalid.**

Specify a valid value for this option. For more information about valid values, refer to the *Pervasive.SQL User's Guide*.

**MKDE–1019: The value specified for Maximum Active Clients is invalid.**

Specify a valid value for this option. For more information about valid values, refer to the *Pervasive.SQL User's Guide*.

**MKDE–1020: The value specified for Worker Threads is invalid.**

Specify a valid value for this option. For more information about valid values, refer to the *Pervasive.SQL User's Guide*.

**MKDE–1021: The value specified for Thread Priority Delta is invalid.**

Specify a valid value. For more information about valid values, refer to the *Pervasive.SQL User's Guide*.

**MKDE–1022: The value specified for Initiation Time Limit is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 0 to 120,000.

**MKDE–1023: The value specified for Disk I/O Wait Time is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 0 to 10,000.

**MKDE–1024: The value specified for Page Write Group Size is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 1 to 500.

**MKDE–1025: The value specified for Auto-Terminate Delay is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 0 to 120,000.

**MKDE–1026: The value specified for Number of Lines Allocated is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 1 to 256.

**MKDE–1027: The value specified for Message Log File Limit is invalid.**

Specify a valid value. In OS/2 the MicroKernel has a range of acceptable values from 1 to 4,096.

**MKDE–1028: The path specified for Home Directory is invalid.**

Specify a valid path. In OS/2 the MicroKernel accepts any valid path.

**MKDE–1030: The value specified for the *option\_name* option is invalid.**

The MicroKernel Database Engine returns this message when the value specified for the *option\_name* option is invalid. Specify a valid value.

**MKDE–1032: The option specified is not a valid option.**

The MicroKernel Database Engine returns this message when the option you specified is not valid. Specify a valid configuration option.

**MKDE–1152: Btrieve cannot get sufficient memory to complete the operation.**

The computer has insufficient memory to allow the MicroKernel to load as it is configured. Configure the MicroKernel to use less memory or unload any unnecessary applications.

**MKDE–1156: The log file filename cannot be written. Check disk space.**

The transaction log file *filename* cannot be written. Check the disk space. If the disk is full, free some space by deleting any unnecessary files. When an error occurs while writing to the transaction log file, the MicroKernel attempts to resume transaction logging after every update to a data file containing system data. Therefore, when the condition causing the error is corrected (for example, when more disk space is made available), the MicroKernel automatically resumes transaction logging.

**MKDE–1161: The file filename is rolling back.**

Rolling back is the procedure the MicroKernel uses to restore a file to a consistent state after a system crash. This procedure undoes the most recent changes to the file (changes in the most recent system transaction).

### **MKDE–1164: Error creating/accessing semaphore.**

The MicroKernel returns this message if an internal error occurs.

### **MKDE–1165: Error creating execution thread.**

Decrease the values for either the Number of Worker Threads or the I/O Threads configuration option.

### **MKDE–1166: Could not open Pervasive Event Log file. This file is created in the home directory. Verify that the directory is present and sufficient space is available.**

The MicroKernel returns this message when it cannot create the Pervasive Event Log (PVSU.LOG) in the directory. Make sure that the disk on your computer is not full.

The home directory is defined as \WINNT\SYSTEM32 (Windows NT), \WINDOWS\SYSTEM (Windows 95) or SYS:\SYSTEM (NetWare server).

### **MKDE–1167: Could not create IPC queue.**

The MicroKernel returns this message if an internal error occurs.

### **MKDE–1168: Clients(s) active - Shutdown anyway?**

The MicroKernel returns this message when you attempt to shut down the MicroKernel while at least one application is still using the MicroKernel. If you choose to shut down anyway, you could lose data.

### **MKDE–1169: Could not open BTI.INI.**

The MicroKernel returns this message when it cannot open the BTI.INI file. Make sure that the file is in the home directory.

### **MKDE–1171: The file BTI.INI is flagged as read only. Btrieve can temporarily override this attribute to save the current settings. Do you wish to continue?**

The MicroKernel returns this message if the BTI.INI file is flagged as read only and you attempt to save changes to the Btrieve configuration settings, which are stored in the BTI.INI file. If you choose to continue, the MicroKernel updates the BTI.INI file but maintains its read only status. If you choose not to continue, the MicroKernel does not update the BTI.INI file. You can still change the configuration settings, but they are only in effect for the current Btrieve session.

### **MKDE–1172: Error writing to BTI.INI.**

Make sure that your computer's disk is not full and that you have the appropriate rights to the file.

### **MKDE–1175: I/O error - filename.**

The MicroKernel returns this message when it encounters an I/O error while reading from or writing to a file. This

message typically accompanies Status Code 2.

### **MKDE–1176: Variable page error - filename.**

The MicroKernel returns this message when it encounters an error on a variable page. This message typically accompanies Status Code 2.

### **MKDE–1177: Help is not available.**

The MicroKernel returns this message when it cannot find the Btrieve help file (BTRIEVE.HLP). Make sure that the file is located in the home directory.

### **MKDE–1178: Resources allocated.**

This is an informative message from the MicroKernel Database Engine. No action is required.

### **MKDE–1179: Resources released.**

This is an informative message from the MicroKernel Database Engine. No action is required.

### **MKDE–1180: System Error: nn.nn.nn.**

The MicroKernel Database Engine returns this message when a system error occurs.

### **MKDE–1181: The message log file exceeds specified limit. Do you wish to truncate the file?**

This message only applies to Btrieve for OS/2 v6.15.

The MicroKernel returns this message when the message log file (MKDEMSG.LOG) becomes larger than the size specified by the MicroKernel configuration. If you choose to truncate the file, the MicroKernel reduces the file to the size specified by the configuration and keeps the most recent messages in the file. If you choose not to truncate, the MicroKernel allows the file to grow beyond the limit specified by the Message Log File Size Limit configuration option.

### **MKDE–1183: Not enough available cache to complete an operation.**

Configure the MicroKernel to use more cache memory or reduce the file's page size.

### **MKDE–1184: Unable to set number of file handles to requested value. Maximum open files reduced to nn.**

The MicroKernel is unable to set the Maximum Open Files to the value you specified; instead the MicroKernel uses the value indicated in the message. To avoid receiving this message, run fewer applications at the same time.

### **MKDE–1192: Settings read from NOVDB.INI - To change to BTI.INI, save current settings.**

The MicroKernel returns this message when it loads and finds a NOVDB.INI file where it expects to find the BTI.INI file. To use the BTI.INI file, open the Settings notebook and click the **Save** button.

### **MKDE–1193: Btrieve is using default settings.**

The MicroKernel returns this message when it loads and cannot find values for one or more configuration options. The BTI.INI file is incomplete or missing. You may update the configuration options using the Settings notebook.

### **MKDE–1194: Additional worker thread spawned.**

The MicroKernel returns this message when it must create more worker threads than the number specified by the Number of Worker Threads configuration option. To avoid receiving this message, increase the value for the Number of Worker Threads option.

### **MKDE–1207: The log file format is invalid.**

The MicroKernel returns this message when the log segment file format is incompatible. Ensure that the version of the MicroKernel doing the roll forward is the same as the version of the MicroKernel that created the log file.

### **MKDE–1208: An Open operation failed on “xxxx” with Btrieve error nn.**

The MicroKernel returns this message during roll forward operations when a Btrieve error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–1209: An unexpected Btrieve error nn occurred on INSERT to “xxxx”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Insert error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–1210: An unexpected Btrieve error nn occurred on UPDATE in “xxxx”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Update error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–1211: An unexpected Btrieve error nn occurred on DELETE from “xxxx”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Delete error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–1212: The system is out of memory.**

The MicroKernel returns this message when the server has insufficient memory to allow the program to operate. Unload any unnecessary applications and retry the roll forward operation.

### **MKDE–1213: An internal error nn occurred.**

The MicroKernel returns this message when an internal error occurs. Try to perform the operation again.

### **MKDE–1214: An unexpected Btrieve error nn occurred on GET EQUAL in “xxxx”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Get Equal error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, "Status Codes"](#) for a description of Status Code *nn*.

