

Centura Team Developer 1.0 Release Notes

About these Notes

This file provides information on Centura Team Developer that is not included in the documentation or online help. It documents general information, installation and migration issues, known problems, and documentation errata.

These release notes are also available as a Microsoft Write file (TDNOTES.WRI) if you wish to print them.

Centura 1.0 up-to-date: Changes and New Features

In Centura online help, there is a help topic called **Centura 1.0 up-to-date**. This topic contains all the latest information on Centura Team Developer. The following subjects are in this help topic and you can find them by searching on the keywords *Centura up-to-date*.

Centura Builder

- SQLWindows 4 scope rules in Centura
- Migrating OLE2 applications from SQLWindows
- SAL functions for idle processing
- Recursive reentry of SQL/API
- 3-Tier Wizard requires Tuxedo Workstation software
- Loss of precision when displaying numbers
- Displaying properties of picture controls
- Displaying properties of QuickTabs
- QuickMail supports MAPI only
- Migrating SQLWindows applications from Solaris
- Migrating SQLWindows 4 applications
- String handling in DLLs

Object Compiler

- Using Microsoft Visual C++ compiler

SQLBase Database Engine

- Support for heterogeneous data replication
- 32-bit SPX Windows 95 client support
- Native Thread Support
- Single user server
- System catalog tables
- Migration: Some load or unload stored commands must be recompiled
- Migrating from one platform to another
- Serverpath keyword semantics
- ALTER TRIGGER command
- ALTER TRIGGER associated changes to system tables
- UNLOAD and LOAD operations and the ALTER TRIGGER command
- Stored procedure provided to disable all triggers on a table
- Stored procedure description

SQLBase and Ranger Replication

- SQLBase and Centura Ranger Replication packages
- Running RMS for the first time
- Centura Ranger to Oracle package: Create the SYSREP user
- Significant performance enhancements
- Remote access connectivity using Microsoft RAS
- Temporary file created by RSA
- AUTO GRANT and subscriber privileges
- Connecting to an Oracle database
- Generating your replication set
- Synchronizing with Oracle

SQLConsole

New Features

Compatibility with SQLConsole 6.0 repository

Information about the following Team Object Manager topics is not included in the documentation, but can be found in the Centura online help. You can locate information on these topics by searching for them in the Index or Find tabs of the Centura Help window.

Team Object Manager

Add PVCS archives

Logging in to Team Object Manager

Adding the Class Browser to the Builder toolbar

Merging files

Merging Conflicts and Warnings

Merging Classes with Properties

Merging classes with child contents

Merging - refreshing

Merging conflict - original copy

Command line options

Scanning for DataLinks

Information Preserved when Importing and Exporting Projects

Exporting and Importing an Entire Repository

Importing a Data Model

Update Key Property

Trace Mode

Default Printer Setting

Outline Browser

Installing Team Developer

Existing SQL.INI is backed up

If setup finds an existing copy of SQL.INI in your installation directory, it will rename it to SQL_OLD.INI before copying the new SQL.INI.

Compatibility with Book Viewer in TICD

Centura products use a newer version of the Interleaf WorldView Books Online viewer than the one which is shipped with the TICD. Both of these versions use an INI file (IVIEW.INI) which resides in the Windows directory. The two versions of the IVIEW.INI file are incompatible with each other, so you won't be able to use them both together. The setup program will detect if you already have an existing IVIEW.INI file and will back it up for you.

Conflict with SQL.INI installed by Sybase System 10/11

Sybase System 10 and 11 also install a SQL.INI file on your system. If you are running Centura Builder and have problems with connectivity, make sure that your PATH variable does not first point to the Sybase SQL.INI file.

Installing Centura Ranger Replication

Automatic conversion of SQLBase 6.0 or 6.0.1 databases. You can use an existing 6.0 or 6.0.1 version database with SQLBase 6.1 server. However, this database is automatically converted to 6.1 format when you connect to it with the 6.1 server. Additionally, it is recommended that you first UNLOAD your earlier 6.x database to a file with your current server so that you have a copy of all your data. Then, you can proceed with using the database with SQLBase 6.1.

The following procedure is recommended:

1. Using the SQLBase 6.0 or 6.0.1 server, use the UNLOAD command to unload your database.
2. Ideally, you should do this before installing the Replication software.
3. Shut down the SQLBase 6.0 or 6.0.1 server.

4. Install and start the SQLBase 6.1 server.
5. Connect to the database with SQLBase 6.1.

If you plan to use replication with a NetWare SQLBase server, you must use the TCP/IP protocol from a Windows NT client (since there is no SPX support from a Windows NT client -this is a Novell limitation).

Centura Builder

The following section contains information about Centura Builder.

3-Tier Wizard requires Tuxedo Workstation software

If you are using the 3-Tier Wizard to build Tuxedo applications, or if you include QCKTUX.APL, you must have Tuxedo Workstation software installed on your client PC. If you do not, you will get an error message indicating that an external library can not be found.

Migrating SQLWindows applications containing cQuickOLE2

Displaying or running a SQLWindows application with a cQuickOLE2 container may display an error:

"An error occurred while attempting to restore the OLE document. This could be caused by attempting to use an older version of the application possibly ported from 16bit SQLWindows. Try re-inserting your OLE objects from the toolbar and saving your application."

Click on the "Choose an OLE Control (OCX)" or "Insert New Compound Document Object" button in the OLE container's toolbar and re-insert the OCX or OLE object. This will fix the problem.

Migrating QuickGraphs from SQLWindows to Builder

Defect# 55073: SQLWindows applications containing instances of cQuickGraph require modifications on migration to Builder. The user must display the property editor for the cQuickGraph instance and re-apply the properties.

New: Saving SAL compiler output to a file

The Output Window that contains the multiple error statements from the SAL compiler now can have its contents sent to an error log. To do this, right-mouse click on the Output Window and choose the menu item to save the contents to a file.

New: Constants for OS platforms

These constants have been added to facilitate cross-platform coding: PLATFORM_Current, PLATFORM_Windows, and PLATFORM_Solaris. Programmer can write code such as:

```
If PLATFORM_Current = PLATFORM_Windows
    ! Do Microsoft Windows specific actions
Else If PLATFORM_Current = PLATFORM_Solaris
    ! Do Solaris/UNIX specific code
Else
    ! Give error about unknown platform
```

Object Compiler

The following section contains information about Object Compiler.

Using Borland C++ compiler

The Object Compiler does not support the Borland C++ compiler.

Dynalink functions

Dynalink functions must be compiled before they can be called by other compiled functions.

Team Object Manager

The following section contains information about Team Object Manager.

Cannot upgrade Repository from early beta versions

There is no automatic migration of a Repository from early beta versions of Team Object Manager to the final released product- you **must** create a new repository database.

If you had created projects using a prior version, you should ensure that a project in the new version does not use the same project directory as a project in the prior version.

Diff/Merge and Object Browser Utilities

These are the file names for the Diff/Merge and Object Browser utilities:

tmdti10.exe	Diff/merge utility
tmtoi10.exe	Object Browser utility

Logging in to Team Object Manager

The first time you log into the repository using Team Object Manager, you should log in as the user "DEMO". (For Sybase System 10, the password is "DEMO10", and for all other databases, the password is "DEMO".) You may create other users using the appropriate options in Team Object Manager.

Creating a Repository on SQLBase

In the Repository Setup Wizard, creation of a database on the server may fail, under certain unlikely conditions. (This is generally the result of conflicting dll file versions with previously installed products.) (This only refers to the physical creation of the database itself - creation of tables, indexes, etc., is not affected.) In this case, create the database first, using the Database Explorer or SQLTalk, and then connect to it using the Repository Setup Wizard.

Display of windows

The Layout tab of a Window shows a "thumbnail" picture of the window. This thumbnail will display, but Team Object Manager may crash on occasions when you try to leave the tab. This problem is intermittent, but is usually related to windows which have a toolbar.

Demo.prj file

A sample project has been supplied with this version of Team Object Manager. To import this file, select Import from the Repository context menu, and select the file "demo.prj".

Registry entries from earlier beta versions

Earlier versions of Team Object Manager created registry entries under the section:

HKEY_CURRENT_USER\Software\Gupta\Team Object Manager\...

This information is now stored in

HKEY_CURRENT_USER\Software\Centura\Team Object Manager\...

You may either:

- Move the existing registry information to the new location using regedit.exe, or
- Simply delete the existing registry information, and allow Team Object Manager to recreate it in the new location.

Failure to delete file

A number of operations (such as checking in or deleting a file) may result in a warning that a cache file (.che) or project file (.pf) could not be deleted. This is usually due to another user currently viewing the file, and a lock is currently being held on the file. (The panes that this occurs on are the outline, window and class hierarchy views of a file.) You may take a note of the filename, and delete this file later, once the other user has finished viewing. Note: be extremely careful to only delete the file that was mentioned in the message - deleting any file other than this one from the project directory will result in repository corruption. If any doubt, rather leave the file undeleted - this will not cause any problems other than taking up unnecessary space.

Repository Setup Wizard - supported databases

This version of Team Object Manager will correctly install a repository on the following databases:

- SQLBase 6.0.1
- SQLBase 6.1 (in beta)
- Oracle 7
- Microsoft SQL Server 6.0
- Sybase System 10.0.3

Storage in database (blobs)

You may experience problems with storage of binary large objects, or blobs (i.e. storage in database) with Oracle, Sybase and/or SQL Server, depending on the version of client, server and SQLRouter software. This is not a problem that is particular to Team Object Manager, rather just a problem related to blobs.

Note: you may not immediately notice any problems with blobs (i.e. storage as database), since Team Object Manager will always use the cache file if it exists, and only extract the blob if the cache file does not exist. To correctly test whether blobs are working correctly, either i) delete all cache files before extracting/checking out (cache files have a .che extension), or ii) extract a version of the file other than the latest version.

To work around any problems you may encounter with blobs, create all your projects using "storage as files". Please refer to the the Frequently Asked Questions of the help file for advice on whether to use storage as files, or storage in the database.

Sybase error #590

There is a known problem with Sybase System 10.0.2 on Windows NT, which may also occur on other platforms. An attempt to select an image column causes the error "Prepared statement 'Q2' does not exist". This is Sybase error #590, and is fixed by EBF #5596. You can also work around this problem by only using storage as files with Sybase.

Creating team members

There are a number of known problems creating or modifying team members. These errors generally occur with databases other than SQLBase.

- Creating a user with shadow user in Sybase fails.
- Deleting a user and then recreating the same user fails in Oracle.
- Creating a user in Sybase succeeds, but the database user is not correctly created. (Create the user manually to work-around this problem.).
- Changing password or privileges of a newly created user fails.

Version Information

This version of Team Object Manager was built and tested with Builder Build 10384 under Windows 95 and Windows NT. It should be run with the DLLs from this version of Centura Builder, or a compatible version. Team Object Manager connectivity to SQLBase, Oracle, Sybase, and SQL Server was tested using files from Build 103847 - it is suggested that SQLBase and connectivity DLLs from this build or later be used. Testing was performed with SQLBase Server for Windows NT and Windows 95, Version 6.1.0 RC17.

Migrating from TeamWindows to Team Object Manager

The most obvious changes between TeamWindows and Team Object Manager are a new user interface, and new features, which are described fully in the documentation. However, some features have changed in a more subtle way, as described in this document.

Terminology

Some terminology has changed. For example:

- 'modules' are now known as 'files'
- file 'versions' are now known by the more specific term 'revisions'
- the 'development directory' is now the 'working directory'
- 'users' are now known as 'team members'
- the 'data dictionary manager' has been replaced by the 'data model viewer'
- the 'trash directory' is now the 'recycle bin path'.

Please refer to the glossary for more information.

Repository tables

The repository database of TeamWindows is not forward compatible with the database of Team Object Manager. The tables which make up the repository have been substantially changed in this version. To ensure that new and old repositories can co-exist on the same database, all the tables in the Team Object Manager repository are owned by the user TMADM, whereas TeamWindows tables were owned by the user TWADM. Any custom queries and reports, or custom applications which were developed against the old repository must be updated.

A Project Conversion Wizard has been provided to convert projects from a TeamWindows repository to a Team Object Manager repository. Please note that any SQLWindows files which are imported using this facility will not be automatically converted to Centura Builder format.

