

Developing High Performance Windows Client Applications for Sybase / MS SQL Server

This is a posting of my experiences (Disclaimer <Copyright and Liability>) of using MS Windows 3.1 with Visual Basic Professional 3.0 / Visual C++ with;

[dbLib / ODBC,](#)

[Jet / ODBC,](#)

[Utilities,](#)

[Sybase / MS SQL Server](#)

There is a new section New Software which contains related news and information on updated software. *Remember you can use the search button to browse the topics in this file.*

What is not covered?

I do not use or discuss any of the ISAM (index sequential access mechanism) database systems, (dbase, Access, etc.) nor anything about Jet (except how to [NOT use Jet](#)). This a question of focus not a comment on other products there are other documents that discuss Jet in detail. I am not interested in advocacy or discussion of these or related topics. This is not an [FAQ](#) but there is a [FAQs](#) section. If you want something added here mail me and I will try and add it to the [FAQ](#) section.

[About this posting...](#)



About this...

The idea behind this posting is to keep developers producing Windows Client applications using dbLib, ODBC, Sybase and Visual Basic up to date, so that we can share our experience. I develop purely for Sybase (and Oracle), so I can not say much about other database systems.

Typically the applications I develop are deployed in an international corporate arena. This means the programs must be of a high standard, and working in the financial sector means that response time and accuracy are key objectives. For this reason you might find my observations very (some might say overly) critical and demanding. This is a reflection of the environment for which I develop. Where ever possible I have included detailed description of bugs and workarounds that I have used.

This an account of some of my experiences with various products. It is not an endorsement of any product, group of products, or method of development. Nor a statement of product suitability. It is not intended as a tutorial or a discussion document about programming style. I am not advocating the use of any product(s) or method(s) nor suggesting that the techniques and ideas discussed are suitable for all forms of development. The information included is only intended to help point out *some* of the issues that *may* be relevant for developing Windows applications for Sybase / MS SQLServer.

I make no conclusions about what is suitable in general only express my personal view based on my experiences.

Seth.

Finally...

Remember please let me know what you think of the posting good or bad, especially technical inaccuracies. I would also appreciate any suitable contributions for the posting from other developers. I am not overly preoccupied about product level advocacy issues. I prefer to focus on performance (Client, Network and Server), ability to deliver and deployment.

Let me know what you think about this posting and any comments / suggestions. At present it is private mail NOT freely distributable.

I do not mind others reading the text, however some of the text may contain undocumented code.

If there is enough interest I will start a freely distributable version and FAQ on the Sybase usenet conference. If anyone wants a copy of this mailing ask them to mail me and will send them a copy.



You can contact me via e-mail if you do please remember the following guidelines.



Let me know about mistakes, especially things that are unclear or wrong with this document.



Provide some background information about yourself, what you do and the environment you work in. (This helps me to compose replies that are more likely to be useful with suitable emphasis)



Avoid posting specific problems (e.g. My program GPFs when calling dbLib, heres the code...!) it is better to phrase it in general terms. (e.g. My program GPFs whenever I call dbCancel in dbLib)



Include all technical information that is relevant (e.g. version of dbLib, network transport, version of server etc..)



By all means include any comments about the contents and how you think things might be improved (without changing the objectives!)



seth@cix.compulink.co.uk





New Software

Windows NT 3.5

This is the all new faster leaner version of Windows NT. The lower memory use means more memory for the SQL Server and so better performance for everyone. The Network Stacks have also been improved and updated. In practice these improvements conspire to produce a tangible difference in performance. Note that MS SQLServer 4.21a is required for NT 3.5.

Sybase System 10.0.2P2 for NT

At last a version of Sybase System 10 for NT that seems to work! At present there are still some problems with the ODBC layers and some of the new data types. At least Sybase are taking NT seriously as a platform. Im still waiting for a decent ODBC driver for System 10. Id like to know of any problems people encounter with System 10 (via dbLib or CTLib) from simple tests, CTLib seems to be slower than dbLib, but my testing is not complete.

MS ODBC Drivers for SQL Server

Microsoft have launched their new ODBC drivers. These include both 16-bit and 32-bit versions for SQLServer. One my performance issues with Microsofts ODBC drivers were that they took a long time to connect. Q+E get around this by only getting information about the server as it is required. MS now do this sort of thing and the result is a much faster connections time. In my experience the Q+E ODBC drivers for SQL Server still seem to be faster and more reliable. I notice that both Q+E and Watcom ODBC drivers have an MS Mode, to work like the MS ODBC drivers to cope with situations where an ODBC driver does things slightly differently to the ODBC specification.