### **MKDE–1215: The log file is corrupt.**

The MicroKernel returns this message when the log segment file is damaged and unreadable. The Log Segment Manager cannot perform roll forward operations using this log segment.

### **MKDE–1216: An error occurred while reading the log file.**

The MicroKernel returns this message when an error prevents the program from reading the log segment file. The MicroKernel cannot perform roll forward operations using this log segment.

### **MKDE–1217: An unexpected Btrieve error *nn* occurred on CREATE INDEX to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a CREATE INDEX operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–1218: An unexpected Btrieve error *nn* occurred on DROP INDEX to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a DROP INDEX operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–1219: An unexpected Btrieve error *nn* occurred on BEGIN TRANSACTION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a BEGIN TRANSACTION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–1220: An unexpected Btrieve error *nn* occurred on END TRANSACTION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on an END TRANSACTION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–1221: An unexpected Btrieve error *nn* occurred on ABORT TRANSACTION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on an ABORT TRANSACTION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–1222: An unexpected Btrieve error *nn* occurred on CLOSE to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a CLOSE operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–1223: An unexpected Btrieve error nn occurred on GET POSITION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a GET POSITION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–1224: An unexpected Btrieve error nn occurred on BEGIN TRANSACTION.**

The MicroKernel returns this message during roll forward operations when an error occurs on a BEGIN TRANSACTION operation. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–1225: An unexpected Btrieve error nn occurred on END TRANSACTION.**

The MicroKernel returns this message during roll forward operations when an error occurs on an END TRANSACTION operation. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–1226: An unexpected Btrieve error nn occurred on ABORT TRANSACTION.**

The MicroKernel returns this message during roll forward operations when an error occurs on an ABORT TRANSACTION operation. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–2001: The value specified for the Cache Allocation option is invalid.**

The MicroKernel returns this message when the value specified for the Cache Allocation option is invalid. Use the Setup utility and specify a valid value for this option.

### **MKDE–2002: The value specified for the Largest Compressed Record Size option is invalid.**

The MicroKernel returns this message when the value specified for the Largest Compressed Record Size option is invalid. Use the Setup utility to specify a valid value for the Largest Compressed Record Size option.

### **MKDE–2004: The value specified for the Number of Open Files option is invalid.**

The MicroKernel returns this message when the value specified for the Number of Open Files option is invalid. Use the Setup utility and specify a valid value for this option.

### **MKDE–2005: The value specified for the Operation Bundle Limit option is invalid.**

The MicroKernel returns this message when the value specified for the Operation Bundle Limit option is invalid. Use the Setup utility and specify a valid value for this option.

**MKDE–2007: The value specified for the Number of Handles option is invalid.**

The MicroKernel returns this message when the value specified for the Handles option is invalid. Use the Setup utility and specify a valid value for this option.

**MKDE–2008: The value specified for the I/O Threads option is invalid.**

The MicroKernel returns this message when the value specified for the number of I/O threads option is invalid. Use the Setup utility and specify a valid value for this option.

**MKDE–2009: The value specified for the Active Clients option is invalid.**

The MicroKernel returns this message when the value specified for the Active Clients option is invalid. Use the Setup utility and specify a valid value for this option.

**MKDE–2010: The value specified for the Sort Buffer Size option is invalid.**

The MicroKernel returns this message when the value specified for the Sort Buffer Size option is invalid. Use the Setup utility and specify a valid value for this option.

**MKDE–2011: The path specified for the Temporary Directory option is invalid.**

The MicroKernel returns this message when the value specified for the Temporary Work Directory option is invalid. Use the Setup utility and specify a valid path.

**MKDE–2012: The value specified for the –option option is invalid.**

The MicroKernel returns this message when the value for an option is not valid. Return to the Setup utility and enter the correct value.

**MKDE–2013: The option specified is not a valid option.**

The MicroKernel Database Engine returns this message when the option specified is not a valid option. Return to the Setup utility and remove the invalid option.

**MKDE–2015: MKDE cannot get sufficient memory to complete the operation.**

The MicroKernel returns this message when the server has insufficient memory to allow the MicroKernel to load as it is configured. Use the Setup utility to reconfigure the MicroKernel to use less memory, or unload any unnecessary applications.

**MKDE–2019: The log file filename cannot be written. Check disk space.**

The MicroKernel returns this message when log file *filename* cannot be written. Check the disk space. If the disk is full, free some space by deleting any unnecessary files. When an error occurs while writing to the transaction log file, the MicroKernel attempts to resume transaction logging after every update to a data file containing system data. Therefore, when the condition causing the error is corrected (for example, when more disk space is made available), the MicroKernel automatically resumes transaction logging.

## **MKDE–2024: The file filename is rolling back.**

The MicroKernel returns this message when the file *filename* is rolling back. Rolling back is the procedure the MicroKernel uses to restore a file to a consistent state after a system crash. This procedure undoes the most recent changes to the file (changes in the most recent system transaction).

## **MKDE–2027: Error creating/accessing semaphore.**

The MicroKernel returns this message if an internal error occurs.

## **MKDE–2028: Error creating execution thread.**

The MicroKernel returns this message when it cannot create additional threads. Decrease the values for either the Worker Threads or the I/O Threads option.

## **MKDE–2029: Could not open Event Log file.**

The MicroKernel returns this message when it cannot create the Pervasive Event Log (PVSU.LOG) in the home directory. Make sure that the home directory exists and that the disk on your computer is not full. Also, ensure that the MicroKernel has rights to the home directory and event log file.

The Home directory is defined as \WINNT\SYSTEM32 (Windows NT), \WINDOWS\SYSTEM (Windows 95), and SYS:\SYSTEM (NetWare servers).

## **MKDE–2039: System Error: nn.nn.nn.**

The MicroKernel returns this message when a system error occurs.

## **MKDE–2053: Additional worker thread spawned.**

This is an informative message the MicroKernel returns when it must create more worker threads than the number specified by the Number of Worker Threads configuration option. To avoid receiving this message, increase the value for the Number of Worker Threads option.

## **MKDE–2055: Error creating named pipe.**

The MicroKernel returns this message when it is unable to create a named pipe for the pipe-manager thread.

## **MKDE–2060: Error reading the registry. Error code: nn.**

The MicroKernel returns this message when it is unable to find a predefined setting in the registry. Use the –REGINSTALL startup parameter next time you start the MicroKernel. This parameter forces the MicroKernel to recreate its section in the registry with the default values for each configuration option.

## **MKDE–2061: Error writing the registry. Error code: nn.**

The MicroKernel returns this message when it encounters an error while trying to write into the registry. Use the –REGINSTALL startup parameter next time you start the MicroKernel. This parameter forces the MicroKernel to recreate its section in the registry with the default values for each configuration option.

## **MKDE–2065: Error calling Service Control dispatcher.**

The MicroKernel returns this message when it encounters an error during the service initialization.

### **MKDE–2067: There is no suitable protocol available on this computer.**

The MicroKernel returns this message when none of the supported protocols are installed on the computer. The Communications Server cannot initialize.

### **MKDE–2068: The communications server is not available.**

The MicroKernel returns this message when it encounters an error loading the Communications Server. Make sure the system path includes the directory where the Communications Server that was shipped with this version of the MicroKernel.

### **MKDE–2075: The transaction log file LAST\_SEG.LOG cannot be opened.**

The MicroKernel returns this message when the log file format is invalid or the Transaction Log Directory does not exist. Ensure that the version of the MicroKernel that returns this message is the same as the version of the MicroKernel that created the log file.

### **MKDE–2076: An Open operation failed on “xxxx” with Btrieve error nn.**

The MicroKernel returns this message during roll forward operations when a Btrieve error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–2077: The log file format is invalid.**

The MicroKernel returns this message when the log segment file format is incompatible. Ensure that the version of the MicroKernel doing the roll forward is the same as the version of the MicroKernel that created the log file.