Revision numbers

In TeamWindows, version labels were restricted to a major and minor version number. In Team Object Manager, the version label is completely free-form text, which may be changed for any version. A sequential revision number is maintained internally by Team Object Manager for its own reference - this is non-editable.

Directories

Many of the directories which had to be specified in TeamWindows are no longer necessary. The only directories which are required are:

- the project directory, which stores files in a project, and
- a working directory for each Team Member, which is the default directory to which files are checked out and extracted.
- a default project directory, which is the default project directory for new projects, and which is set up when the repository is installed.

Other directories are required for builds, and other areas, but in most cases, their use is quite obvious from the context.

Project directory

A project directory is now required for both 'storage as files' and 'storage in database'. In the case of 'storage in database', the files in the project directory are not the 'master' copies of the files, but merely cache files used to improve checkout speed and allow on-line browsing. These files should never be modified directly.

Security

In TeamWindows, security was managed globally, and there were only four security groups. Security is now managed at the project level. You can create multiple team roles (or groups), each of which can have different privileges.

Security is also extended to the file level. Each file in a project belongs to a file security group, and the privileges of each file security group can be defined independently for different team roles. Please refer to the documentation for a full discussion of these features.

Promotion levels

Promotion levels are no longer fixed - you can have as many or few promotion levels as you desire, and name them whatever you want. TeamWindows had alternate promotion paths (test and shared) depending on whether the file was marked as sharable or not. This complexity has been removed.

Promotion

TeamWindows performed some fairly complex operations when files were promoted, including building of executables, and copying of files to designated directories. These operations depended on whether the file was marked as 'shared' or not, and what level the file was being promoted to. This feature, though quite powerful, is difficult to understand and use, and is inappropriate with multiple user-definable promotion levels.

This feature has been replaced by two simple features in Team Object Manager:

- The ability to place a copy of a file in a designed directory when it is promoted to a specified promotion level. You can control exactly which files are copied, and which promotion levels perform the copy.
- The build facility, which allows you to perform a build of all files at a particular promotion level (as well as other criteria).

The combination of these two new features provides much more power and flexibility, while being much easier to understand.

Archives

The concept of "module archives", both automatic and manual, has disappeared. Instead, Team Object Manager simply stores every single revision a file in a linear fashion, ordered by the date that the file was checked in. Each revision of a file may be promoted or demoted to any valid promotion level. Any revision of a file may be deleted if it is no longer needed. There is also an automatic purge option, which deletes the oldest revision automatically when a new revision is checked in.

Sharing files

TeamWindows supported file-sharing by copying files marked as 'shared' to a 'shared directory' when the file was promoted to the 'shared' promotion level. This feature is still supported, but is now much more general, because it can be defined at any of the user-definable promotion levels.

Another method of sharing files has been added to Team Object Manager - you simply mark a file as sharable in one project, and add the shared file to a second project. This file is then available for use in the second project, although it is still owned by the first project.

Folders

Team Object Manager allows you to define sub-folders in the repository, allowing you to group your files by any criteria. When you check out or extract a file, the same sub-directory structure will be created in the working directory.

It is possible to create folders in the repository, and still have all files checked out to the root of the working directory (see "preserve folder hierarchy" in the documentation). This is to allow you to define a meaningful structure in the repository, but still extract all the files to a single directory. (Centura Builder is easier to set up with all necessary source files in a single directory.)

Long file names

The names of files checked into the repository may be longer than eight characters, and may be of mixed case. If you wish to use long/mixed case file names, you should ensure that all team members of your project are extracting/checking-out to volumes which support long file names.

It is not necessary for the project directory to reside on a volume which supports long file names (even if long file names are used in the project), since Team Object Manager creates 8 character filenames for files in the project directory.

Teamwin.ini file and the registry

Team Object Manager now stores information which is local to a workstation in the registry. Much of the information in the registry is similar to the information which was stored in the teamwin.ini file. Team Object Manager will attempt to read the information from the old teamwin.ini file, and use this information as the basis for registry information. The old teamwin.ini file will not be modified or deleted, so old versions of Team Object Manager can run successfully after the Team Object Manager has been run.

Export file format

Because of the changes to the repository tables, .prj files which have been exported from TeamWindows may not be imported into Team Object Manager.

Converting projects

You may convert a project in a TeamWindows into a project in the Team Object Manager repository by using the

Project Conversion Wizard. This process prompts you to log in to the old TeamWindows repository, and select a project. It then imports all revisions of all files to the new project, as well as any logs.

If you do not want to convert the entire project, you may simply use TeamWindows to extract the latest revision of each file to a temporary directory, and then use Team Object Manager to add all these files to a new project.

Converting data-dictionary information

Data dictionary information in TeamWindows projects can also be imported into Team Object Manager. You can:

- use the "Import from TeamWindows" option in the data model viewer menu, or
- use the "Import from database" option to import directly from an instance of the database.

File names

The names of all the Team Object Manager files have been changed from prior versions. All Team Object Manager executables and support files have a prefix of 'tm'. Help files, report files, and documentation files do not have the 'tm' prefix.

Branching

A simple form of branching was available in TeamWindows. File could be maintained independently at Development, Test, and Production versions. However, this was also potentially dangerous, since if you checked out a Production version, and checked it in, it would become the new Development version, and you would lose any work that had been performed at the Development level. The flexibility of this branching approach was limited to only three branches.

Team Object Manager takes a structured, project-oriented approach to branching. To avoid accidentally overwriting work, only the most recent version of a file may be checked out. And to allow greater flexibility, any number of branches of an existing project may be created at any time, using a range of criteria to specify which revision of each file appears in the branched project.

Two different types of branching are supported. The first is appropriate for a new version of a project, and the second is appropriate for customizations of a project. The main difference between these two types of branches is whether changes to files in the branch project are synchronized to the files in the base project. Please see the documentation for further information.

SQLRouter/Sybase

The following sections contain information about SQLRouter/Sybase.

SQLRouter/Sybase Connectivity

SQLRouter/Sybase provides native connectivity from Centura Team Developer to Sybase System 10 and Sybase System 11 using Sybase CT-Lib.

SQLRouter/Sybase compatibility

This version of SQLRouter/Sybase is certified with:

Operating Systems

- Microsoft Windows 95
- Windows NT 3.51

Centura software

- Centura Builder v1.0

Sybase System 10 software:

- Sybase SQL Server Release 10.0.2.3/P EBF 5539 Rollup for Sun-4/OS 4.1.x
- Sybase Open Client Client-Library Release 10.0.3 for Windows 95/ Windows NT
- Sybase Open Client Net-Library Release 10.0.3 for Windows 95/ Windows NT

Sybase System 11 software

- Sybase SQL Server Release 11.0 for Sun-svr4/OS 5.4 (Solaris)

Configuring SQL.INI for Sybase

To connect to Sybase System 10 or System 11, your SQL.INI file should have these sections:

```
[win32client.dll]
    comdll = sqlsyb32
```

```
[sybgtwy]
```

```
; see p. 6-17 of Ch. 6, Connecting to Sybase, in "Connecting Centura Objects
; to Databases" book.
```

Known Problems with SQLRouter/Sybase

IMAGE and TEXT input bind variables

Image and text input bind variables cannot be supported due to limitations within Sybase Client Library software. Please refer to the topic SQL Server Restrictions, in Sybase's Open Client Client-Library/C Reference Manual. Use SybWriteText to insert/update image or text data. For details on how to write and retrieve image and text data, see page 6-38 of Chapter 6, Connecting to Sybase, in Connecting Centura Objects to Databases.

"With cursor" extension

Defect #57020: When using the "with cursor" clause, which instructs the SQLRouter/Sybase to declare a cursor for stored procedures, input parameters to the stored procedure, if any, must be passed using bind variables. Input parameters to stored procedures cannot be passed using literals when using the "with cursor" clause. Please refer to the topic Using cursors with stored procedures, on page 6-29 in Chapter 6, Connecting to Sybase, in the book, Connecting Centura Objects to Databases, for details on the "with cursor" clause.

Raiserrors

Raiserrors are user-defined errors, typically added to stored procedures. They are considered informational and non-fatal. SQLRouter/Sybase returns the raiserror, if the error number used in the raiserror statement is greater than 17,000. For example, you could create the stored procedure as:

```
SqlPrepareAndExecute(hSql,
    'create procedure showtable_sp @tablename varchar(18) as
      if not exists( select name from sysobjects
        where name = @tablename)
        begin
          raiserror 17999 "Table %1! not found.", @tablename
        end
      else
        begin
          select sysobjects.name, type, crdate, indid
            from sysindexes, sysobjects
            where sysobjects.name = @tablename
            and sysobject.id = sysindexes.id
        end
```

You may execute the stored procedure as:

```
SybExecuteProc(hSql, 'exec showtable_sp table1', ':df1, :df2, :df3, :df4')
```

If the table *table1* does not exist, you will get the raiserror. Alternatively, you could use

```
When SqlError
    Call SqlGetError(hSql, nError, sError)
    Call SybGetServerMsg(hSql, l, nMsg, sMsg, nSeverity)
    Call SalMessageBox(sMsg, 'SqlError', MB_Ok)
    Return TRUE
```

before the SybExecuteProc statement. When you get the raiserror, the raiserror text is displayed.

SQLRouter/Oracle

The following section contains information about SQLRouter/Oracle.

SQLRouter/Oracle Compatibility

This version of SQLRouter/Oracle is certified with:

Centura Software

- Centura Team Developer 1.0
- SQL/API 6.1.0

Oracle Software

- Oracle 7.1 on OS/2
- Oracle 7.1 on SunOS
- Oracle 7.2 on Windows 95 and Windows NT
- SQL*Net 1.1.x and SQL*Net 2.x

Operating Systems

- Windows NT 3.51
- Windows 95

Network Software

- TCP/IP, SPX/IPX

Configuring SQL.INI for Oracle

The COMDLL for Centura Team Developer to use the SQLRouter/Oracle is SQLORA32. Please specify this in the [win32client.dll] section in the SQL.INI file as follows:

```
[win32client.dll]
comdll=SQLORA32
```

The router specifies entries, like **longbuffer**, in sql.ini file under **[ORAGTWY]** section. For example,

```
[ORAGTWY]
longbuffer=64000
```

To enable the Oracle router log, please specify the following in the SQL.INI:

```
[win32client.ora32]
log=c:\temp\ora32.log
```

SQL*Net 2.1

The following Oracle DLL has been added to this release:OCIW32.DLL

Presently there are three versions of the SQL*Net DLL in use:

- ORA7NT.DLL - Oracle 7.0.x
- ORANT71.DLL - Oracle 7.1.x
- ORA72.DLL - Oracle 7.2.x

To make an OCI application work with any of the above versions, Oracle has released OCIW32.DLL. We certified Oracle connectivity with OCIW32.DLL. Before you attempt to use the Oracle router add the following entry to the Windows 95 or Windows NT Registry under HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE section:

```
ORAOCI=ORA7NT.DLL ; For Oracle 7.0
ORAOCI=ORANT71.DLL ; For Oracle 7.1
ORAOCI=ORA72.DLL ; For Oracle 7.2
```

To connect to Personal Oracle7 on Windows 95, add an entry in the Windows 95 Registry under

HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE section as follows:

```
LOCAL = 2:
```

In your sql.ini file add the following line:

```
[ORAGTWTY]
remotedbname=ora95,LOCAL
```

Known problems with Oracle connectivity

- O.S. authentication of users when connecting to Oracle with the dummy/dummy userid/password combination fails.
- Oracle 20936 error and 145 error when trying to use SUBSTITUTE=@VALUE,TO_NUMBER in sql.ini and issue select @VALUE(column) from table.
- Error - 'SQL Error 25533 not found in ERROR.SQL file. 00133 SQL IVE ERROR -2553 NOT FOUND in ERROR.SQL' returned, while trying to execute a Stored Procedure with 254 arguments.

SQLRouter/ODBC

The following section contains information about SQLRouter/ODBC.

ODBC connectivity

SQLRouter/ODBC provides connectivity between Centura software and databases for which an ODBC driver is available. SQLRouter/ODBC works with ODBC drivers that implement Microsoft ODBC specification version 2.0 and above.

These notes list the drivers which have been certified by Centura Software Corporation. Separately, these notes list the drivers distributed with this release of Centura Team Developer. If a driver you want to use is not distributed by Centura, you must obtain it yourself from an appropriate vendor. Any driver to be used with SQLRouter/ODBC must conform to at least level 1 of the ODBC API. All drivers certified by Centura meet this requirement.