This new lease of life in the MS driver and the fact that Q+E has started charging a per user licence for its ODBC drivers. If cost is a major consideration then it is worth checking a vendors licensing policy.

VB 4.0

Since MS plan to release VB4.0 after Win 95 is out (due September 95) all the VB add-ons vendors are having to re-think their plans.

There are problems with VB4.0 that mean any control that adds more than 6 custom controls to the tools palette wont work. But this is a VB4.0 issue. OCXs are available and seem to work with the beta of VB 4.0. We do not use this in anger, I tried this out one rainy afternoon.

OCXs

OCXs are the replacements for VBXs. They will become important as Win 95 and Windows NT 3.5 are released.

The official line from Q+E was that no major new version of Multi-link is planned. The next major release will be an OCX. However they have shipped me a bug fixed version, and seem to be improving the Multi-link product and driver set. This may be subject to the release date of Win 95 and VB.4.

The VBX to OCX move will be simplified by a Wizard in Microsoft VC++ version 2.00 that does the conversion automatically. The only additional function is one to export the types used by the VBX. This Wizard will also be made available for MS VC++ v 1.5. That said the Q+E people at the OCX developers

conference seemed to know more about data aware OCX issues than some other developers.

MS SQLServer Version 4.21a for Windows NT 3.5

This is much more than it might sound like. Firstly this is the ONLY version of MS SQLServer certified to run on Windows NT 3.5. This new release of Windows NT is much less memory hungry than NT 3.1. The result is that existing databases will be able to use the memory that was previously used by the OS.

Additionally the 4.21a release of MS SQLServer can also support replication services *from* Sybase System 10! This means that a network of MS SQL Servers can be fed via a back bone of Sybase System 10 boxes. At present the replication is only *to* an MS SQLServer but that may be enough. If you want to get data off an NT box for replication then you will have to install Sybase System 10 for NT. SQL Server 95 the next release of MS SQL Server will have many of the System type functionality. The competition will certainly make things interesting.

Q+E Buyout by Intersolv

I have used Intersolv products before and was impressed with their company & products (even though they were expensive). The Q+E buyout give Q+E a larger capitalisation and in my opinion makes it more able to compete with other players in the market. Intersolv s products work Multi-server multi-OS and they seem to be independent of all vendors, while remaining market leaders in some market sectors.

In practical terms the Q+E name is now Intersolv Q+E, the service seems to have improved. I reported bugs and got confirmation. I mentioned competing products that did what I wanted. Q+E bought the products to see what the issues were and then came up with an even better solution, which they implemented and seems to work well.

PER USER LICENSE

Intersolv have announced plan to charge per user for their ODBC drivers. The current ODBC drivers are royalty free, however future version of Q+E ODBC will not be royalty free. THIS MAY EFFECT YOUR CHOICE OF Q+E MULTILINK VB.

Multi-link 2.02/2.03 (Beta)

New version make sure you have the latest version. The known bugs I have found are included in this posting. In general faster more features etc.. The new driver seem to be 20% faster in most of our applications. System 10 Support is available, however I have not used this in anger yet. I have set up a System 10 SQL Server, Replication Server & OMNI-Server I will keep you updated on this. Refer to [Common Questions](#) section: [Which dbLib should I use?](#)

Other Sybase Utilities

Have you found a utility for Sybase that might be useful for other developers, let me know and I will add your text to the next posting..

[Making your own VBX](#)

[Sombrero](#)

[VBSQL](#)

[ADE ODBC Tools](#)

Making you own VBX for Sybase



We have developed software that allows us to by pass ODBC for some tasks (e.g. operations that are not cursor or large result set based, computed columns etc..). It allows you to connect to Sybase directly (via dbLib). This talks via dbLib so is it very fast. 0.1 of a second to retrieve 400 customer names and populate a VB. tabbed list box control and about 60 individual (single row) SQL statements executed per second (Client to server and back).

This has been useful for running stored procedure, computed columns and other features poorly or not supported by ODBC.

It also allows us to populate text boxes, the clipboard and listboxes with simple SQL directly. Combo boxes (even tabbed listboxes) can also be populated with regular SQL. (E.g. `SELECT * FROM sysobjects`) will produce a dropdown with the sysobjects. Again all this is possible via ODBC / Multilink VB but we need improved performance.