### **MKDE–2078: An Open operation failed on “xxxx” with Btrieve error nn.**

The MicroKernel returns this message during roll forward operations when a Btrieve error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–2079: An unexpected Btrieve error nn occurred on INSERT to “xxxx”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Insert error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–2080: An unexpected Btrieve error nn occurred on UPDATE in “xxxx”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Update error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–2081: An unexpected Btrieve error nn occurred on DELETE from “xxxx”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Delete error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a

description of Status Code *nn*.

### **MKDE–2082: The system is out of memory.**

The MicroKernel returns this message when the server has insufficient memory to allow the program to operate. Unload any unnecessary applications and retry the roll forward operation.

### **MKDE–2083: An internal error *nn* occurred.**

The MicroKernel returns this message when an internal error occurs. Try to perform the operation again.

### **MKDE–2084: An unexpected Btrieve error *nn* occurred on GET EQUAL in “XXXX”.**

The MicroKernel returns this message during roll forward operations when a Btrieve Get Equal error occurs. The MicroKernel returns the specified status code for the specified file. Refer to [Chapter 1, “Status Codes”](#) for a description of Status Code *nn*.

### **MKDE–2085: The log file is corrupt.**

The MicroKernel returns this message when the log segment file is damaged and unreadable. The Log Segment Manager cannot perform roll forward operations using this log segment.

### **MKDE–2086: An error occurred while reading the log file.**

The MicroKernel returns this message when an error prevents the program from reading the log segment file. The MicroKernel cannot perform roll forward operations using this log segment.

### **MKDE–2087: An unexpected Btrieve error *nn* occurred on CREATE INDEX to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a CREATE INDEX operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

### **MKDE–2088: An unexpected Btrieve error *nn* occurred on DROP INDEX to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a DROP INDEX operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

### **MKDE–2089: An unexpected Btrieve error *nn* occurred on BEGIN TRANSACTION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a BEGIN TRANSACTION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, “Status Codes”](#).

### **MKDE–2090: An unexpected Btrieve error *nn* occurred on END TRANSACTION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on an END TRANSACTION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–2091: An unexpected Btrieve error nn occurred on ABORT TRANSACTION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on an ABORT TRANSACTION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–2092: An unexpected Btrieve error nn occurred on CLOSE to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a CLOSE operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–2093: An unexpected Btrieve error nn occurred on GET POSITION to filename.**

The MicroKernel returns this message during roll forward operations when an error occurs on a GET POSITION operation for the specified file. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–2094: An unexpected Btrieve error nn occurred on BEGIN TRANSACTION.**

The MicroKernel returns this message during roll forward operations when an error occurs on a BEGIN TRANSACTION operation. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–2095: An unexpected Btrieve error nn occurred on END TRANSACTION.**

The MicroKernel returns this message during roll forward operations when an error occurs on an END TRANSACTION operation. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

### **MKDE–2096: An unexpected Btrieve error nn occurred on ABORT TRANSACTION.**

The MicroKernel returns this message during roll forward operations when an error occurs on an ABORT TRANSACTION operation. You may have a corrupted file, or your log file may be corrupt. For more information about Status Code *nn*, refer to [Chapter 1, "Status Codes"](#).

## NWBSRVCM

# NetWare Btrieve Server Communications Manager Messages

### **NWBSRVCM-4: Btrieve Communication Manager is using default settings.**

The Btrieve Communications Manager displays this message when it cannot find the configuration file (BTI.CFG). It will use the default configuration settings.

### **NWBSRVCM-5: Btrieve Communication Manager detected an invalid command line option.**

Specify a valid option.

### **NWBSRVCM-6: Btrieve Communication Manager terminating - could not initialize properly.**

The Btrieve Communications Manager could not load the TCP/IP or SPX communication module. Verify that the configuration file (BTI.CFG) entry for Supported Protocols is correct.

### **NWBSRVCM-7: Btrieve Communication Manager terminating - insufficient memory to execute.**

An error occurred during a memory allocation operation. Free memory by unloading NLMs or reconfiguring NLMs to use less memory. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

### **NWBSRVCM-8: Btrieve Communication Manager could not load/initialize BTCPCOM.NLM.**

An error occurred during the loading or initialization of BTCPCOM.NLM. The operating system error number is displayed as errno. Refer to your operating system documentation for the definition of the errno value.

### **NWBSRVCM-9: Btrieve Communication Manager could not load/initialize BSPXCOM.NLM.**

An error occurred during the loading or initialization of BSPXCOM.NLM. The operating system error number is displayed as errno. Refer to your operating system documentation for the definition of the errno value.

### **NWBSRVCM-10: Btrieve Communication Manager terminating - invalid Supported Protocols specified in configuration file (BTI.CFG).**

Verify that the configuration file (BTI.CFG) entry for Supported Protocols is correct.

## NWSSRVCM

# NetWare Scalable SQL Server Communications Manager Messages

### **NWSSRVCM-4: Scalable SQL Communication Manager is using default settings.**

The Scalable SQL Communications Manager displays this message when it cannot find the configuration file (BTI.CFG). It will use the default configuration settings.

### **NWSSRVCM-5: Scalable SQL Communication Manager detected an invalid command line option.**

Specify a valid option.

### **NWSSRVCM-6: Scalable SQL Communication Manager terminating - could not initialize properly.**

The Scalable SQL Communications Manager could not load the TCP/IP or SPX communication module. Verify that the configuration file (BTI.CFG) entry for Supported Protocols is correct.

### **NWSSRVCM-7: Scalable SQL Communication Manager terminating - insufficient memory to execute.**

An error occurred during a memory allocation operation. Free memory by unloading NLMs or reconfiguring NLMs to use less memory. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

### **NWSSRVCM-8: Scalable SQL Communication Manager could not load/initialize STPCOM.NLM.**

An error occurred during the loading or initialization of BTPCOM.NLM. The operating system error number is displayed as errno. Refer to your operating system documentation for the definition of the errno value.

### **NWSSRVCM-9: Scalable SQL Communication Manager could not load/initialize SSPXCOM.NLM.**

An error occurred during the loading or initialization of BSPXCOM.NLM. The operating system error number is displayed as errno. Refer to your operating system documentation for the definition of the errno value.

### **NWSSRVCM-10: Scalable SQL Communication Manager terminating - invalid Supported Protocols specified in configuration file (BTI.CFG).**

Verify that the configuration file (BTI.CFG) entry for Supported Protocols is correct.

## SSPXCOM

### Scalable SQL SPX Communications Module Messages

The following messages are specific to the Scalable SQL Communications Manager.

#### **SSPXCOM-2: The server has insufficient memory to execute SSPXCOM.NLM**

An error occurred during a memory allocation operation. Free memory by unloading NLMs or reconfiguring NLMs to use less memory. Retry the operation. If the error persists, contact Pervasive Software Customer Support.

#### **SSPXCOM-3: An internal error has occurred. SSPXCOM detected a semaphore allocation failure.**

An internal error occurred during a semaphore allocation operation. Retry the operation. If the error persists, contact Pervasive Software Customer Support.

#### **SSPXCOM-4: The Service Request Block (SRB) function code nn contains invalid data. Check for an incompatible version of the file SSPXCOM.NLM.**

Ensure that SSPXCOM's version is compatible with the version number of the workstation's Scalable SQL Requester.

#### **SSPXCOM-6: Another NLM is using the socket number reserved for SSPXCOM.NLM.**

To solve this condition, unload all other NLMs. Load MKDE.NLM and NWSSRVCM.NLM, and then reload the other NLMs. This process reveals the NLM that is using SSPXCOM.NLM's socket number.

#### **SSPXCOM-7: An SPX receive I/O error (nn) has occurred. The connection has been lost.**

An error occurred during an SPX-level receive operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **SSPXCOM-8: An SPX send I/O error (nn) has occurred. The connection has been lost.**

An error occurred during an SPX-level send operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **SSPXCOM-10: Bad connection ID detected on receive. The SPX connection was lost after the initial request began.**

If this message appears frequently, increase your requester's SPX Timeout parameter, which is controlled with the Configuration (Setup) utility. Also, check for NLMs that are monopolizing the CPU time.