Note: This release of Centura Team Developer allows you to certify drivers that Centura Software Corporation has not certified using DCC (Database Connectivity Configuration). Read the Team Developer online help for more details on DCC.

Once you have installed the drivers you need, verify the connection between your client and a data source before starting to use Centura Team Developer.

SQLRouter/ODBC Compatibility

This release has been tested and certified with the following products:

Operating Systems

- Microsoft Windows 95
- Windows NT 3.51

Centura software

- Centura Team Developer 1.0
- SQL/API 6.1.0
- dBASE ODBC driver - Centura dBASE driver 2.11 (GPDBF08.DLL) which supports:dBASE III, dBASE IV and dBASE V filesClipper filesFoxPro and FoxBase files

Microsoft software

- Microsoft ODBC 2.50.2923 Driver Manager

ODBC drivers certified

- dBASE ODBC driver from Centura v2.11
- Microsoft Access driver from Microsoft v3.40.2204.
- Microsoft SQL Server 6.0 driver from Microsoft v6.00.126.

- Oracle 7 driver from Visigenic v1.10.0004.

Configuring SQL.INI for ODBC

Please note that the following should be part of your SQL.INI file to connect to any ODBC data source successfully.

The **win32client.dll** section should have the following :

```
[win32client.dll]
comdll = sqlodb32
```

There should be a section **odbcrt** with optional keywords in it.

```
[odbcrt]
```

Please refer to the book "Connecting Centura Objects To Databases" for more details on all the keywords.

Truncation of error messages

At runtime, messages longer than a Centura-defined limit (roughly 220 bytes) are truncated. Because multiple messages from a single ODBC function call are always concatenated into a single message before being returned to the application, any message that starts beyond the limit is lost.

Informational and warning messages ignored

SQLRouter/ODBC ignores the informational and warning messages available when an ODBC function returns SQL_SUCCESS_WITH_INFO.

Table name maximum length

In Centura Team Developer, tables with names longer than 18 characters cannot be accessed. Any attempt to access such a table results in the error "Delimited identifier is too long".

Visigenic Oracle Driver

Requirements

The current version of Visigenic Oracle 7 ODBC driver will only work with Oracle SQL*Net TCP/IP v1.1 running on Windows 95. The following files are needed in a separate directory from Centura, e.g. c:\orant:

```
corent23.dll -- 1/5/95 2:12pm 775,168bytes
msvcrt10.dll -- 7/23/93 5:31pm 210,944bytes
orant71.dll -- 2/27/95 10:47pm 251,904bytes
sqltcpt.dll -- 2/13/95 12:31pm 25,600bytes
```

The search path in autoexec.bat needs to be added with the path of c:\orant, in this example.

Another file is needed in the \windows\system directory:

```
ct13d32.dll : ver. 2.31.0000
```

An entry is needed in file 'services', which is in \windows directory:

```
orasrv          9999/tcp
```

9999 is the port number on which the server is listening.

Known Problems with SQLRouter/ODBC

Oracle PL/SQL support deferred

The current Visigenic Oracle 7 ODBC driver does not support PL/SQL stored procedures. Please contact Visigenic for the update to the driver which will support PL/SQL stored procedures.

Positioned Updates not supported for Oracle

The current Visigenic Oracle 7 ODBC driver supports only static ODBC cursors. Because of the problems involved in using static ODBC cursors, Positioned Updates (update/delete ... where current of ...) are not supported in this release.

StarWare DB2/400 ODBC Driver

Contact Information

The DB2/400 ODBC driver is certified with Centura Team Developer by StarWare Corporation. If you need to connect to DB2/400 from Centura Team Developer, you need to purchase the ODBC driver from StarWare. StarWare can be contacted at:

2150 Shattuck Ave, Suite 600
Berkeley, CA 94704
Phone: +1(510) 704-2000

Table and column name lengths in DB2/400

The DB2/400 database engine supports table names up to 128 characters and column names up to 30 characters. However, Centura Team Developer limits both table names and column names to 32 characters. This limitation allows maximum-length column names, but does not allow table names longer than 32 characters.

SQLRouter/Microsoft SQL Server 6

The following section contains information about SQLRouter/Microsoft SQL Server 6.

SQLRouter/ Microsoft SQL Server.

SQLRouter/ Microsoft SQL Server supports Microsoft SQL Server 6, using its native call-level interface-ODBC.

SQLRouter/Microsoft SQL Server Compatibility

This release of the SQLRouter/Microsoft SQL Server has been tested and certified with the following products:

Operating Systems

- Microsoft Windows 95
- Windows NT 3.51

Centura software

- Centura Team Developer 1.0
- SQL/API 6.1.0

Microsoft software

- Microsoft ODBC 2.50.2923 Driver Manager

Microsoft SQL Server software

- Microsoft SQL Server 6.0.0126 (Service Pack 2) on Windows NT

Microsoft SQL Server client software

- Microsoft SQL Server ODBC driver version 6.0.0126, ODBC version 2.5

Known Problems with SQLRouter/Microsoft SQL Server 6

Information and warning messages ignored

SQLRouter/Microsoft SQL Server ignores all information and warning messages associated with ODBC calls that return SQL_SUCCESS_WITH_INFO.

Displaying Indexes with Database Explorer

The Database Explorer is not able to display a list of indexes or index information for SQL Server databases before version 6.5 of SQL Server. There is no known work around for this problem. Index viewing will work correctly in SQL Server 6.5 or later databases.

IDENTITY column not listed in QuickObjects Key Fields list box

QuickObjects uses an IDENTITY column (if available) to identify unique rows in a table. However, unless the column is used to create a unique index, it is not listed in the Key Fields list box.

Maximum MONEY value that can be entered

The Microsoft SQL Server money data type is mapped on the client to a decimal data type that has a maximum precision of 15 decimal digits; this is less than the maximum precision of the money data type.

Positioned updates do not work

The Microsoft ODBC driver reports an error when your application attempts to do a positioned update

(CURRENT OF *cursor*), even if result set mode has been enabled and there exists a unique index for the table being updated. The error message states "FOR UPDATE clause allowed only for DECLARE CURSOR", and the attempt to update is unsuccessful. (Microsoft has been advised of this problem.)

NOTE: This problem has been fixed in the Service Pack 1. You must have the SQL Server Service Pack 1 for the positioned updates to work. To obtain this pack, please contact Microsoft.

Use Microsoft SQL Server utilities to execute sp_droplogin

The Microsoft ODBC driver reports an error when you attempt to execute **sp_droplogin**. The error message states:

This login is the owner of 11 Task(s). You must drop or re-assign these Tasks before the login can be dropped.

Use Microsoft SQL Server utilities (such as **isql/w**) to execute **sp_droplogin**. (This problem has been reported to Microsoft.)

Avoid multiple connections without autocommit

If two applications running on a single client attempt to access the same table, or a single application attempts to access the same table through two different data sources, and autocommit is off, the error "Communication link failure" is generated and the applications abort -- no recovery is possible.

Data sometimes cannot be accessed in result set mode

Assume a table that contains columns, one or more of which are of type LONG (that is, TEXT or IMAGE data) and others which are not. If you insert NULLs into the LONG column(s) and non-NULL data into the others, then attempt to view the data with result set mode enabled (scrollable cursors), the non-NULL data cannot be retrieved. (This is a problem with the Microsoft ODBC driver for SQL Server, and has been reported.)

Note: This problem does not occur when result set mode is disabled (forward-only cursors). However, you cannot disable result set mode in Quest. To correctly view a table in Quest with LONG columns that have NULLs, restrict your view of the table to the non-LONG columns.

Decimal or numeric columns where precision equals scale

In Database Explorer, you cannot insert or update a numeric or decimal column where the precision equals the scale with a value of zero. In Centura Team Developer, you cannot insert or update a numeric or decimal column where the precision equals the scale using a bind variable with a value of zero. If you attempt to do so, you get the error "Numeric value out of range", and the operation fails. (This is a problem with the Microsoft ODBC driver for SQL Server, and has been reported.)

Default autocommit behavior

The default Autocommit behavior for Microsoft SQL Server is ON. But, if you make a call to SqlGetParameter (... DBP_AUTOCOMMIT...) it will falsely return OFF.

Index Operations not supported for Microsoft SQL Server

Index related operations will not work in Database Explorer when connected to Microsoft SQL Server. This is a restriction of the server and there is no workaround.

SQLRouter/Informix

The following section contains information about SQLRouter/Informix.

SQLRouter/Informix Compatibility

This version of SQLRouter/Informix is certified with:

Operating Systems

- Microsoft Windows 95
- Windows NT 3.51

Centura software

- Centura Team Developer 1.0
- SQL/API 6.1.0

Informix software

- Informix 5.01
- Informix 6.0 and 7.1 (Partial)
- Informix-Net TCP/IP 5.01 TE 1 for Windows NT
- Informix-Star TCP/IP 5.01 Sun OS 4.1.2
- Informix-SE 5.01 Sun OS 4.1.2
- Informix-ONLINE 5.01 Sun OS 4.1.2
- Informix-SE 6.0 Solaris 2.3
- Informix-ONLINE 6.0 Solaris 2.3
- Informix-SE 7.1 Solaris 2.4
- Informix-ONLINE 7.1 Solaris 2.4

Network Software

- TCP/IP

Restrictions of SQLRouter/Informix

The following are restrictions of SQLRouter/Informix.

Windows 95 clients not supported

Windows 95 client connectivity to Informix is not supported with this release. Only Windows NT clients are supported.

Configuring SQL.INI for Informix

The **COMDLL** for Centura Team Developer to use the SQLRouter/Informix is **SQLIRT32**. Please specify this in the **[win32client.dll]** section in the SQL.INI file as follows:

```
[win32client.dll]
comdll=SQLIRT32
```

The router specifies entries, like **longbuffer**, in sql.ini file should be under section **[INFOGTWY]**. For example,

```
[ INFOGTWY]
longbuffer=64000
```

To enable the Informix router log, please specify the following in the SQL.INI:

```
[win32client.irt32]
log=c:\temp\irt32.log
```

SQLBase Database Engine

The following sections cover migration issues from SQLBase 5.x and 6.0.x to SQLBase 6.1.

Automatic conversion of SQLBase 6.0 or 6.0.1 databases.

An existing SQLBase 6.0 or 6.0.1 database is automatically converted to 6.1 format when you connect to it. We recommend that you UNLOAD your earlier 6.x database to a file before you connect to a 6.x database.

Migrating 5.x databases to SQLBase 6.1

Run the SQLBase Conversion Utility to convert local SQLBase 5.x databases to SQLBase 6.1 format.

Viewing load and unload messages

In SQLBase 6.x, load and unload messages are no longer logged on the client and server screens. If you want to display all the load and unload messages printed in the message logs (which could have performance implications), set the message level for the Process Activity display screen to 1 (or higher; level 1 is the lowest level at which the messages appear). For more information, refer to "Process Activity display screen (F2)" in Chapter 5 of the *Database Administrator's Guide*.

SQLTalk USE command

If you had batch files that you used with SQLBase 5.2's WinTalk, be aware that the syntax of the USE command is different between WinTalk and SQLBase 6.x's SQLTalk.

WinTalk supported USE <database_name>.

In SQLBase 6.x SQLTalk, the syntax is USE <cursor_number>, as documented in the *SQLTalk Command Reference*. You cannot execute a USE <database_name> statement.

Default for GROUPCOMMIT keyword

The default value for the GROUPCOMMIT keyword in SQL.INI has changed between SQLBase 5.2 and SQLBase 6.x. In SQLBase 6.x the default is now 1.

In SQLBase 5.x, the GROUPCOMMIT default was 10 for NetWare, and 2 for other platforms.

GROUPCOMMIT specifies the maximum number of COMMITs that SQLBase groups together before physically writing them to disk.

Refer to the *Database Administrators Guide* for an explanation of the GROUPCOMMIT keyword, and information on where to configure the keyword in SQL.INI.

Change to auto-launch behavior

For both Windows 95 and NT, auto-launch behavior of the server under anonymous pipes has changed:

- In SQLBase, 6.0.1 sqlapipe.dll automatically launches DBNTRSRV.EXE (by relying on the path being set so that the operating system can locate the executable).
- In SQLBase 6.1, sqlapipe.dll tries to launch DBNT1SV.EXE using this search order:
 - a. The current directory.
 - b. The system directory.
 - c. On Windows NT only: The 16-bit system directory.
 - d. The Windows directory.
 - e. The directories listed in the PATH environment variable.