We are using it for parts of a program where we need the raw speed, and the control does not have to be linked to a Q+E query.

I mention this to illustrate that dbLib programming under Windows is relatively painless. It will allow you to enter SQL directly to dbLib and have the results sent direct to any Windows control. This cuts out the bottleneck of VB., while allowing VB. to access all the data returned.

You can also add more sophisticated things like caches, time-outs (close connections that have not been used for a while) and hinting (an SQL Comment that means the SQL will NOT be cached). It might be worth doing as it stops hitting the server and reduces network i/o. Which can only improve response time and server & network usage.

A cache might contain the following components;

Router Routes incoming SQL and decides where to send the request to the caches or to the SQL Server. *Is it cacheable?* (Checks for URGENT hints, GetDateTime, INSERT, UPDATE, SELECT INTO, there are functions in dbLib that return true if the SQL changes.)

Requester Requests information form the cache. *Is it in the Cache?*

Manager Is the cache full (MFU algorithm looks the best, easy to implement)

Sombrero

Sombrero VBX for dbLib this is a commercial product that does the dbLib sort of things. Seems to be a simple wrapper to make dbLib callable from VB. Available as an OCX soon. OK if you do not have your own wrapper. Also works with both Sybase and MS SQLServer.

VBSQL

This is part of the Microsoft SQLServer Programmers toolkit for Visual Basic. Seems to be a simple wrapper to make dbLib callable from VB, unlike Sombrero this **may** not work with future versions of Sybase (MS might not support for CT-Lib and new datatypes support e.g. numeric data types).

ADE ODBC Tools

A quaint set of tools that allow access to ODBC API via regular bound controls. A SQL Admin. Control, an Editable Grid, and some other interesting things. Nice product, lots of potential. We do not use it because of lack of error handling and customisability. Very helpful tech -support well worth a look.

One unique feature of the ADE Tools it that they have a bound control that talks directly to ODBC (by passing Jet). This seems to work very well, though I have not used this functionality in anger.

Q+E MultiLink VB (aka DataDirect)

This is a set of ODBC drivers and Custom Controls for VB. The drivers cover several databases. I have been using the drivers and controls with Windows 3.1 and MS SQL Server. For this configuration I have found the combination of Q+E Multi-Link VB and the associated drivers to be very fast and reliable. This approach the most popular way of developing database applications using VB. However this combination of controls and drivers has allowed me to develop applications that out perform systems using standard VB database construction methods.

See also:

[Known Bugs.](#)

[Debugging.](#)

[Tips](#)

Problems With Q+E Multi-link VB and Sybase (NOT fixed in Multi-Link VB Version 2.0.3)

Introduction

The Q+E Multi-link VB software provides some excellent functionality, indeed the QueryGrid alone is the fastest datagrid by far.

This is a summary of the problems I have encountered;

Logging:pOneHstmtOption = QE_HSTMT_LOCAL *(Fixed March 95 not generally available, should be in next full version)*

pMasks (Currently being completely rewritten by Q+E)

pKey (Can be worked around)

Connections & Multiple .EXEs

Problem : pOneHstmtOption = QE_HSTMT_LOCAL

This does not work on Sybase/MS SQLServer databases. To reproduce the error, try and perform a query on a large table. While this query is active, issue a second query. Q+E will then attempt to download the contents of the first (large) database to your PC.

Status: Fixed not generally available.

I have called this problem the Local logging problem (seems to effect Sybase 4.x & SQLServer). QELINK.VBX 2.0.3 (Dated 7th July 1994) and new driver DLL seemed to fix the problems with Q+E Multi-link attempting to log all records in the current query into a local file. Even if you have 100,000 records in the current query. This fix was dropped from subsequent code.

The QELINK.VBX 2.0.3 (Dated 7th July 1994) allows update, new, delete operations however the fFind still causes a log to be generated. Reported this to Q+E.

A QueryGrid that has an active query to a large table (say 100,000 records 4k per record) and then start a query on a form to add a record to a small (say 5 record) table. Q+E seems to try and read the contents of the large table into a local file. (This is like the problem that we had originally with queries, which was fixed and then the fix dropped. Currently the logging problem occurs with Query Controls and Query Grid controls.

Q+E Responses:

Setting this option should prevent a new query from effecting other active queries. We will also correct any errant behaviour in MultiLink with this property setting.

This issue has been recreated and filed as Failure # 58155.