#### **SSPXCOM-11: Bad connection ID detected on send. The SPX connection was**

## **lost after the initial request began.**

If this message appears frequently, increase your requester's SPX Timeout parameter, which is controlled with the Configuration (Setup) utility. Also, check for NLMs that are monopolizing the CPU time.

## **SSPXCOM-12: An error (nn) was detected while trying to establish an SPX session requested by a remote workstation.**

An error occurred during an SPX-level session establishment operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

## **SSPXCOM-13: The session was rejected because the session limit was reached. Increase the value specified for the Number of Remote Sessions option.**

Increase the value specified for the Number of Remote Sessions configuration option, as described in the *Pervasive User's Guide*.

## **SSPXCOM-14: The request for statistics from the Monitor utility was not recognized. Check for an incompatible version of the utility or SSPXCOM.NLM.**

Ensure that your SSPXCOM.NLM version is compatible with the version number of the Monitor utility.

## **SSPXCOM-15: An internal error has occurred. BSPXCOM did not recognize the GET\_EIM\_STATS function.**

Ensure that the version of SSPXCOM.NLM is compatible with the version number of the Monitor utility.

## **SSPXCOM-16: Duplicate database name found on network. '<name>' will not be advertised.**

The database name is already being advertised on the network. Change the specified database name so that it is unique on the network.

## **SSPXCOM-19: A thread initialization error has occurred.**

An error occurred during BSPXCOM.NLM thread initialization. Refer to the Pervasive Software Event Log (PVSU.LOG) for error messages preceding this one. Correct the problem and retry the operation.

## **SSPXCOM-20: A BeginThread() error (nn) has occurred.**

An error occurred during a BeginThread operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

# SSQL

## Scalable SQL Engine Messages

This section lists the messages that the Scalable SQL engine generates.

### **SSQL-0: An error occurred while filename was being loaded. Insufficient memory is available, or the executable file is corrupt.**

Under Windows, the operating system does not have enough memory to load the specified DLL for the Scalable SQL for Windows engine, or the DLL file on the disk is corrupt. For insufficient memory, try using a larger swap file or adding more memory. For a corrupt file, reinstall the engine files.

### **SSQL-1: An error occurred while filename was being loaded. The file was not found.**

The Scalable SQL for Windows engine could not load the specified DLL or EXE file because the file or one of its dependent DLLs was not in the path. Ensure that the file or DLL is in the path.

### **SSQL-2: An error occurred while filename was being loaded. The path was not found.**

The Scalable SQL for Windows engine could not load the specified file because the file or one of its dependent DLLs was not in the path. Ensure that the file or DLL is in the path.

### **SSQL-3: An error occurred while filename was being loaded. A file sharing error occurred.**

The Scalable SQL for Windows engine encountered a file sharing error while trying to load the specified file. Try reloading the file.

### **SSQL-4: An error occurred while filename was being loaded. Error 6.**

The specified file is corrupt. Reinstall the specified file.

### **SSQL-5: An error occurred while filename was being loaded. Insufficient memory is available to load the DLL.**

Windows does not have enough memory to load the DLL for the Scalable SQL for Windows engine. Try using a larger swap file or adding more memory.

### **SSQL-6: An error occurred while filename was being loaded. The incorrect MS Windows version is in use.**

The Scalable SQL for Windows engine requires Windows v3.1.

### **SSQL-7: An error occurred while filename was being loaded. The executable file is invalid.**

The Scalable SQL for Windows engine could not load the DLL or EXE file because the file on the disk is corrupt.

Reinstall the engine files.

**SSQL–8: An error occurred while filename was being loaded. The incorrect operating system is in use.**

The Scalable SQL for Windows engine requires Windows v3.1.

**SSQL–9: An error occurred while filename was being loaded. The system is designed for use with MS-DOS v4.0.**

The version of MS DOS must be v4.0 or later.

**SSQL–10: An error occurred while filename was being loaded. The type of executable is unknown.**

The Scalable SQL for Windows engine encountered an error while trying to load the specified file. Try the operation again.

**SSQL–11: An error occurred while filename was being loaded. The real-mode executable cannot be loaded.**

The Scalable SQL for Windows engine encountered an error while trying to load the specified file. Try the operation again.

**SSQL–12: An error occurred while filename was being loaded. Error 16.**

The specified file is corrupt. Reinstall the specified file.

**SSQL–13: An error occurred while filename was being loaded. The executable is compressed.**

The Scalable SQL for Windows engine encountered an error while trying to load the specified file. Ensure that the file is not compressed and retry the operation.

**SSQL–14: An error occurred while filename was being loaded. The DLL is invalid.**

The Scalable SQL for Windows engine could not load the specified file because the file on the disk is corrupt. Reinstall the engine files.

**SSQL–15: An error occurred while filename was being loaded. MS Windows 32-bit extensions are required.**

The Scalable SQL for Windows engine encountered an error while trying to load the specified file.

**SSQL–16: An error occurred while filename was being loaded. An unknown load error occurred.**

The Scalable SQL for Windows engine encountered an error while trying to load the specified file. Try loading the file again. If necessary, stop and restart the Scalable SQL for Windows engine and try to load the file.

## **SSQL–17: Scalable SQL cannot link to the needed module: filename.**

A Scalable SQL component tried to link to *filename* but *filename* does not export a needed symbol. A common problem is if Novell's NETAPI.DLL is being loaded instead of Microsoft's NETAPI.DLL. Ensure that you are using the latest version of *filename*.

## **SSQL–18: DefineEntryPoints failed.**

Scalable SQL v3.x has insufficient memory to complete process initialization. Try using a larger swap file or adding more memory.

## **SSQL–19: The file filename is not loaded.**

This message appears after other error messages to inform you that the file is not loaded.

## **SSQL–20: The process could not be initialized.**

Scalable SQL v3.x has insufficient memory to complete process initialization. Try using a larger swap file or adding more memory.

## **SSQL–21: The MicroKernel Database Engine is not loaded.**

The Scalable SQL engine could not load or access the MicroKernel. Ensure that the following MicroKernel components are in your path:

- For NetWare, NWMKDE.NLM, which must be loaded before loading Scalable SQL.
- For Windows NT, NTMKDE.DLL, which is loaded by Scalable SQL.

## **SSQL–23: Insufficient memory is available to load messages.**

The operating system does not have enough memory to load the messages from the resource DLL. Try using a larger swap file or adding more memory.

## **SSQL–24: This Scalable SQL engine must be run with the fully-licensed MicroKernel Database Engine!**

The Scalable SQL engine is designed to run only with the fully-licensed and distributable MicroKernel for the corresponding platform. Make sure this version of the MicroKernel is loaded. The Scalable SQL engine includes a copy of the correct MicroKernel. Purchasing the Scalable SQL engine gives you the right to copy and distribute both engines with an application.

## **SSQL–25: This Scalable SQL development engine must be run with the MicroKernel Database development engine!**

The Scalable SQL development engine is designed to run only with the MicroKernel development engine for the corresponding platform. Make sure the MicroKernel development engine is loaded. The Scalable SQL Developer Kit includes a copy of the Scalable SQL development engine and the MicroKernel development engine. These two engines are licensed to support five concurrent users in a development environment and are not licensed for distribution.

## **SSQL–26: Unexpected MicroKernel Database Engine status nn during engine initialization!**

Scalable SQL failed to initialize because the MicroKernel Database Engine returned an unexpected status code indicating that it cannot be used. Refer to [Chapter 1, "Status Codes"](#) for a description of Status Code *nn*.

### **SSQL–28: The trace file filename cannot be opened.**

Scalable SQL could not open the specified trace file. Ensure that the path and filename are valid. Also ensure that you have enough free disk space and that the file is not read-only.