If this fails to locate DBNT1SV.EXE, SQLBase looks up the registry entry. If it finds an entry, it looks for the default value and uses the path specified there to launch the executable. This is similar to the mechanism used by the Windows 95 RUN command.

Specific migration and compatibility issues for users upgrading from SQLBase 6.0.0

The following information applies ONLY to users upgrading from SQLBase 6.0.0 (*not* 6.0.1) to SQLBase 6.1.

If you never installed SQLBase 6.0.0, you may disregard this section..

If you have already upgraded from SQLBase 6.0.0 to SQLBase 6.0.1, you were notified of these issues in the 6.0.1 release notes. No additional action is required to upgrade to 6.1 and you may disregard this section.

Compatibility issues for triggers created in SQLBase 6.0.0

SQLBase version 6.0.1 introduced the ability to use triggers to set column values. This feature can cause compatibility problems with triggers or triggered stored procedures created in SQLBase 6.0.0 (if you recompile them in 6.0.1 or higher and the procedure uses receive parameters).

For more information on this feature, refer to page 2-12 of the *SQLBase New Features Guide*.

Certain SQLBase 6.0.0 stored procedures must be recompiled

If you have stored procedures that use a bind variable multiple times in a SQL statement in SQLBase 6.0.0, you must recompile the stored procedures in order to correct a bug (#21341) that occurred in SQLBase 6.0.0.

The fix for the bug was introduced in SQLBase 6.0.1 (the fix is also in SQLBase 6.1). You must recompile affected procedures in order to take advantage of the fix.

For example, the following statement within a stored procedure is legitimate but needs to be recompiled for SQLBase 6.0.1 or later. SQLBase 6.0.0 handled the statement incorrectly because the same variable is used multiple times in the update (nNewQty and nNewStatus are used multiple times).

```
call SsqlImmediate( \
'update part set \
qty_on_hand = qty_on_hand + :nNewQty, \
qty_available_iss = qty_available_iss \
+ @if(@exact(:sNewStatus, 'A'), :nNewQty, 0), \
qty_available_mrp = qty_available_mrp \
+ @if(@exact(:sNewStatus, 'U'), 0, :nNewQty2) \
where id = :sNewPartID')
```

In contrast, if all bind variables in the update statement were unique, SQLBase 6.0.0 handled the statement correctly.

Note that this situation also applies to procedures used by triggers. That is, if a trigger's inline procedure uses a bind variable multiple times, you must recompile it.

Change in behavior of SQLBase 6.0.0 null passed to a stored procedure

The behavior of a null value passed into a stored procedure has changed from SQLBase 6.0.0. In SQLBase 6.0.1 and higher, a null value passed into a stored procedure remains null. Previously, when a null was passed into a procedure, SQLBase assigned a default value to the parameter. Now ONLY uninitialized local variables assume default values, and all parameters passed into a procedure keep the values that were passed in, whether null or not null.

SQLBase 6.1 Certification Matrix

The following tables illustrate the test environments that have been used to certify SQLBase using various operating systems and network protocols. This Product Certification table reflects combinations of clients, servers, and network protocols that passed our most rigorous tests. However, there are other combinations of clients, servers, and network protocols that are supported.

Underneath each server and client entry in parenthesis will be the Centura COMMDLL associated. Please refer to the section below that describe the operating environment in more detail.

The network used was 802.3/Ethernet based.

Server Operating Systems	Database Server Column Format	Network Protocol	Client Operating System
Windows 95 (SQLAPIPE.DLL)	FAT	Anonymous Pipes	Windows 95 SQLAPIPE.DLL)
Windows NT 3.51 (SQLAPIPE.DLL)	NTFS	Anonymous Pipes	Windows NT 3.5.1 SQLAPIPE.DLL)
Windows NT 3.5.1 (SQLWS32.DLL)	NTFS	TCP/IP	Windows NT 3.5.1 SQLWS32.DLL)

Note: Microsoft Windows 95 and Windows NT (Server and Workstation) are not able to support SQLBase as a server listening on IPX/SPX. Each operating system environment listed above is shown in more detail below. During testing we focused on one protocol at a time. However, all protocols were installed and loaded at the same time. Concurrent use of all protocols is supported.

32-bit SPX Windows NT client support

See the “Up-to-date” topics in the Team Developer help file for 32-bit SPX Windows 95 client support.

32-bit SPX support for Windows NT is also available and requires Microsoft’s implementation of WINSOCK 1.1 along with Microsoft’s Client for NetWare**.

To enable 32-bit SPX support for Windows NT, specify “comdll=sqlwsspx” in the [win32client.dll] section of SQL.INI.

**Note that currently Novell’s Client32 for NetWare operates on Windows 95 only, and not Windows NT. On the other hand, Microsoft’s Winsockets implementation currently only works with Windows NT and not Windows 95, due to a Microsoft bug in their Win95 Winsockets implementation.)

ALTER TRIGGER command

This augments the “Up-to-Date” topic in the Team Developer help. The parameters for the ALTER TRIGGER command are:

- * *Triggername* - This is the name of the trigger to be disabled or enabled. The trigger name can be qualified by prefixing the creator’s name. If no creator’s name is specified, the current user is assumed to be the creator.
- * ENABLE - Enables the trigger
- * DISABLE - Disables the trigger

Change SYSADM and SYSSQL passwords

SYSADM is the most powerful SQLBase user. SYSADM can see everyone’s password and can set up one or more users with DBA, RESOURCE, or CONNECT authority. SYSSQL is the owner of the server-independent views on the system catalog.

The default password for SYSADM is SYSADM and the default password for SYSSQL is SYSSQL. Both these passwords should be tightly controlled. You should change these passwords soon after installation to prevent unauthorized access to the database. For more information on both security and the system catalogs, refer to *the Database Administrator’s Guide*.

Changes to sort behavior for national language settings

Changes have been made to Sort behavior. For an English setting (the default), SQLBase uses the Postman Sort algorithm for sorting. However, when it detects a national language setting in COUNTRY.SQL, SQLBase switches to an alternative sorting algorithm to deal with non-English ASCII/ANSI sort orders. This detection is done automatically and you are not required to do anything to enable this behavior. In this mode the keyword *sortbuffer* is irrelevant to the alternative sort algorithm and is ignored.

SQL.INI settings for users of non-English code page

If you use a code page other than the standard English code page, you must specify the COUNTRY keyword in BOTH the [sqltalk] and [sqlapiw] sections of the SQL.INI. If you do not specify the keyword in each section, SQLTalk may load the incorrect code page.

For example, the following settings in SQL.INI will result in correct code page execution:

```
[sqltalk]
country=japan
```

```
[sqlapiw]
country=japan
```

For more information on COUNTRY.SQL and the code page, refer to page 11-3 in the "National Language Support" chapter of the Database Administrator's Guide.

SERVERPATH keyword semantics and the use of the wildcard character

The serverpath keyword in the client's TCP/IP section in the SQL.INI tells the client how to find the server and the databases to which it needs to connect. It uses the host server's name or IP address to indicate the location of the server on the network.

A client may need to connect to more than one server using TCP/IP. To allow for this, you can specify more than one SERVERPATH keyword in the client section. Each keyword defines the location of the server and the databases available on it. The wildcard '*' can be used to specify any database in the keyword. This means that if you have multiple databases on the host server, you need not list each one of them by name. You can use * to indicate all databases on that server.

There are a few restrictions apply to the keyword as follows:

- 1) SQLBase does not allow the wildcard character in more than one SERVERPATH keyword entry. That is, if you have two or more SERVERPATH keywords in your Winsock TCP/IP section, then only one of them can use the '*' wildcard character.
- 2) If a wildcard is used to denote the database name using the serverpath keyword, then it should be used alone and not as a template. That is, the following serverpath statements are incorrect:

```
SERVERPATH=server1,host1/DEMO*  
SERVERPATH=server1,host1/*EMO  
SERVERPATH=server1,host1/D*MO
```

The correct usage is

```
SERVERPATH=server1,host1/*
```

- 3) The wildcard should not be used for servername, hostname, or port number. For example, the following statements are incorrect:

```
SERVERPATH=*,host1/DEMO  
SERVERPATH=server*,host1/DEMO  
SERVERPATH=*ver,host1/DEMO  
or  
SERVERPATH=server1,*/DEMO  
SERVERPATH=server1,*ost/DEMO  
SERVERPATH=server1,host*/DEMO  
or  
SERVERPATH=server1,host1,*/DEMO.
```

- 4) The serverpath keyword should not contain any embedded spaces between the servername to the database name. If a space character is encountered then the string is terminated at the first occurrence of the space, and an error message is returned for missing parameters.

- 5) The wildcard character poses a problem when you are using more than one protocol DLL with your client and the TCP/IP protocol is the first protocol of choice. In this scenario, if you specify a wildcard in your serverpath keyword, SQLBase will never use the other protocols to search for the specified database.

For instance, consider the following in your client SQL.INI :

```
[win32client.dll]
```

```
comdll=sqlws32
comdll=sqlapipe
```

```
[win32client.ws32]
serverpath=server1, 198.206.11.20/demo1, *
```

If you have a local database DBLOCAL using Anonymous Pipes, an attempt to connect to it using 'CONNECT DBLOCAL 1' will fail. This is because SQLBase will first use Windows Sockets to locate the database and will match the name with the wildcard in the serverpath keyword. Since it will not find the database on server1, it will return the error 'Cannot open database'. It will not attempt to use Anonymous Pipes to locate the database. To avoid this problem, you should always specify TCP/IP as the last protocol. That is, in this case you should specify the following:

```
[win32client.dll]
comdll=sqlapipe
comdll=sqlws32
```

A future release of SQLBase will resolve some of these issues. However, for now, we strongly recommend that you avoid using the wildcard character when possible. By explicitly listing the databases on the server in the serverpath keyword statement, you can avoid these problems altogether.

Timer events currently enabled only through SQLConsole

SQLBase does not support creating and executing timer events on the server in this release. This supersedes all documentation about events in the SQLBase 6 documentation set. SQLBase Events will be enabled in a future SQLBase release.

You can, however, still create and schedule events on the client side through SQLConsole's Scheduling Manager feature. The Scheduling Manager allows you to execute scripts or other programs, as well as schedule backup and maintenance events. The main difference between using SQLBase events and SQLConsole events is that you must maintain a client connection to the machine where SQLConsole is running to execute your SQLConsole event.

The SQLConsole Scheduling Manager is documented in Chapter 10 of the *SQLConsole Guide*.

Memory requirements for Window 95 and NT

SQLBase is able to operate under the minimum recommended memory requirements for Windows 95 (8 MB) and Windows NT (12 MB).

The following table shows the results of our internal tests running the unlimited user SQLBase Server on Windows NT (DBNTRV.EXE). These figures show the memory required to bring up the server and issue connections (these results apply only to Windows NT). The figures do not include space allocated for the cache.

Activity	KB Required
To bring up base server	1972 (KB)
First connect	313 (additional KB)
Additional connect	53 (additional KB)

Known Problems with SQLBase 6.1

This section describes known problems and limitations of SQLBase 6.1:

Defect ID 52332

Users and Tables nodes in the Replication Management System (RMS) do not have expandable indicators set upon entry to RMS. After expanding them, they are set.

Defect ID 52515

The Window menu in SQLConsole does not always list all the open windows.

Defect ID 53087

Painting problem in SQLTalk. When you move certain dialog boxes or windows over the tool bar and then close the dialog box or window, the tool bar icons that were under the dialog box or window disappear. This happens in Win32 but not in Win16. Workaround: If SQLTalk is minimized and then maximized, the tool bar is fine.

Defect ID 53455

SQLTalk: Using the ReplaceAll command (under the Edit menu) causes an infinite loop if the replace string you specify is two instances of the find string. For example, Find [omed] - ReplaceAll [omedomed] causes an infinite loop.

Defect ID 53590

Under Replication for horizontal partitioning, SQLConsole accepts any invalid SQL statement in the WHERE clause. Later the Sync Agent fails to Sync but it is difficult to find out why it failed; there is no information pointing to the bad WHERE clause

Defect ID 53786

Breach of security under Replication: after you create your replication sets any user can start replication. Security is there only the first time and not after that.