Problem : pMasks.

Basically in its current state this is unusable, it does not work like a Windows Edit box and the behaviour of the these controls is unpredictable / unreliable.

pMask should allow;

Text masks Uppercase /Lowercase / mixes

Date mask and optional popup calendar (like a date in Spread VBX)

Max. Length (This is a standard Windows function on all Edit Boxes)

Correct behaviour with VBs clipboard functions

In Find Mode the mask should additionally allow operators (<,>!=, etc..) to be entered.

Q+E Response: pMask code / functionality is being reviewed and will be re-written.

Problem: pKey

From the ML Help file qelink.hlp;

You can specify two field controls as primary keys so that MultiLink/VB. searches two database fields for the unique combination of the **two** values.

Setting pKey = True allows ML to construct a Where clause for the update statement with only the fields where you have specified pKey = True.

This does **not** seem to be correct. For example...

We use 4 pKeys (set to true) and this seems to work even though the documentation says 2 is the maximum.

If I do not set any pKey fields then, everything works fine for tables with one or two primary fields in the primary key.

However doing this for a table with three or four fields as primary keys, I get an error "more than one record matches the record you are trying to edit / update". And I have to set Pkey to true for all the fields used in the primary key.

Sybase / MS SQLServer if passed a where clause with several fields WILL select on primary key fields first REGARDLESS OF ORDER. (This is a feature of the SQLServer Optimiser). As this is the case, and all record being edited by definition will have all the primary fields completed why not use all the fields (that are filled in) in the where clause. This will allow ML to support any # of pKey fields and the programmers would not have to tell ML which fields were primary keys.

Has the limit of 2 pKey fields been removed? It should be unlimited. Also Sybase can tell the ODBC driver what the primary index fields are so why should the programmer have to provide this information. Let me know if you need some SQL to detect primary keys and index fields by using the system tables, this is a legal and fully supported way of extracting this information and will work under Sybase 4.x, Sybase system 10 and MS SQL Server.

Problem: pKey not used by Q+E for fFind

When performing a search the information provided in the **pFindValue** for primary key fields is not used by a call to **fFind**. This means that Q+E will retrieve all records in the database and search them on the client! To get around this bug I have had to write code that scans the **pFindValue** of the controls and generates a where clause which I have to apply to the query controls **pWhere**. This makes the server perform the filtering, which is hundreds of times faster.

I have written some code to perform this generically for a form. Please let me know if you need a copy of this to include in your sample programs. Really this sort of performance improvement should be built into your product. I had a user who said a search (using the primary key ID number) was taking 30 minutes. By setting the pWhere the search now takes 3 seconds.

Fix: This can be fixed by using QBE mode that performs the filter on the server.

Problem: pWhere Clause > 100 characters

Setting a pWhere clause that is over 100 character in length causes an error. This is a problem for German, Dutch and Japanese clients names where the business name and legal entity name can exceed 100 characters. This is not a limitation of Sybase / MS SQLServer it seems to be a bug in Q+E MultiLink VB.

Problem: Clearing / Resetting a Query Grid without calling a remote database

Let me start by saying that this is the best data control I have ever seen or used in term of performance and usability. The query Grid control seems to be unbelievable fast even on huge tables. We have literally every data aware grid on the market and the Q+E QueryGrid is the fastest by far.

Are there plans to add tot he functionality of the query Grid?

Despite is amazing speed there is no;

-direct print or clipboard facility (I have to copy the contents of the QueryGrid cell by cell! To another grid like Spread.VBX and print using Spread.VBX. This is slow as the QueryGrid only supports cell by cell retrievals) It would be nice if the QueryGrid had some sort of block mode / report writer / print options.

-option to size a columns to fit the data / column header

-way of clearing a QueryGrid once populated without having to issue a query that will no retrieve any records. E.g. Setting the pWhere to 1=0. This means that clearing a QueryGrid involves a call to the server. This generates an unnecessary hit on the server, and makes the client software look slow. (E.g. we use QueryGrids to show child records on a form). However clearing the screen requires calling fNew, and then re-querying the QueryGrid as well.)

-The query Grid in its uninitialised state only has two columns and looks like an error to most users. We can hide the QueryGrid in some applications, but this leave a whole in the screen that is filled later. It would be nice to draw the field names etc. on the grid without have any data there when the grid is initialised or cleared.

-The grid is not editable. This can be worked around with forms but users want to bulk load data via an editable grid control. I have both Formula One and Spread VBX both are good but need the daft MS data controls that are slow and clunky.