### **SSQL–29: The Scalable SQL Engine cannot load configuration information.**

Scalable SQL returns this message for one of the following reasons:

- In Windows NT, the Registry key used to store settings cannot be created. Use RegEdit to determine if the following keys have been created. If they have not, try to create them. The keys are as follows:
  - HKEY\_LOCAL\_MACHINE, SOFTWARE, Pervasive Software, Scalable SQL Server Engine, Version 4.00, Settings
  - HKEY\_LOCAL\_MACHINE, SOFTWARE, Pervasive Software, Database Names, Version 4.00, Settings
- In NetWare and DOS, the BTI.CFG file cannot be opened or locked. Ensure that BTI.CFG is writable and that BTI.LCK is not present.

### **SSQL–31: An error occurred while filename was being loaded. xxxx.**

A fatal error occurred while the Database Service Manager was loading the *filename*. Refer to the detailed message (*xxxx*) to resolve this error.

### **SSQL–32: A Database Service Manager error occurred. xxxx.**

A fatal error has occurred in the Database Service Manager. A detailed description (*xxxx*) follows the error message.

### **SSQL–33: The xx network protocol failed to initialize. Error code: nn.**

Ensure that the *xx* protocol has been configured correctly. Once you receive this message, the *xxx* network protocol is no longer available. The following are possible error codes:

- 1 Wsock32.dll is not available or could not be loaded.
- 2 Too many instances. This is an internal error and should not occur.
- 3 Out of memory.
- 4 Invalid Instance ID. This is an internal error and should not occur.
- 5 Not initialized. This is an internal error and should not occur.
- 6 Protocol Unavailable. The Scalable SQL engine is configured to use a protocol that is not available on this system. Use the Setup Utility to check the setting for the Scalable SQL Communications Manager, Server

Communication Configuration and Supported Protocols. Available protocols are SPXII and TCP/IP.

- 7 Protocol not supported. The Scalable SQL engine is configured to use a protocol that is not available on this system. Use the Setup Utility to check the setting for the Scalable SQL Communications Manager, Server Communication Configuration and Supported Protocols. Available protocols are SPXII and TCP/IP.
- 8 Error opening socket. Windows NT system resources may be low. Try rebooting the system on which you are loading the Scalable SQL engine.
- 9 Error Creating an I/O port. Windows NT system resources may be low. Try rebooting the system on which you are loading the Scalable SQL engine.
- 10 Error Creating Thread. Windows NT system resources may be low. Try rebooting the system on which you are loading the Scalable SQL engine.
- 11 Error Initializing Protocol. Possible causes are:
  - Out of memory or error creating a thread. Windows NT system resources may be low. Try rebooting the system on which you are loading the Scalable SQL engine.
  - The Windows Sockets implementation may have detected that the network subsystem has failed. Verify that the network subsystem is operational.
  - This machine's name and address cannot be determined. Verify that your name services are configured correctly and are operational. Depending on the configuration of your computer, this could include DNS, WINS, DHCP, and a static HOSTS file.
- 12 Invalid Communications Buffer Size. Use the Setup Utility to check the setting for the Scalable SQL Communications Manager, Server Communication Configuration and Communications Buffer Size. It should not be 0.
- 13 Invalid Parameter. This is an internal error and should not occur.
- 14 Engine Callback not specified. This is an internal error and should not occur.
- 15 Error not listening. This is an internal error and should not occur.
- 16 Error not working. This is an internal error and should not occur.

- 17 Timeout on cleanup. This is an internal error and should not occur.
- 18 Insufficient Buffer. This is an internal error and should not occur.
- 19 Thread Initialization error. Windows NT system resources may be low.  
Try  
rebooting the system on which you are loading the Scalable SQL engine.
- 20 Initialization is not reentrant. This is an internal error and should not occur.
- 21 Too many threads. This is an internal error and should not occur.
- 22 Error Creating Semaphore. Windows NT system resources may be low.  
Try  
rebooting the system on which you are loading the Scalable SQL engine.
- 23 Error initializing protocol, retrying. This is an informative status that should cause Scalable SQL to display message SSQL-34.

### **SSQL-34: The xx network protocol failed to initialize. Error code: nn. Retrying...**

Ensure that the xx network protocol has been configured correctly. The Scalable SQL engine will continue to attempt initializing the xx network protocol. Refer to the text for SSQL-33 for a list of communications status codes.

### **SSQL-35: The Scalable SQL Engine is running with no network protocol support.**

All selected network protocols failed to initialize or no network protocols have been configured. Remote access to the Scalable SQL engine is unavailable. Ensure that the selected network protocols have been configured correctly or use the Setup utility to select a different protocol.

### **SSQL-36: The Scalable SQL Engine has not been configured to use network protocols.**

No network protocols have been selected to run with the Scalable SQL engine. Use the Setup utility to select one or more network protocols.

### **SSQL-37: The Communications Server is not available.**

The Scalable SQL engine failed to load the Communications Server. Remote access is disabled. Ensure that the selected network protocols have been configured correctly.

**SSQL–38: The filename component is version n.n and must be version y.y or later.**

The *filename* module is an older version. Ensure you have the correct version of this module in your path.

**SSQL–39: An error occurred while using filename. Error code: nn.**

There was a problem using the *filename* module. Refer to error code *nn* for detailed information.

**SSQL–40: An error occurred while using filename. xxxx.**

There was a problem using the *filename* module. Refer to the message following the error message for detailed information.

**SSQL–41: The Scalable SQL Engine is running with a user count of zero.**

A problem occurred while obtaining the user count because the user count is not configured correctly or it has not been incremented. Use the User Count utility to configure the user count. Because the user count is zero, you only have local access; remote access is disabled.

**SSQL–42: Unexpected status “nn” received creating system threads.**

During the Scalable SQL engine’s initialization process, an error occurred in creating system threads. The system resources may be low. Try rebooting the system on which you are loading the Scalable SQL engine.

**SSQL–43: Unexpected status “nn” received creating worker threads.**

During the Scalable SQL engine’s initialization process, an error occurred in creating worker threads. The system resources may be low. Try rebooting the system on which you are loading the Scalable SQL engine.

**SSQL–44: Unable to access named database information.**

Scalable SQL engine can neither open nor create the DBNAMES.CFG file.

**SSQL–45: An error occurred while loading Database Services.**

A fatal error occurred while loading Database Services.

**SSQL–46: An error occurred while loading the Scalable SQL engine. Error code: xx.**

- 1 This is an internal error and should not occur.
- 2 Out of memory. Windows NT system resources may be low. Try rebooting the System on which you are loading the Scalable SQL engine.
- 3 Error creating Semaphore. Windows NT system resources may be low or a previous copy of DBSvcMgr did not exit cleanly. Try rebooting the system on which you are loading the Scalable

SQL engine.

- 4 DBSvcMgr has received a new request but it was already signaled to shutdown. Wait for the Scalable SQL engine or MicroKernel to shutdown before starting another engine.
- 5 This is an internal error and should not occur.
- 6 The Scalable SQL engine module is already loaded. This is an informative status that should cause Scalable SQL to display Message SSQL-WNT-27.
- 7 This is an internal error and should not occur.
- 8 This is an internal error and should not occur.
- 9 Invalid DBSvcMgr command sent to DBSvcMgr. This could be seen if a newer stub is loading an older DBSvcMgr. Verify that the correct versions of all components are being used and are on your path. If the problem persists reinstall the engine files.
- 10 This is an internal error and should not occur.
- 11 This is an internal error and should not occur.
- 12 Components out of sync. Verify that the correct versions of all components are being used and are on your path. If the problem persists, reinstall the engine files.
- 13 This is an internal error and should not occur.
- 14 This is an internal error and should not occur.
- 15 This is an internal error and should not occur.
- 16 The Scalable SQL engine could not be loaded because the DLL could not be found. Verify that the correct versions of all components are being used and are on your path. If the problem persists, reinstall the engine files.
- 17 This is an internal error and should not occur.
- 18 Scalable SQL engine files could be corrupted. Verify that the correct versions of all components are being used and are on your path. If the problem persists, reinstall the engine files.
- 19 Scalable SQL engine files could be corrupted. Verify that the correct versions of all components are being used and are on your path. If the problem persists, reinstall the engine files.