Defect ID 54742

SQLBase user with DBA authority cannot see any tables owned by others in the Tables tree view or in Define Query.

Defect ID 56077

Using Anonymous Pipes protocol on Windows NT when an application is running as an "NT service" does not work. The workaround is to configure the client and server to use Winsock TCP/IP.

Defect ID 56197

On Windows 95 only, we have seen cases where a user cannot use the file browser dialog under the Audit Manager to locate an audit file on a drive different from the drive on which the server is running. If you experience this problem, it is possible to cause the audit file to be created on a different drive, but it will not be viewable using the Audit Viewer feature of the Audit Manager. To avoid this problem, make sure that audit files are stored on the same drive as the server.

The same problem makes it impossible to use the file browser dialog in database load and unload operations if the unload file is on the server and in a drive different from the drive the server is in. In these cases, full functionality is available because the relevant information can be typed in.

Defect ID 56575

By default, SQLTalk shows 10 characters (including the decimal) for a FLOAT. However, with defect 56575, SQLTalk truncates FLOAT data from non-SQLBase backends. The problem is most troublesome with backends that store the data in scientific notation. As a workaround, use SQLTalk's COLUMN WIDTH option to correct the defective behavior (for example, set "COLUMN WIDTH 1 20" to force column 1 of a SELECT statement to show 20 characters).

Defect ID 56829

If you launch the 16-bit SQLTalk and then later try to launch the 32-bit version, you will not be able to do so. The workaround is to reboot and then launch the 32-bit version of SQLTalk.

Defect ID 57503

SQLConsole: When using the Statistics tool for a database to view its locks in a situation where available memory is very low (less than 5 MB) and where the database is very large (many tables and/or columns) SQLConsole will sometimes terminate with a general protection fault. If this happens, make sure that as much memory is available to SQLConsole as possible by closing all other running applications, or by using a bigger machine.

Centura Ranger and SQLBase Replication

The following topics contain information about SQLBase and Centura Ranger Replication.

After you install SQLBase and Ranger Replication

Change the SYSREP password. For SQLBase manager (not subscriber) databases, you must change the SYSREP password before you run the Replication Management System (RMS) for the first time. You can do this with SQLTalk or with a SQLConsole workspace. After you successfully change the default password, it is stored in the SQLCON.INI file, which is then used on subsequent entries so that you are no longer prompted for the password.

To change the password with SQLConsole, put a check in the System Objects field on the SQLConsole toolbar, expand the appropriate database node to view the Users node under it, right-click the Users node, select **Alter**, and change the password. To change the password from SQLTalk, use the GRANT CONNECT TO command, which overrides any existing password. Do not use the ALTER PASSWORD command since it requires you to know the current password.)

Create the SYSREP user manually. If you use the Ranger-to-Oracle package, you must create the SYSREP user manually, with DBA authority. Otherwise, if you do not use Ranger-to-Oracle, SYSREP is automatically created.

Additional privileges required if using Oracle. If you use Oracle manager database support, you must use the WITH GRANT OPTION clause of the GRANT command to allow SYSREP to grant privileges on objects. You must do this before you create a replication set, otherwise Oracle will not permit SYSREP to assign required privileges to subscribers.

For example, assume you want to include the following two tables in a replication set:

A.T1
B.T1

Assume also that you have the following two subscribers:

User1
User2

Before a replication set is created, table owners A and B must use the WITH GRANT OPTION clause to grant privileges to SYSREP. This will allow SYSREP to give privileges on T1 and T2 to subscribers User1 and User2, which is necessary to replication.

This information supplements the general information on privileges and security on page 2-13 of the *Replicating Data with SQLBase and Centura Ranger* manual.

Important new features in Replication

Please note the following new features for this release of SQLBase and Centura Ranger Replication. These features are described in more detail in the following sections. Also refer to the SQLBase Database Engine and SQLConsole topics in these release notes for additional information.

- Automatic conversion of 6.0 and 6.0.1 databases to 6.1 upon upgrade to a 6.1 server.
- Oracle Manager database support provided in the Centura Ranger-to-Oracle package.
- New packaging and installation program.
- All components installed now to \CENTURA directory.
- RSA is available both as a 16-bit and 32-bit client.
- SQLConsole 6.1 is now available (on Win32 platforms only).
- Commit logging support (refer to the *SQLBase New Features Guide* for details).

- Database sequence object (refer to *SQLBase New Features Guide* for details).
- ALTER TRIGGER command (refer to this topic in these release notes for details.)
- Simulated replication example and sample application are provided online.

Online replication example and sample application

An online “screen-cam” simulation of a replication example is provided in this release. It shows the creation of an example replication set using SQLConsole. Note that this simulation software does not actually use a database, the RMS, or RSA, but rather it is a record of their use. Also provided is the database unload file and application for the screen cam demo.

Also included is a running sample application which uses the replication set created at the beginning of this simulation, together with the synchronization of changes. The samples of the RSA OCX use SQLWindows, Centura Builder, Visual Basic (both 16- and 32-bit) and C++.

The screen cam simulation provides pop-ups with information about what is going on as the demo progresses. Use the ENTER key after reading each pop-up so that you continue to progress through the simulation.

The simulation is located in the root directory of the CD under SAMPLES\DEMO. Run DEMORUN.EXE to view the demo.

Temporary file created by RSA

This augments the “Up-to-date” information in the Team Developer Help File.

User Scripts. Note that user scripts are limited to 32K for 16-bit platforms.

Catalog tables. You should not attempt to replicate any database system catalog tables. Doing so will cause run time errors during synchronization. These errors will be related to security violation on the system catalog (or a related error).

Manual reinitialization. If a subscriber wants a site reinitialized as the result of a change in a manager replication set attribute (such as a Before Init script, or Bind values for row filtering changes) where the RMS does not update the replication set version number automatically (which causes subscriber reinitialization on the next synchronization), the subscriber can use the Force-Reinit option when synchronizing.

Timeout values. The RSA currently uses the server time-out value for all SQL lock time-out operations. For SQLBase, this value is set for the server via the SQL.INI file in the appropriate server section. The default value is 5 minutes. The administrator may want to adjust this value according to whether longer or shorter timeouts are desired.

Timestamps recorded in the SYSREPSYNCLOG table. The SYSREPSYNCLOG table contains begin, end, and restart time columns, where timing information is recorded by the RSA during each phase of synchronization. Note that phases 1, 7, and 8 (subscriber phases) record timestamps from the subscriber PC, while phases 3, 4, 5 (manager phases) record timestamps from the machine where the manager database resides. Therefore if the clocks on the two machines are not synchronized, the time information between subscriber and manager phases may be misleading. SYSREPSYNCLOG time stamps are used for reporting or for performance measurement purposes.

Using the same replication schema for multiple manager sites. For testing or deployment purposes, some administrators may try to create a replication set at one manager database using the RMS, and then unload the database to capture the replication meta data schema and load that schema at one or more other manager databases. If this technique is used, you must modify database-specific attributes that are stored in the replication meta data. Specifically, the following values should be updated at the target manager database after the schema has been unloaded and before any subscribers have synchronized:

- The DBName column value in the SysRep.SysRepDatabase table should contain the name of the manager

database.

- If this unload/load technique is being used in a production (versus test) environment, the DatabaseID column value must be unique among all other manager databases in the network. This column value is a part of the key in most SysRep meta data tables.

System requirements for Replication

This section describes the Replication system requirements.

Manager database site: SQLBase Server

- SQLBase 6.0 or later.

Note: This product installs SQLBase version 6.1. If you are upgrading to either a 6.0 or 6.0.1 to SQLBase 6.1, all databases are automatically converted to 6.1 format when you connect to them with SQLBase 6.1. (The database schema is converted; user data is not touched.) To keep a copy of your current data for backup, it is recommended that you first UNLOAD your earlier 6.x databases to a file before you connect to them with the 6.1 server.

- Applicable operating system version for your platform.
- Appropriate software and SQL.INI configuration to accept connections from the Replication Synchronization Agent (RSA) and Replication Management System (RMS) clients. See the *SQLBase New Features Guide* and also the *Database Administrator's Guide* for information on connecting to and configuring a Windows and Win32 client.
- If you have a remote database, you also need appropriate communications software and hardware to connect to the manager site, such as a modem.
- If you use the Ranger-to-Oracle package, you need Oracle 7.0 or later, and the applicable operating system version for your platform. For more information, refer to your Oracle installation documentation.

Manager administration site: SQLConsole and the Replication Management System

The RMS is an extension of SQLConsole, and is installed on a client machine of the database server that will act as the manager database. Note that a local SQLBase 32-bit server is installed on this client to support SQLConsole functions. (Keep in mind that SQLConsole 6.1 only runs on 32-bit Windows platforms.)

The RMS machine has the following requirements. Note that the RMS can reside on the same machine as the manager site database and server for Windows NT and Windows 95.

- Win32 environment (Windows NT or Windows 95)
- Centura Team Object Builder deployment files (installed automatically).
- SQLConsole 6.1
- For the Centura Ranger-to-Oracle package: Centura Oracle router software. (This software, which is installed automatically, is needed for the RSA components to connect to the Oracle Server during data synchronization.)
- Appropriate software and SQL.INI configuration to connect to the manager database. See the *SQLBase 6.1 New Features Guide* and also the *Database Administrator's Guide* for information on connecting to and configuring a Win32 client.

Note that you can have the subscriber database residing on the same machine as the manager database. That is, it is possible to do replication between two databases on the same machine.

Subscriber database site: Replication Synchronization Agent (RSA)

The subscriber machine has the following requirements:

- 16- or 32-bit Microsoft platform (such as Windows 3.1, Windows 95, or Windows NT). Note: Currently the RSA works on Windows NT version 3.5.1, but not on version 3.5.0.
- SQLBase 6.1.

- For the Centura Ranger-to-Oracle package: Centura Oracle router software, which is installed automatically, is needed for the RSA components to connect to the Oracle Server during data synchronization.
- Appropriate software and SQL.INI configuration to accept connections from the RSA. See the SQLBase New Features Guide and also the Database Administrator's Guide for information on connecting to and configuring a Windows and Win32 client.
- Remote LAN software such as RAS or Shiva for a detached remote subscriber.
- The RSA software components (installed automatically).

Centura Ranger-to-Oracle Replication

Refer to the SQLConsole help file, the *Replicating Data with SQLBase and Centura Ranger* guide and the *Connecting Centura Objects to Databases* guide if you will be using Centura Ranger-to-Oracle replication.

Known limitations of Oracle replication

The following section lists some of the known limitations in replicating Oracle data with Centura Ranger-to-Oracle.

- Any indexes defined on the *Char/Varchar2* columns of size > 254 characters will not be replicated.
- Any indexes defined on columns of type *Raw/Rowid/MLSLLabel* will not be replicated.
- Differences in granularity may result in reduction of precision and scale of Oracle Number data types.
- Oracle *Date* data types do not support microseconds. These columns are replicated as SQLBase *Timestamp* data types. Such replicated columns on SQLBase may contain microsecond information. But any updates to microsecond information on the subscriber site are lost when replicated back to the Manager site.
- Oracle *Rowids* are replicated as *Char* data types on SQLBase and as such they will not have any special Rowid characteristics on the subscriber site; this is also true for *MLSLLabel* data types.

Known problems with Oracle replication

The following topics describe known problems with Oracle replication.

Oracle Replication: Database connection with SQLTalk scripts

We discourage the use of connect database to the subscriber/manager database operations within the Before- and After- Scripts that are invoked by the RSA through SQLTalk during synchronization. Connect operations may fail due to timeouts or other SQL errors, and these errors will be reported to the RSA. If the StopOnScriptError replication set attribute is on, then any SQLTalk script error reported to the RSA will cause the current synchronization to fail.

Oracle Replication: Inconsistent Endtime column value

The EndTime column value is recorded in the SysRep.SysRepSynclog table at the subscriber site for the Manager Copy phase (Phase Type 5) if it doesn't match the Endtime value for the same row in the manager site. Currently the SysRep.SysRepSyncLog row for Manager Copy is copied to the subscriber site prior to the time when the manager's version of the row is updated with the timestamp for the end of phase. During this copy operation, the timestamp reflecting the copy time is used as the subscriber's EndTime column value. A similar anomaly also exists for the Endtime column value for the Subscriber Copy phase (PhaseType 1); the EndTime value in the subscriber's version of the SysRep.SysRepSyncLog is correct, whereas the manager's version of this row for PhaseType 1 is slightly different.

This inconsistency problem will most likely only affect users who collect timing information from the SYSREP.SYSREPSYNCLOG table for performance measurement purposes. If you collect timing information, you should be sure to collect beginning and end timestamps from the SYSREP.SYSREPSYNCLOG locale to where each phase of replication occurred (e.g. timing information for the Manager Copy phase should always be collected from the manager SYSREP.SYSREPSYNCLOG.)

Oracle Replication: Horizontal partition behavior

When a horizontal partition references non-replicated columns, and an update to that table is done at the

manager site where the before-update value is inside the partition range, and the after-update value is outside the partition range, the "delete" of that row is not propagated to the subscriber during synchronization. (defect ID #54035)

Oracle Replication: Security rights not revoked when subscriber is dropped

Security GRANTS that result from using GRANT USER RIGHTS ON in a replication set definition are not revoked at the time the subscriber is dropped from the replication set. This means a user who is no longer a subscriber will still be able to perform DML operations on replication meta data tables, according to those granted rights. (defect ID #51546)

Oracle Replication: Inconsistent name generation for auxiliary tables

In some instances, the names generated for auxiliary logging tables are not generated consistently between instances of replication set definition through the RMS. In other words, the new delta and remote delta table names that are generated differ from those generated by previous invocations. For example, with one usage, the name suffixes were in ascending numeric sequence corresponding to the alphabetic order of table names (Account =repsdt1, Balance=repsdt2...). With a subsequent usage, the following was generated (Account =repsdt2, Balance=repsdt1). Therefore the consequence is that it throws off test results that look for specific delta and remote delta table names. (defect #53799)

Oracle Replication: Using DDL commands against replication tables

Do NOT run any DDL commands against tables defined in your replication set without making corresponding changes to the replication set definition through the RMS. For example, if you drop a replication table from the manager database, be sure to also drop that table from the replication set definition through the RMS. If an administrator issues a DDL command such as drop, alter or rename against a replicated base table, a run time error will occur during synchronization. The error will most likely be a SQL error reporting that the table couldn't be found or is inconsistent.

Oracle Replication: Manager sites databases that are pre-SQLBase 6.1

Manager sites databases that are pre-SQLBase 6.1 need to create a SYSREP user with DBA privileges. This is necessary because the RMS component requires the user to connect as the SYSREP user before accessing any Replication manager functions. If the SYSREP user doesn't exist at the manager site, the RMS will be unable to connect as this user, and access to the RMS will be disallowed.

Oracle Replication: Dropping tables from replication sets

If you drop a table from a replication set after a subscriber has initialized, the subscriber receives an error when it next tries to synchronize. To clear this problem the subscriber must manually re-initialize by using the Force ReInit RSA option in the Replication Synchronization Agent application during the next synchronization.

Oracle Replication: Performance and indexes

Since this release does not replicate all indexes on a replicated table, you should keep a record of what indexes need to be created at the subscriber site. You can then create an AFTER-INIT script with the relevant CREATE INDEX commands. However, note that this release *will* replicate the unique index for the columns designated as the primary key in the replication table's definition. Other than this index, no other auxiliary indexes are replicated.

Oracle Replication: Updating primary key columns during Manager Apply phase of synchronization

The RSA may fail during the Manager Apply phase of synchronization if the manager base table(s) have referential constraints which prohibit updating primary key columns. The Manager Apply phase updates all replicated columns in the manager base table, including primary key columns, if those columns are defined as updatable at the subscriber database. If you wish to avoid this referential constraint violation on manager-replicated tables whose columns are used as primary keys at the subscriber, you should define those columns as

non-updatable at the subscriber site (which means that no changes to those columns are propagated to the manager site). This is accomplished by using the column properties in the RMS.

Oracle Replication: Replicating manager tables with delete cascade referential integrity constraints

You must use the REFRESH option rather than the DELTAS option, which is the default during replication set creation. If the DELTAS option is specified for the replicated table and delete cascade constraints exist, the RMS will not be able to create the triggers required for DELTAS mode. The following error is reported:

08412 EX2 TGL Delete Cascade constraint already defined on table.

Reason: Cannot create a Delete trigger on a table that has a Delete Cascade constraint defined on it.

This will halt the replication definition process because you cannot create a trigger on a table that has any foreign keys with a delete rule of ON DELETE CASCADE. Since the replication process requires building triggers on the table if the DELTAS replication table option is used, you will have to change the delete rules for any foreign keys defined as such. (Note that this is not a limit specific to the replication manager; you cannot do so anywhere within SQLBase.).

Specifically, you cannot create a delete trigger on a table with a rule of ON DELETE CASCADE nor can you always create an update trigger on a table with a rule of ON DELETE SET NULL (only if the column list referenced in the trigger definition contains a column in the foreign key).

This rule is enforced by SQLBase and conforms to the ANSI SQL standard with respect to the relationship between trigger definitions and referential integrity. By defining your replication table to use REFRESH mode, you avoid the need to create triggers and thereby avoid this restriction altogether. The other alternative is to remove the RI constraints on the manager base table.

Oracle Replication: Using the SQLConsole workspace interface

Using the SQLConsole workspace interface is not always consistent with using SQLConsole's Replication Management System (RMS) interface. For example:

- You cannot right-click on a node in the RMS to access the Explore menu item.
- In a SQLConsole workspace, you can place the cursor anywhere on the line associated with a listbox item and clicking the mouse (either right- or left-clicking) brings it into focus. In the RMS, you must be on the graphic or text item in order for it to come to focus.

Oracle Replication: SYSREP and SYSSQL users

Currently the SYSREP and SYSSQL special users are automatically created by SQLBase, but can be dropped by a user with SYSADM authority. Since SYSREP is a special user identity used with replication, and is assumed to exist, you should never drop this user from either the subscriber or manager database.

At the subscriber database, Ranger replication assumes that the SYSREP user has the password with which it was originally created. Therefore do not change the SYSREP user's password at the subscriber site.

At the manager database, before the RMS (in SQLConsole) is used to manage replication sets, the password of the pre-created SYSREP user should be changed for security reasons. (See the information about changing the SYSREP password earlier in these release notes, under "Change the SYSREP password".) You can accomplish this by the SYSADM user issuing the following statement in SQLTalk:

```
GRANT CONNECT TO SYSREP IDENTIFIED BY newpassword;
```

Oracle Replication: Option for installing the 16-bit RSA

The user has the option to install the 16-bit RSA. Note that currently the 16-bit VB application will run under Windows NT only if you've applied Service Pack 3. (This is a problem related to VB, *not* the 16-bit control.)

SQLConsole

The following topics contain functional information for SQLConsole.

Compatibility with SQLConsole 6.0 Repository

This augments the "Up-to-date" topic in the Team Developer Help file.

At startup, SQLConsole fails to automatically connect to local server. SQLConsole fails to automatically connect to local server. When you first start SQLConsole, it asks whether you want to connect to your local server. Even if you say Yes, SQLConsole fails to connect to the local server (as specified by the servername keyword in SQL.INI). This is a known problem (defect ID 55499). The workaround is to use the Connect Manager to connect to the local server, as documented below:

1. Select **Workspace, Open**.
2. Select DEFAULT, then **OK**.
3. Double-click the Connect Manager node (which is under the Tools node) to open it. (You may first need to double-click the Tools node to expand it.)
4. Click the **New Server Entry** push button. (This is the push button below the word "Servers" in the Connection Manager screen, next to the push button with an "X." Use tools tip help to identify this push button, or, for a picture, refer to the section in the manual on connecting to a server.)
5. Enter the server name in the Server field.
6. If the server has a password, type that in the Password field. If you do not have a server password defined in the SQL.INI file, then any value in the password field will do, including a null password.
7. Click **OK**.
8. Exit and restart SQLConsole to see the new server connection. For more information, refer to the section in the SQLConsole Guide on connecting to a server.

SQLBase 32-Bit ODBC Driver

The following topics contain information about the 32-bit SQLBase ODBC Driver.

Version and Platform Support

Portions of this software are copyright 1991-95 by Intersolv, Inc.

The Centura SQLBase ODBC Driver includes and works with the Microsoft 2.5 Driver Manager.

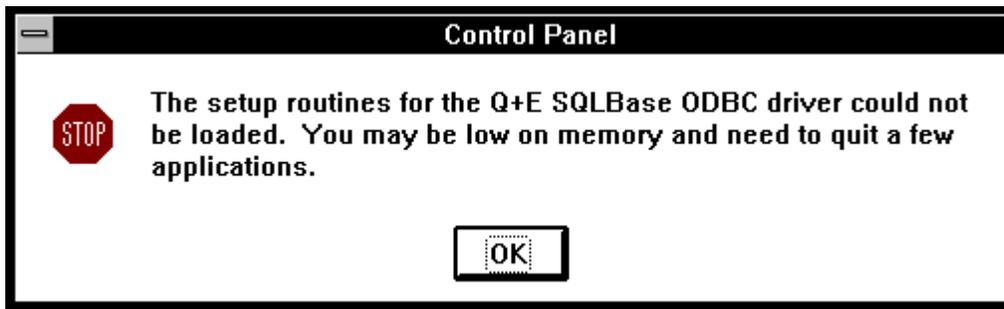
The Centura SQLBase ODBC Driver supports SQLBase version 5.2 and 6.x. The Centura SQLBase ODBC Driver is called GPGUP08.DLL.

System Requirements

The Centura SQLBase ODBC Driver requires either Microsoft Windows 95 or Microsoft Windows NT.

You **must** install Windows Client Support. If you install the SQLBase ODBC driver into a different directory than your Centura client software, you must have the SQLBase client software on your path.

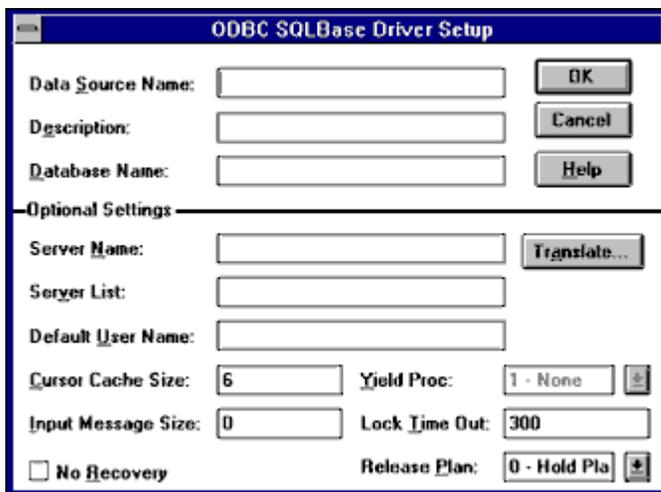
If you attempt to configure a data source and you have not installed Windows Client Support or you do not have Centura client software on your path, you receive the following message:



Configuring Data Sources

To configure a data source, perform the following:

1. Start the ODBC Administrator. A list of data sources appears.
2. Indicate either a new or existing data source, as follows:
 - If you are configuring a new data source, click **Add**. A list of installed drivers appears. Select SQLBase and click **OK**.
 - If you are configuring an existing data source, select the data source name and click **Setup** (in some environments this button is called **Modify**).
 - The Setup dialog box appears.



3. Specify values as follows:

Data Source Name: Identifies a single connection to a database system. You enter any descriptive string. For example, "Accounting" or "SQLBase-Serv1."

Description: An optional long description of a data source name. For example, "My Accounting Database" or "SQLBase on Server number 1."

Database Name: The name of the database system to which you want to connect, or the actual local SQLBase database name.

Server Name: Optionally, identifies the name of the server that services the desired database system. Enter the word LOCAL if you are using the local server.

Translate: Displays the Select Translator dialog box to allow you to perform a translation of your data from one character set to another. Choose the INTERSOLV OEM ANSI translator to translate your data from the IBM PC character set to the ANSI character set.

Server List: A comma-separated list of servers that will appear in the Logon dialog box. Enter LOCAL to add the local server.

Default Logon ID: The default Logon ID used to connect to your database system. Your ODBC application may override this value or you may override this value in the Logon dialog box. A Logon ID is required only if security is enabled on your database system.

Cursor Cache Size: The number of cursors the cursor cache can hold. The default is 6 cursors.