**SSQL-47: An error occurred while locating <component>. Error code: <nn>. Binding rule: <rule>.**

Scalable SQL was attempting to load a component and could not complete the task. The component that the engine cannot find is listed in the message. Verify that this file is on your system. Contact Pervasive Software Customer Support for more information.

**SSQL-48: Your Pervasive.SQL license expired on <date>. Please contact Pervasive Software for more information.**

Your user license for Pervasive.SQL was temporary and has expired. Contact Pervasive Software at the telephone number listed on this manual to obtain a permanent license key.

**SSQL-49: Your Pervasive.SQL license will expire on <date>. Please contact Pervasive Software for more information.**

This message displays as a warning to notify you that your database engine will not continue to run unless you purchase a permanent license key. Your user license for Pervasive.SQL is temporary and will expire at the date shown on the message. Contact Pervasive Software at the telephone number listed on this manual to obtain a permanent license key. If the license expires, you will no longer be able to use the Pervasive.SQL engine.

**SSQL-DOS-19: The file SSQL.EXE will be unloaded.**

This message appears after other error messages to inform you that the file is not loaded.

**SSQL-DOS-27: SSQL.EXE is already loaded.**

Scalable SQL returns this message if you try to load Scalable SQL when it is already loaded at the DOS client machine.

**SSQL-DOS-28: Use SQLSETUP to modify options in BTI.CFG.**

Scalable SQL returns this message if you try to set options at the command line. You should use the DOS SQLSETUP.EXE to configure the Scalable SQL for DOS engine.

**SSQL-NW-19: The Scalable SQL Engine, nn, has shutdown.**

While loading, the Scalable SQL engine received an error and cannot finish initializing. The error is displayed before message SSQL-NW-19. Correct the cause of the error, unload and then reload the Scalable SQL engine.

## STPCOM

### Scalable SQL TCP/IP Communications Module Messages

#### **STPCOM-1: The option specified is not a valid option.**

Specify a valid option.

#### **STPCOM-2: The server has insufficient memory to execute STPCOM.NLM.**

An error occurred during a memory allocation operation. Free memory by unloading NLMs or reconfiguring NLMs to use less memory. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **STPCOM-3: An internal error has occurred. STPCOM.NLM detected a semaphore allocation failure.**

An error occurred during a semaphore allocation operation. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **STPCOM-4: The Service Request Block (SRB) function code nn contains invalid data. Check for an incompatible version of the file STPCOM.NLM.**

Ensure that STPCOM's version is compatible with the version number of the workstation's Scalable SQL Requester.

#### **STPCOM-5: A TCP/IP level receive I/O error nn has occurred. The connection has been lost.**

An error occurred during a TCP/IP level receive operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **STPCOM-6: A TCP/IP level send I/O error nn has occurred. The connection has been lost.**

An error occurred during a TCP/IP level send operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **STPCOM-7: A TCP/IP accept error (nn) has occurred.**

An error occurred during a TCP/IP level accept session establishment operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation. If the error persists, contact Pervasive Software Customer Support for more information.

#### **STPCOM-8: The session was rejected because the session limit was reached. Increase the value specified for the Number of Remote Sessions option.**

Increase the value specified for the Number of Remote Sessions option. This option is described in the *Pervasive User's Guide*.

**STCPCOM-9: The request for statistics from the Monitor utility was not recognized. Check for an incompatible version of the utility or STCPCOM.NLM.**

Ensure that the STCPCOM.NLM version is compatible with the version number of the Monitor utility.

**STCPCOM-10: An internal error has occurred. STCPCOM did not recognize the GET\_EIM\_STATS function.**

Ensure that the STCPCOM.NLM version is compatible with the version number of the Monitor utility.

**STCPCOM-12: A TCP/IP select() error (nn) has occurred while trying to establish a session requested by a remote workstation.**

An error occurred during a TCP/IP-level select session establishment operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

**STCPCOM-13: A thread initialization error has occurred.**

An error occurred during STCPCOM.NLM thread initialization. Refer to the Pervasive Software eventlog (PVSU.LOG) for error messages preceding this one. Correct the problem and retry the operation.

**STCPCOM-15: A TCP/IP socket() error (nn) has occurred.**

An error occurred during a TCP/IP-level socket operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

**STCPCOM-16: A TCP/IP setsockopt() error (nn) has occurred.**

An error occurred during a TCP/IP-level setsockopt operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

**STCPCOM-17: A TCP/IP bind() error (nn) has occurred.**

An error occurred during a TCP/IP-level bind operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

**STCPCOM-18: A TCP/IP listen() error (nn) has occurred.**

An error occurred during a TCP/IP-level listen operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation again. If the error persists, call Pervasive //Software technical support for more information.

**STCPCOM-19: A TCP/IP gethostname() error (nn) has occurred.**

An error occurred during a TCP/IP-level gethostname operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

### **STCPCOM-20: A TCP/IP NWgethostbyname() error (nn) has occurred.**

An error occurred during a TCP/IP-level NWgethostbyname operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

### **STCPCOM-21: A TCP/IP NetDBgethostbyname() error (nn) has occurred.**

An error occurred during a TCP/IP-level NetDBgethostbyname operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

### **STCPCOM-22: A BeginThread() error (nn) has occurred.**

An error occurred during a BeginThread operation. Refer to your operating system documentation for the definition of the errno value. Retry the operation again. If the error persists, call Pervasive Software technical support for more information.

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](http://www.pervasive.com) All rights reserved.

## UCINIT

### User Count Initialization Utility Messages

The following messages apply to the User Count Initialization utility.

#### **UCINIT-100: Abnormal completion, status = nn.**

The User Count Initialization utility returns this message when an operating system or User Count Manager error occurs. If the specified status is -1, the error relates to the operating system. A message reflecting the system error number is displayed just above this one on the terminal. Refer to your operating system documentation for the definition of the error number value. If the specified status is in the range of 7000 to 7099, the error relates to the User Count Manager. A message with the User Count Manager error is displayed just above this one on the terminal. Refer to the User Count Manager error messages for the definition of the error.

#### **UCINIT-203: Keyboard input error.**

The User Count Initialization utility displays this message when an operating system error is reported during keyboard input. Retry the keyboard input operation.

#### **UCINIT-308: The User Count Manager cannot migrate the Btrieve user count because the Btrieve file is invalid, file = %s.**

The User Count Initialization utility displays this message when you try to migrate a Btrieve user count with an invalid Btrieve v6.10 database engine file (BTRIEVE.NLM).

#### **UCINIT-309: Error adding license key while migrating Btrieve v6.1x user count (UCMGR status code = nn).**

The User Count Initialization utility displays this message when attempting to migrate a Btrieve v6.1x user count. For NetWare server engines, ensure that the Btrieve file is v6.10 or later. For Windows NT server engines, ensure that the Btrieve file is v6.15 or later.

#### **UCINIT-400: File does not exist, file = filename, (errno = nn).**

The User Count Initialization utility displays this message when trying to access a file that does not exist. Ensure that the specified file exists and retry the operation.

#### **UCINIT-404: File open error, file = filename, (errno = nn).**

The User Count Initialization utility displays this message when it cannot open the specified file. The operating system error number is displayed as `errno=nn`. Refer to your operating system documentation for the definition of the `errno` value. Ensure that your disk is not full and that you have appropriate access rights to the specified file.

#### **UCINIT-407: File read error, file = filename, (errno = nn).**

The User Count Initialization utility displays this message when it cannot read from the specified file. The operating system error number is displayed as `errno=nn`. Refer to your operating system documentation for the definition of the `errno` value. Ensure that the specified file is not damaged. If the file is damaged, delete it and re-execute the utility.

#### **UCINIT-411: File seek error, file = filename, (errno = nn).**

The User Count Initialization utility displays this message when it cannot seek in the specified file. The operating

system error number is displayed as `errno=nn`. Refer to your operating system documentation for the definition of the `errno` value. Ensure that your disk is not full and that you have appropriate access rights to the specified file.