Yield Proc: This value can be 0 (peek and dispatch), 1 (no yielding), 2 (the SQLBase SQL/API default yield procedure), or 3 (dispatch via Windows Yield function). It is recommended that you choose the value 1.

Input Message Size: The number of bytes in the input message buffer. The default is 4000 bytes.

Lock Time Out: The number of seconds to wait for a lock to be freed before raising an error. Values can be -1 (wait forever) to 1800; default is 300.

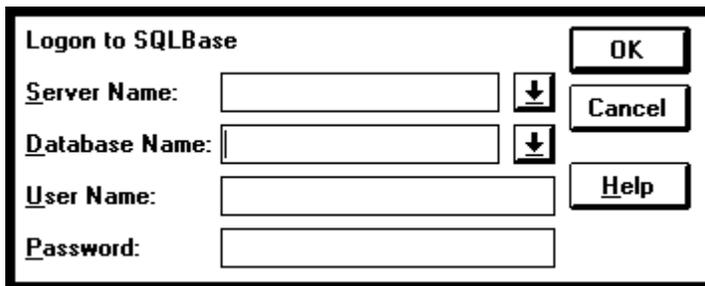
No Recovery: Selecting this check box disables transaction recovery. Selecting this box is dangerous because your database system can become inconsistent in the event of a system crash.

Release Plan: A number that determines whether a lock is maintained on a table when the cursors accessing the table are freed. Freeing the lock on the table results in a request to the server, which can decrease performance. No locks are freed when this option is set to the default, 0. Locks are freed when this option is set to 1.

4. Click **OK** to write these values to ODBC.INI. These values are now the defaults when you connect to the data source. You can change the defaults by configuring your data source again.

Connecting to a Data Source Using a Logon Dialog Box

Some ODBC applications display a Logon dialog box when you connect to a data source. The dialog box is as follows:



In this dialog box, perform the following:

1. Enter the server name.

This is the name of the server that services the database tables you want to access. You can type the name or select the name from the Server Name box, which displays the server names you specified in the Setup dialog box. To access a local SQLBase database system, enter the word LOCAL.

2. Enter the database name.

The local SQLBase database name.

You can type the name or select it from the Database Name box.

3. Enter your user name.
4. Enter your password.
5. Click **OK** to complete the logon and to update the values in ODBC.INI.

Connecting to a Data Source Using a Connection String

If your application requires a connection string to connect to a data source, you must specify the data source name that tells the driver which ODBC.INI section to use for the default connection information. Optionally, you may specify *attribute=value* pairs in the connection string to override the default values stored in ODBC.INI.

You can specify either long or short names in the connection string. The connection string has the form:

```
DSN=data_source_name[;attribute=value[;attribute=value]...]
```

An example of a connection string for SQLBase is:

```
DSN=SQLBASE_TABLES;SRVR=QESRVR;DB=PAYROLL;UID=JOHN;PWD=XYZZY
```

Table 1-1 gives the long and short names for each attribute, as well as a description.

Table 1-1. SQLBase Connection String Attributes

Attribute	Description
DataSourceName (DSN)	A string that identifies a single connection to a SQLBase database system. Examples include "Accounting" or "SQLBase-Serv1."
Database (DB)	The name of the database system to which you want to connect.
ServerName (SRVR)	The name of the server that services the SQLBase database tables you want to access. Optional.
Servers (SRVRLIST)	A comma-separated list of servers with which to prompt the user in a Logon dialog box.
LogonID (UID)	The default Logon ID used to connect to your SQLBase database system. A Logon ID is required only if security is enabled on your database system. If so, contact your system administrator to get your Logon ID.
Password (PWD)	Case-sensitive password.
CursorCacheSize (CCS)	Controls the number of cursors in the cursor cache. The cursor cache increases performance and uses few database resources.
YieldProc (YLD)	This option, which determines if you can work in other applications when the Centura SQL/API is busy, is useful to users of ODBC applications.
0	(peek and dispatch) causes the driver to check the Windows message queue and to send any messages to the appropriate Windows application.
1	(no yielding) does not let you work in other applications.
2	(default yield procedure) uses the SQLBase SQL/API's default yield procedure.
3	(dispatch via Windows Yield function) turns control over to the Windows kernel. The Windows kernel checks the message queue and sends any messages to the appropriate application window. It is recommended that you use YLD=1. YLD=0, YLD=2, and YLD=3 do not work with all Windows applications.
InputMessageSize (IMS)	Controls the number of bytes of the input message buffer. Increasing this value retrieves more records across the network in a single fetch.
LockTimeOut (LTO)	The number of seconds SQLBase should wait for a lock to be freed before raising an error. Values can be -1 to 1800; -1

means wait forever.

NoRecovery (NR)

Enables or disables transaction recovery. NR=0 enables recovery; NR=1 disables recovery. NR=1 improves performance but is dangerous because your database can become inconsistent in the event of a system crash. See your SQLBase documentation for information on this option.

Data Types

Table 1-2 shows how the SQLBase data types map to the standard ODBC data types.

Table 1-2. SQLBase Data Types

<u>ODBC</u>	<u>SQLBase</u>
SQL_DATE	Date
SQL_DECIMAL	Decimal
SQL_DOUBLE	Double Precision
SQL_INTEGER	Integer
SQL_LONGVARCHAR	Long Varchar
SQL_REAL	Real
SQL_SMALLINT	Smallint
SQL_TIME	Time
SQL_TIMESTAMP	Timestamp
SQL_VARCHAR	Char
SQL_VARCHAR	Varchar

Isolation and Lock Levels Supported

SQLBase supports isolation levels 1 (read committed) and 3 (serializable). The default is 1. SQLBase supports page-level locking.

Connections and Statements Supported

SQLBase supports multiple connections and multiple statements per connection.

Windows 95 and Windows NT Multithread Applications

The SQLBase driver does not allow concurrent access by multiple threads. Applications must guarantee that only one thread access the database at all times.

Special Notes

SQLBase stored procedures are now supported.

The driver name is the same for Windows 95 and Windows NT. The actual binary files are different. Separate installation disks are available.

Books and Online Help

The following sections describe errata in the Centura Team Developer books and online help.

Introducing Centura book

The *Introducing Centura* book, on page 4-11, in step 11, states "Click the After Synch tab." This should read "Click the After Init tab". If uncorrected, it will cause replication to fail.

Developing with Centura Builder book

In the section "How to compile a dynalib" on page 17-10 of the book *Developing with Centura Builder*, it says you can compile a dynalib from the command line using the executable swi60.exe. The correct executable to use is

cbi60.exe.

Reporting with Centura Builder book

The *Reporting with Centura Builder* book, on page 1-5, states that a right-click on the QuickReport object will open the cQuickReport customizer. The correct procedure is to double-click on the object to open the customizer.

Using and Extending QuickObjects book

The *Using and Extending QuickObjects* manual, on page 10-25, refers to USER.DLL. The correct name of this library is USER32.DLL. Refer to Migrating Applications to Centura for a list of correct library names.

The *Using and Extending QuickObjects* manual, on page 10-26, includes the multiline field class: clnstructions as the first line of the Class Definitions. The first line is actually the functional class: cSizeChildren.

Managing Teams and Objects with Centura book

The following areas of the printed documentation have been changed or are incorrect. Please set the online documentation for up-to-date information. (Press F1 from Team Object Manager, or double click on TeamObj.hlp.)

- The notes regarding PVCS in Appendix C are incorrect. Please see the on-line help for further information on PVCS integration.
- The information on manually building a repository in Appendix B is incorrect. Please refer to the on-line help for corrected information.
- The Class Browser which is included with Team Object Manager can be installed into the Centura Builder development environment, so that class browsing is available while you are editing your Centura Builder application. Instructions are contained in the on-line documentation.
- Merging is a complicated process, which should be undertaken with care. A number of warnings, caveats, and suggestions are provided in the on-line help to assist with merging. Please read these before using the merge tool, or performing a merge in Team Object Manager, or using the concurrent checkout facility.
- Team Object Manager supports a number of command line options. Please refer to the on-line documentation for details.
- Please refer to the online documentation for information on DataLinks. A DataLink is a link between a Centura Builder window and a table/column in a database. DataLinks appear on the DataLinks tab of Window components. Team Object Manager locates DataLinks by scanning windows for table-based QuickObjects, or for special comments that the programmer has embedded into the form. (Scanning for DataLinks implicitly scans for Components - it is not necessary to scan for Components if you scan for DataLinks.)
- Team Object Manager supports a trace mode, where relevant debug information is written to a trace file. Please refer to the on-line documentation if you wish to turn trace mode on. (Please note that turning trace mode on will adversely affect performance.)

Migrating Applications to Centura book

The *Migrating Applications to Centura* book on page 1-4 has a table showing the 32-bit DLLs in Centura and the corresponding 16-bit DLLs in SQLWindows. The last entry is incorrect; it should read:

SQLSYBW.DLL	SQLSYB32.DLL
-------------	--------------

This table is also missing three entries:

SQLIRTW.DLL	SQLIRT32.DLL
SYBSAL.DLL	SYBSAL32.DLL
ODBSAL.DLL	ODBSAL32.DLL

Online Help

For best results while creating the Find word list in the **Find Setup Wizard**, we recommend selecting *Maximize search capabilities*.

Printing with the Books Online Viewer

A limitation in the Interleaf Worldview books online viewer means you must print a topic in the online book collection as follows: Copy the topic, then paste it to another file, then print.

SQLBase New Features Guide book

COMDLL value: Windows 95 and Netbios

The comdll keyword table on page 1-5 of the *New Features Guide* should contain an additional row for the Windows 95 platform as follows:

Platform: Windows 95
Protocol: NetBIOS
File name: SQLNTNBI.DLL
Comdll value: sqlntnbi

SQL Language Reference book

The following items update the *SQL Language Reference*.

ADDENDUM TO TRIGGERS DOCUMENTATION

The following is an addendum to the documentation on triggers in the *SQL Language Reference*:

- You cannot alter a table that has a trigger defined on it.
- You cannot pass SQL @ functions as parameters to procedures that are referenced by triggers.
- In certain situations SQLBase allows you to drop a table that a trigger or stored procedure has actions on, without warning you that a dependent trigger or stored procedure exists. Instead, you receive a runtime error later, when the trigger or stored procedure executes and the table on which it acts no longer exists.

For example, assume you create two tables, A and B. Then create an update trigger (TRIG_A) on table A. Now assume that TRIG_A calls a stored procedure (named SP_A) to insert data into table B.

If you try to drop table B, SQLBase allows this without first warning you about TRIG_A, which is a dependent object of table B.

If you drop table B and then update table A, the trigger fails at run-time with the error that table B does not exist. The transaction is rolled back and you cannot update table A until you recreate table B or drop the trigger.

A similar situation occurs if you erase stored procedure SP_A.

- The following paragraphs clarify information on page 3-69 regarding using the "BEFORE UPDATE" and "AFTER UPDATE" clauses:

SQLBase can verify that updated data adheres to integrity constraint rules before performing an UPDATE. If you use the REFERENCING NEW AS <new values tablename>.clause of the CREATE TRIGGER command with the BEFORE UPDATE clause, then the updated values are accessible to the triggered SQL statements.

Similarly, if you use the REFERENCING OLD AS <old values tablename>... clause of the CREATE TRIGGER command with the AFTER UPDATE clause, then the values that existed prior to the invoking update are accessible to the triggered SQL statements.

- Do not use triggers to issue INSERT, UPDATE, or DELETE actions indirectly (such as through a referential integrity action like DELETE CASCADE) or directly against the table on which the trigger is defined. If you use a trigger to perform such actions, you get error 848 ("Row being processed for a DELETE/UPDATE was modified by triggered actions") when the invoking DML is executed.

- If you perform an INSERT/UPDATE/DELETE with a subselect or WHERE clause, and if this DML causes a triggered procedure to run, we do not support that triggered procedure modifying any table referenced in the subselect or WHERE clause of the invoking DML statement. This restriction exists because actions of the trigger can cause the subselect or WHERE clause of the invoking DML statement to return unpredictable results.

For example, if the following statement invokes a triggered procedure, we do not support the triggered procedure

modifying the table referenced in the WHERE clause (table T1):

```
DELETE FROM T1 WHERE C1 > 3;
```

Note that in SQLBase 6.0.1 and later, you can use BEFORE triggers to set values for columns with INSERT and UPDATE statements. For more information, refer to page 2-12 of the SQLBase New Features Guide.

- Page 3-66 of the *SQL Language Reference* incorrectly states the conditions necessary to create a trigger that uses a stored procedure. In order to create a trigger that uses a stored procedure, one of the following must be true: 1) You are SYSADM or have DBA authority, 2) you own the table and the stored procedure, or 3) you own the table and you have been granted execute authority for the stored procedure.

- You can define a maximum of eight triggers on a table: two INSERT, two DELETE, and four UPDATE triggers. This corrects information on page 3-70.

You cannot create a trigger on a system catalog table.

ADDENDUM TO PROCEDURES DOCUMENTATION

The following is an addendum to the documentation on procedures in the *SQL Language Reference*:

- You cannot replace an existing procedure with a procedure that uses the same stored name. You must first erase the first procedure and then store the second instance.

- Static procedures must be stored before they can be executed. Dynamic procedures can be executed whether or not they are stored.

- Static procedures cannot perform DDL commands except for a LOAD...ON SERVER... command.

- Recursion and nesting limits of procedures are bound by your system.

- A procedure cannot perform SQL commands that require a SET SERVER command (for example, CREATE DATABASE).

- The documentation on the customizer's ROW field is incorrect (page 7-51 of the *SQL Language Reference*). The documentation should state that you set the "maximum rows in memory to 32000" (instead of "rows to 32000").

- If you have a procedure that performs some action on a table, you are able to drop or alter the table the procedure references. This results in a runtime error when the procedure is executed. You should note this behavior, because you may expect users to be unable to drop or alter a table that a procedure references.

- The SqlRetrieve function documentation on page A-32 of the *SQL Language Reference* should reiterate the fact that you must preface the strBindList and strIntoList parameter values with a colon.

- The section on Indentation (page 7-14 of the *SQL Language Reference*) should state that you should not mix spaces and tabs for indentation. For example, four spaces might appear to have the same indentation as a tab in your on-line editor, but the four spaces represents four levels of indentation, while the tab only represents one.

- The documentation for the SqlCommit function (page A-4 of the *SQL Language Reference*) should state that in a stored procedure, if you have a SqlPrepare function called in an On Procedure Startup section, and a SqlCommit function called in a subsequent On Procedure Execute section, the commit will destroy the cursor of the SqlPrepare function.

Subsequent executions will fail because the cursor's "preparation" has been lost. To preserve the cursor's `SqlPrepare` statement, you must first turn on cursor context preservation. To do this, use the `SqlSetParameter` function in conjunction with the `DBP_PRESERVE` parameter.

- Page 7-10 of the *SQL Language Reference* should state that (as of release 6.0.1) use of the `NUMBER_Null` system constant is supported in procedures.

USING PROCEDURES WITH SQLWINDOWS APPLICATIONS

Because `SQLBase` now provides scrollable result sets for procedures, you can now disregard the following previous restrictions:

- On page 7-4, disregard the first paragraph; programs calling `SQLBase` procedures can now use scrollable cursors.

- On page 7-51, the second paragraph is now incorrect; beginning with `SQLBase 6.0.1`, you CAN use `SqlFetchPrevious` when the source of data is a `SQLBase` procedure.

- On page 7-51, the fourth paragraph, disregard the statement that begins "Turn off Results Sets with `SqlSetResultSet...`" By default `SQLWindows` turns results sets on, so it is no longer necessary to turn them off. Additionally, you are no longer restricted to setting rows to 32000 or setting `discardable` to `NO`. Instead, adjust these customizer settings as your application requires.

- Page 7-52, in the third code example, eliminate the following line of code: `Call SqlSetResultSet(hCur1, FALSE)`.

OWNERSHIP OF DATABASE OBJECTS (IF UNQUALIFIED)

If you are create an object in someone else's name (for example, if you are creating an object for `USER1` and you are either `SYSADM` or have `DBA` privileges), `SQLBase` assumes that unqualified names in your statement belong to the user, not you. For example, assume you execute the following command as `SYSADM`:

```
CREATE TRIGGER A.TRIG BEFORE UPDATE ON T1...
```

Since table `T1` is unqualified, `SQLBase` assumes that the qualified table name is `A.T1`, not `SYSADM.T1`

This behavior applies for `CREATE VIEW`, `CREATE TRIGGER`, and `PROCEDURE` commands.

DATA VALIDATION OF DML STATEMENTS

The following is an addendum to the *SQL Language Reference*. It lists the execution order for data validation, trigger execution, and integrity constraint checking for a single DML statement:

1. Check for number of bind data.
2. Validation of values if values are part of the statement (that is, not bound). This includes null value checking, type checking (numeric), etc.
3. Security checks.
4. Execute `BEFORE` statement trigger.
5. Loop for each row affected by the SQL statement.

For each row:

a) Validation of values if they are bound in. This includes null value checks, type checking, and size checking (character string too long). NOTE: Size checking is done here even for values that are not bound.

b) `BEFORE ROW` trigger fires.

- c) Duplicate value checking.
 - d) Referential integrity checks on invoking DML.
 - e) INSERT/UPDATE/DELETE.
 - f) AFTER ROW trigger fires.
6. Execute AFTER statement trigger.

This is the execution model for any DML statement. Triggers themselves may cause DML to be executed, and those actions will also follow the steps shown in this model.

LOADS AND OBJECTS REFERENCED BY TRIGGERS OR PROCEDURES

The following information updates information on LOAD, triggers, and procedures in the *SQL Language Reference*:

If you alter or drop objects that are referenced by triggers or procedures, you can receive errors when you later attempt to load the trigger or procedure. To avoid the error, you must either:

- Recreate any referenced object that you dropped

OR

- Restore any referenced object you modified back to its original state (the state known by the procedure or trigger)

SELECT...FOR UPDATE OF CLAUSE

The description of the FOR UPDATE OF clause of the SELECT statement (page 3-143 of the *SQL Language Reference*) fails to state that the behavior occurs when you use named cursors.

The corrected paragraph is:

"If you are using a named cursor, this locks parts of a table so that a subsequent UPDATE or DELETE will not cause a deadlock between concurrent users. This clause is compatible with DB2."

Also, the first three sentences at the top of page 3-144 (also in the description of the FOR UPDATE OF clause) fail to state that update locks are incompatible with other update locks. The corrected version is:

"When you use the FOR UPDATE OF clause, SQLBase uses update locks. An update lock reduces the possibility of deadlocks. Update locks are compatible with shared locks, but not with other update locks or exclusive locks."

UPDATING THE STATISTICS FOR STORED COMMANDS

The "UPDATE STATISTICS" documentation in the SQL Language Reference manual states that you update the statistics associated with your stored commands and thereby improve performance. This section should also state that you should then run the Centura-supplied RECOMPILE procedure. For more information, refer to Appendix B of the Database Administrator's Guide.

DATA TYPE CORRECTION (INTEGER PRECISION)

On page 2-13 of the *SQL Language Reference*, the second example of DECIMAL definitions erroneously states that 5 is the precision of INTEGER:

"DECIMAL (5.0), (same as INTEGER precision)"

The corrected statement is:

"DECIMAL (5.0), (same as SMALLINT precision)"

NUMBER DATATYPE

The description of the NUMBER datatype on page 2-12 of the *SQL Language Reference* is misleading. The description should state the following:

NUMBER is a superset of all SQLBase numeric types and supports the widest range of precision and scale (up to the maximum allowed by SQLBase numeric types). The NUMBER datatype supports 23 digits of precision, but one of these "digits" must be a plus (+) or minus (-) sign.

GRANT EXECUTE ON

The *SQL Language Reference* erroneously omits the GRANT EXECUTE ON command. This command grants execute privilege on stored procedures to other users. For information on the command, refer to the *SQLTalk Command Reference*.

REVOKE EXECUTE ON

The *SQL Language Reference* erroneously omits the REVOKE EXECUTE ON command. This command revokes a user's execute privilege on a stored procedure.

For information on the command, refer to the *SQLTalk Command Reference*.

SQLTalk Command Reference book

The following items update the *SQLTalk Command Reference*.

SQLTALK BAT...CACHE=N OPTION NO LONGER VALID

The SQLTalk BAT...CACHE=n option documented on page 1-14 and 1-15 of the *SQLTalk Command Reference* is no longer valid. If you set this option you do not receive an error, but SQLTalk ignores the setting.

SET TRACE AND SET TRACEFILE

The documentation for the SET TRACE and SET TRACEFILE commands (page 2-111 of the *SQLTalk Command Reference*) should state that traces are cursor-specific.

Database Administrator's Guide book

The following items update the *Database Administrator's Guide*.

GROUPCOMMITDELAY keyword

The *Database Administrator's Guide* should include the following documentation on the GROUPCOMMITDELAY keyword. (Note: for information on new SQL.INI keywords, refer to the *SQLBase New Features Guide*.)

Purpose. Specifies the maximum number of system ticks that SQLBase waits until actually performing COMMITs. The meaning of "tick" is platform-dependent, but is approximately 1/20 of a second. SQLBase physically performs COMMITs when either of the following conditions is met:

The number of grouped COMMITs exceeds the value of groupcommit.

or:

The number of ticks that have elapsed since the last physical COMMIT is greater than the value of GROUPCOMMITDELAY.

Default value. The default setting is 1 tick.

Example. To set GROUPCOMMITDELAY to 20, specify:

GROUPCOMMITDELAY=20

The best setting for maximum performance is application dependent. You should experiment with various settings to find the right setting for your environment.

Optional. Configuration of this keyword is optional.

Where configured. Configure this keyword in any single-user engine or multi-user server section of the configuration file (SQL.INI).

SET INPUT MESSAGE BUFFER SIZE THROUGH SQL/API, NOT SQL.INI

Page 3-37 of the *Database Administrator's Guide* erroneously states that you can use the INMESSAGE keyword to set the size of the input message buffer. Setting this keyword actually does NOT affect the input message buffer size. Instead, use the SQL/API SQLIMS (Input Message Size) function to set input message buffer size.

SET OUTPUT MESSAGE BUFFER SIZE THROUGH SQL/API, NOT SQL.INI

Page 3-37 of the *Database Administrator's Guide* erroneously states that you can use the OUTMESSAGE keyword to set the size of the output message buffer. Setting this keyword actually does NOT affect the output message buffer size. Instead, use the SQL/API SQLOMS (Output Message Size) function to set output message buffer size.

AUTOMATIC CONFIGURATION OF GUI KEYWORDS

Chapter 3 (Configuration) of the *Database Administrator's Guide* contains incorrect information regarding configuring the GUI keywords. The document should state that SQLBase automatically configures the GUI keywords for you when you use SAVE SETTINGS from the FILE menu. The GUI keywords are:

- DBWIN
- MAINWIN
- PROCWIN
- STATWIN
- SYSWIN

LISTENRETRY KEYWORD

Chapter 3 of the *Database Administrator's Guide* incorrectly states that the LISTENRETRY keyword is configured only in the NetBIOS communications section of the SQL.INI file.

LISTENRETRY is actually protocol-independent and is configured in any multi-user server section of SQL.INI.

PROCESS ACTIVITY DISPLAY PARAMETERS

The "Process Activity display screen (F2)" section of Chapter 5 of the *Database Administrator's Guide* fails to describe the following parameters that appear under certain (high) message levels. The parameters are:

PNM = client process number
SNM = client communications session number
CFL = a flag byte value
FCD = function code

COMDLL SETTING FOR NETWARE 4.X SPX COMMUNICATION LIBRARY

Page 2-33 of the *Database Administrator's Guide* contains incorrect information regarding the COMDLL setting for the NetWare 4.x communication library. In the "SPXDLL40.NLM" section, the sample COMDLL setting should state:

```
comdll=sqlspx
```

ACCESS TO THE SYSTEM CATALOG

Page 7-24 of the *Database Administrator's Guide* contains incorrect information on system catalog access. The corrected paragraph is:

"By default, SELECT privileges on the system catalog tables are granted to PUBLIC (all users). The exceptions are SYSCOLAUTH, SYSCOMMANDS, SYSEVENTS, SYSTABAUTH, SYSTRGCOLS, SYSTRIGGERA, and SYSUSERAUTH, which only SYSADM and DBA can view..."

SYSTEM CATALOG SUMMARY

On page A-21 of the *Database Administrator's Guide*, the description of the GROUPNUM column in the SYSADM.SYSTABLE system table is incorrect. The column description for GROUPNUM should state: "The group number for the table."