### **UCINIT-503: Terminating, cannot access internal message table.**

The User Count Initialization utility displays this message when it cannot access its internal message table. This message is only returned on NetWare systems. Ensure that the utility is either installed in the `SYS:SYSTEM` directory, or that the directory from which you loaded the utility is in your path.

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](http://www.pervasive.com) All rights reserved.

## UCMGR

### User Count Manager Messages

The following messages apply to the User Count Manager.

#### **UCMGR–301: The key number used for increasing the user count is invalid.**

The User Count Manager returns this message when the key number used for increasing the user count is invalid. Provide a valid key number.

#### **UCMGR–302: Invalid key file check sequence. The integrity of the user count key file has been compromised.**

The User Count Manager returns this message when the integrity of UCMGR.SYS has been compromised. Reinstall Scalable SQL.

#### **UCMGR–303: Invalid key number (key number is of the wrong type).**

The User Count Manager returns this message when you attempt to provide a key number that is of the wrong type. Ensure that you entered the key number correctly.

#### **UCMGR–305: Duplicate key number error. You cannot increase the user count with the same key twice.**

The User Count Manager returns this message when you attempt to provide a key number that has already been used to increase the user count. Ensure that you entered the key number correctly.

#### **UCMGR–306: The diskette key has already been applied. You cannot increase the user count with the same diskette key twice.**

The User Count Manager returns this message when you attempt to apply a diskette-based key that has already been applied. Ensure that you use a diskette-based key that has not been previously applied.

#### **UCMGR–307: The diskette key file is invalid, file = filename, (errno=nn)**

The User Count Manager returns this message when an invalid diskette key file has been detected. Retry the operation using another diskette key file.

#### **UCMGR–310: Serial number synchronization error. Unable to find valid stamp area.**

The User Count Manager returns this message when the serial numbers in Scalable SQL components do not match. Ensure that you installed Scalable SQL correctly.

#### **UCMGR–311: Serial number synchronization error. Invalid stamp area read.**

The User Count Manager returns this message when the serial numbers in Scalable SQL components do not match. Ensure that you installed Scalable SQL correctly.

#### **UCMGR–312: Serial number synchronization error. One or more stamp**

### **numbers read do not match.**

The User Count Manager returns this message when the serial numbers in Scalable SQL components do not match. Ensure that you installed Scalable SQL correctly.

### **UCMGR–313: Serial number synchronization error. Could not write stamp number to one of the files. Make sure the database engine is not loaded and that you have access rights to the files.**

The User Count Manager returns this message when the serial numbers in Scalable SQL components do not match. Ensure that you installed Scalable SQL correctly.

### **UCMGR–314: Serial number synchronization error. Files have not been stamped and the user count key file exists.**

The User Count Manager returns this message when the serial numbers in Scalable SQL components do not match. Ensure that you installed Scalable SQL correctly.

### **UCMGR–315: Serial number synchronization error. The key file stamp number does not match the stamp number in the other files.**

The User Count Manager returns this message when the serial numbers in Scalable SQL components do not match. Ensure that you installed Scalable SQL correctly.

### **UCMGR–316: Serial number synchronization error. Files have been stamped but the user count key file does not exist. If you have made a backup copy of your user count key file (<plat>UCMGR.SYS), please restore it to the database engine directory.**

The User Count Manager returns this message when the serial numbers in Scalable SQL components do not match. Ensure that you installed Scalable SQL correctly.

### **UCMGR–317: Serial number synchronization error. During a user count migration update, could not find a valid stamp area.**

The User Count Manager returns this message when the serial numbers in Scalable SQL components do not match. Ensure that you installed Scalable SQL correctly.

### **UCMGR–318: Serial number synchronization error. During a user count migration update, an invalid stamp number was found.**

The User Count Manager returns this message when the serial numbers in Scalable SQL components do not match. Ensure that you installed Scalable SQL correctly.

### **UCMGR–319: Serial number synchronization error. During a user count migration update, one or more stamp numbers do not match.**

The User Count Manager returns this message when the serial numbers in Scalable SQL components do not match. Ensure that you installed Scalable SQL correctly.

### **UCMGR–320: Serial number synchronization error.**

The User Count Manager returns this message when the serial numbers in Scalable SQL components do not match. Ensure that you installed Scalable SQL correctly.

### **UCMGR–321: You have previously migrated xx users and have tried to migrate xx users at this time. This is an invalid operation because the migrated user count is less than the previously migrated user count.**

The User Count Manager returns this message when you attempt to migrate a number of users that is different from the number of users you previously migrated.

### **UCMGR–401: Cannot find diskette key file, file = filename, (errno = nn).**

The User Count Manager returns this message when you attempt to access a diskette-based key file that does not exist. Ensure you have a diskette-based key file and retry the operation.

### **UCMGR–402: Invalid file pathname = path, file = filename, (errno = nn).**

The User Count Manager returns this message when it determines that a file pathname is invalid. The operating system error number is specified by *errno*. Refer to your operating system documentation for more information about the *errno* number. Ensure that the pathname is correct and that you have appropriate access rights.

### **UCMGR–403: File change mode error, file = filename, (errno=nn).**

The User Count Manager returns this message when trying to change the file attributes of the specified file. The operating system error number is specified by *errno*. Refer to your operating system documentation for more information about the *errno* number. Ensure that the specified file is not already open and that you have appropriate access rights to the specified file.

### **UCMGR–404: File open error, file = filename, (errno=nn).**

The User Count Manager returns this message when trying to open the specified file. The operating system error number is specified by *errno*. Refer to your operating system documentation for more information about the *errno* number. Ensure that your disk is not full and that you have appropriate access rights to the specified file. When this error occurs, ensure that the database engine is not loaded.

- In Netware, enter the following commands at the server console:

```
sqlstop
```

```
bstop
```

- In Windows NT, enter the following command at the Scalable SQL console:

```
SSQL_v4.0> quit
```

### **UCMGR–405: Error writing to user count key file, (errno=nn).**

The User Count Manager returns this message when trying to write to the user count key file (UCMGR.SYS). The operating system error number is specified by *errno*. Refer to your operating system documentation for more information about the *errno* number.

### **UCMGR–406: Error writing to file, file = filename, (errno=nn).**

The User Count Manager returns this message when trying to write to the specified file. The operating system error number is specified by *errno*. Refer to your operating system documentation for more information about the *errno* number. Ensure that your disk is not full and that you have appropriate access rights to the specified file.

### **UCMGR–408: Error reading from user count key file, (errno=nn).**

The User Count Manager returns this message when trying to read the user count key file (UCMGR.SYS). The operating system error number is specified by *errno*. Refer to your operating system documentation for more information about the *errno* number. Ensure that your disk is not damaged. If the disk is damaged, reinstall Scalable SQL.

### **UCMGR–409: Error reading from file, file = filename, (errno=nn).**

The User Count Manager returns this message when trying to read from the specified file. The operating system error number is specified by *errno*. Refer to your operating system documentation for more information about the *errno* number. Ensure that the specified file is not damaged. If the file is damaged, reinstall Scalable SQL.

### **UCMGR–410: Unexpected end-of-file found, file = filename, (errno=nn).**

The User Count Manager returns this message when trying to read the specified file. The operating system error number is specified by *errno*. Refer to your operating system documentation for more information about the *errno* number. Ensure that the specified file is not damaged. If the file is damaged, reinstall Scalable SQL.

### **UCMGR–411: File seek error, file = filename, (errno=nn).**

The User Count Manager returns this message when trying to seek in the specified file. The operating system error number is specified by *errno*. Refer to your operating system documentation for more information about the *errno* number. Ensure that your disk is not full and that you have appropriate access rights to the specified file.

### **UCMGR–412: File delete error, file = filename, (errno=nn).**

The User Count Manager returns this message when trying to unlink the specified file. The operating system error number is displayed as *errno*. Refer to your operating system documentation for the definition of the *errno* value. Ensure that you have appropriate access rights to the specified file.

### **UCMGR–413: File rename error, file = filename, (errno=nn).**

The User Count Manager returns this message when trying to rename the specified file. The operating system error number is displayed as *errno*. Refer to your operating system documentation for the definition of the *errno* value. Ensure that you have appropriate access rights to the specified file.

### **UCMGR–415: DOS Partition Services are not present.**

This message is returned only on the NetWare platform. The User Count Manager returns this message when trying to access the User Count diskette key file (UCMGR.KEY) and the DOS partition services are not present in NetWare. Reboot your NetWare server and ensure that the DOS partition services are present.

### **UCMGR–416: DOS Copy error. The User Count Manager cannot copy the diskette key file, (errno=nn).**

This message is returned only on the NetWare platform. The User Count Manager returns this message when it cannot access the User Count diskette key file (UCMGR.KEY) on NetWare. The operating system error number is displayed as *errno*. Refer to your operating system documentation for the definition of the *errno* value. Ensure you have a valid diskette-based key file and retry the operation.

### **UCMGR-417: DOS Open error. The User Count Manager cannot open the diskette key file, (errno=nn).**

This message is only returned on the NetWare platform. The User Count Manager returns this message when it cannot access the User Count diskette key file (UCMGR.KEY) on NetWare. The operating system error number is displayed as errno. Refer to your operating system documentation for the definition of the errno value. Ensure you have a valid diskette-based key file and retry the operation.

### **UCMGR-418: DOS Write error. The User Count Manager cannot write to the diskette key file (errno=nn).**

This message is only returned on the NetWare platform. The User Count Manager returns this message when it cannot access the User Count diskette key file (UCMGR.KEY) on NetWare. The operating system error number is displayed as errno. Refer to your operating system documentation for the definition of the errno value. Ensure you have a valid diskette-based key file and retry the operation.

### **UCMGR-500: The User Count Manager is not initialized. Unload or stop your database engine and run the User Count Initialization utility (UCINIT.EXE, NTUCINIT.EXE, or NWUCINIT.NLM).**

The User Count Manager returns this message when the manager has not been initialized prior to issuing requests for its services.

### **UCMGR-501: The User Count Manager is busy. Only one instance of the User Count Manager can be in use at one time.**

The User Count Manager returns this message when the Manager is already in use. Try to use the User Count Manager later.

### **UCMGR-502: The User Count Manager cannot determine the platform.**

The User Count Manager returns this message when it cannot determine the correct platform during an open operation. Ensure that the MicroKernel is in the target directory and retry the operation.

### **UCMGR-503: Terminating, cannot access internal message table.**

This message is returned only on the NetWare platform. The User Count Manager returns this message when it cannot access its internal message table. Ensure that the utility is either installed in the SYS:SYSTEM directory, or that the directory from which you loaded the utility is in your path.

# UCUTIL

## User Count Utility Messages

The User Count utility generates the following messages.

### **UCUTIL–100: Abnormal completion, status = nn.**

The User Count utility returns this message when an operating system or User Count Manager error occurs. To recover from this error, use the specified status number, as follows:

- If the specified status code is in the range of 7000 through 7099, the error relates to the User Count Manager. A message reflecting the User Count Manager error is displayed just above this one at the console. Refer to the UCMGR messages for more information.
- If the specified status code is -1, the error relates to the operating system. A message reflecting the operating system error number (*errno*) is displayed just above this one at the console. Refer to your operating system documentation for more information about the *errno* number.

### **UCUTIL–201: Invalid key number length.**

The User Count utility returns this message when you enter a key number that is not 16 characters long. Ensure that you have entered the key exactly as Pervasive Software provides it.

### **UCUTIL–203: Keyboard input error.**

The User Count utility returns this message when an operating system error is reported during keyboard input. Retry the keyboard input operation.

### **UCUTIL–204: Invalid product code.**

The User Count utility returns this message when you enter an invalid product code. The product code for Btrieve is 1, and the product code for Scalable SQL is 2.

### **UCUTIL–205: Invalid diskette drive specification.**

The User Count utility returns this message when you specify an invalid diskette drive. The valid diskette drives are a: and b:.

### **UCUTIL–400: File does not exist, file = filename, (errno = nn).**

The User Count utility displays this message when trying to access a file that does not exist. Ensure that the specified file exists and retry the operation.

### **UCUTIL–404: File open error, file = filename, (errno = nn).**

The User Count utility displays this message when it cannot open the specified file. The operating system error number is displayed as *errno=nn*. Refer to your operating system documentation for the definition of the *errno* value. Ensure that your disk is not full and that you have appropriate access rights to the specified file.

### **UCUTIL–407: File read error, file = filename, (errno = nn).**

The User Count utility displays this message when it cannot read from the specified file. The operating system error number is displayed as *errno=nn*. Refer to your operating system documentation for the definition of the *errno* value.

Ensure that the specified file is not damaged. If the file is damaged, delete it and re-execute the utility.

### **UCUTIL-411: File seek error, file = filename, (errno = nn).**

The User Count utility displays this message when it cannot seek in the specified file. The operating system error number is displayed as `errno=nn`. Refer to your operating system documentation for the definition of the `errno` value. Ensure that your disk is not full and that you have appropriate access rights to the specified file.

### **UCUTIL-503: Terminating, cannot access internal message table.**

This message is returned only on the NetWare platform.

The User Count utility returns this message when it cannot access its internal message table. Ensure that the utility is either installed in the `SYS:SYSTEM` directory, or that the directory from which you loaded the utility is in your path.

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](http://www.pervasive.com) All rights reserved.



## Changes Since Btrieve 6.15

A number of status codes have been disambiguated in Pervasive.SQL. This section contains a list of status codes that in Pervasive.SQL are supplemented with more specific codes.

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## Status Code 3 replacements

The following status codes are returned instead of Status Code 3 in Pervasive.SQL.

[3006: The MicroKernel router detected an invalid session.](#)

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## Status Code 20 replacements

The following status codes are returned instead of Status Code 20 in Pervasive.SQL.

[2009: Cannot load MicroKernel Router component.](#)

[2010: Thunk not possible - the Win32 MicroKernel router could not be found or is incompatible.](#)

[3001: Local access is unavailable to the MicroKernel router.](#)

[3003: The MicroKernel router detected an incompatible network component.](#)

[3008: Invalid configuration for MicroKernel router.](#)

[3009: NETAPI.DLL is not loaded.](#)

[3010: NetWare API DLLs are not loaded.](#)

[3011: Thunk not possible - the Win32 MicroKernel router could not be found or is incompatible.](#)

[3013: The remote engine is inaccessible to the MicroKernel router because the networking component is not loaded.](#)

[3014: The MicroKernel router cannot find an engine.](#)

[3020: An error occurred while loading the MicroKernel.](#)

[3103: Server name not found by Pervasive Network Services Layer.](#)

[3104: The Pervasive Network Services Layer encountered a permission error.](#)

[3105: No available transport protocol for the Pervasive Network Services Layer.](#)

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## Status Code 75 replacement

The following status code is returned instead of Status Code 75 in Pervasive.SQL.

[3108: The Pervasive Network Services Layer detected an invalid session.](#)

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## Status Code 91 replacements

The following status codes are returned instead of Status Code 91 in Pervasive.SQL.

[3106: The Pervasive Network Services Layer encountered a connection failure.](#)

[3115: Pervasive Network Services Layer encountered a transport failure.](#)

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## Status Code 95 replacements

The following status codes are returned instead of Status Code 95 in Pervasive.SQL.

[3110: The network layer is not connected.](#)

[3111: Failure during send to the target server.](#)

[3112: Failure during receive from the target server.](#)

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## Status Code 109 replacements

The following status codes are returned instead of Status Code 109 in Pervasive.SQL.

[3004: The MicroKernel is not responding to the MicroKernel router.](#)

[3019: The MicroKernel router encountered a semaphore error.](#)

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

## Status Code 2003 replacement

The following status code is returned instead of Status Code 2003 in Pervasive.SQL.

[3012: Local engine is not accessible to the MicroKernel router.](#)

---

Send comments to [docs@pervasive.com](mailto:docs@pervasive.com). Copyright © 1998 [Pervasive Software Inc.](#) All rights reserved.