-The select options on the QueryGrid do not work as expected, e.g.;

i)Selecting a column or a row does not highlight the required area of the grid.

ii)It is not possible to drag a column to change the order of columns. There are no events in the control to

add this functionality either.

Problem: fDoQuery followed by fFind, displays the results of the fDoQuery and only then starts the fFind.

This means that when the users enter some information to do a search, it first record in the results set is always displayed. This take time and confuses the users. What we have had to do is to set the forecolor of all DBText boxes to the same value as there bgcolor during this time. Then restore the colours once the search has got some results. This problem is compounded by the fact that fFind logs all data locally and then searches this on the client. Since we set the pWhere by using the pFindValue prior to the fDoQuery now this is not a major problem.

Fixed: Use QBE mode or setting the pWhere. This is not clear in the documentation. If you use Ffind the client performs the filter, using QBE set the WHERE clause and the Server performs the Filter.

Problem : All Q+E Combo box controls fire off SQL to populate themselves when they are displayed.

This might seem like a small overhead but on a form with say 10 combo boxes with say 50 users this can lead to 500 queries hitting the server just to draw the form.

Q+E Response: This issue has been raised by several customers and is being considered for the next major release of Multilink. A work around to enable this functionality is to control the query that fills the list code. This is done by setting Query2.pAutoQuery = None (where Query2 will actually bring back the list values). Then, in the Got Focus event of the db_combo box, add 2 lines of code executing fDoQuery and fEndQuery. This will force the combo box to only populate itself when someone clicks on the combos drop down arrow.

Q+E Response: Enh Request # 58173. *[Combo boxes should be loaded only on click event & not pre-loaded. Try placing code for fDoQuery(ListQExpr_Query) in the click event of the combo box, while at the same time setting ListQExpr_Query.pAutoQuery = False.]*

Combo boxes should allow more than one field to be included in the select but only one nominated one to be put back into the database text field. Another field should be able to be passed back to another VB. control.

Q+E Response: Enh Request # 58174.

Problem: Integrating other data aware controls with Q+E query controls. Is this possible?

Q+E Response: No, not with the current release of Multilink.

Problem : Bugs with DB Text Box

There are a number of problems I have found with the DBText boxes in Multilink VB. version 2.0.3. QELINK.VBX [QELINK.VBX, 340720, 13/09/94, 0:00]

Problems retrieving selected text to clipboard

Multi-Line / Maxlength

GPFs in QELINK.

Printing DBTextboxes

Retrieving Selected Text To Clipboard

The first is that the clipboard functions do NOT work all the time. Selecting some of the text in a DBText box and then calling SelText can return an empty string. Because of this bug I have no choice but to retrieve all the text.

```
Clipboard.Clear Clear Clipboard.  
If Screen.ActiveForm.ActiveControl.SelText <> Then  
Clipboard.SetText Screen.ActiveForm.ActiveControl.SelText, 1  
Else This is to handle the bug in Q+E DBTextBoxes  
Clipboard.SetText Screen.ActiveForm.ActiveControl.Text, 1  
End If  
Debug.Print Screen.ActiveForm.ActiveControl.SelLength  
Debug.Print Screen.ActiveForm.ActiveControl.SelStart  
Debug.Print Screen.ActiveForm.ActiveControl.SelText
```

The problem with this fix is that it does not allow the selected text to be copied because the SelText function does not work with DBTextboxes.

Multi-Line / Maxlength

Another problem is that setting the multi-line property **and** the Maxlength on a DBText box can stop keyboard presses being processed by the DBText control in question.

E.g. Set a DBTextbox Maxlength to 420, and the multi-line to true. Now paste text that exceeds 420 chars (or something that make the contents exceed 420).

The DBTextbox will no longer process key presses (e.g. BACKSPACE key!)

However clipboard based cut copy paste will work. Once the length of the text drops below the Maxlength limit.

Maxlength

The Maxlength feature is very useful as are multi-line textboxes. The only way I have found of working around the problem is to set Maxlength to zero. This means that if the user exceeds the database definition of the string the server returns an error.

Maxlength Vs Find Mode

Setting the Maxlength works however during Find mode (on a form with Multilink controls) there may not be room for an operator and a space e.g. note the space). It would be nice if the max. length property could be disabled during a search. We can do this in code e.g. pass the form name to a function that enumerate the DBText controls and checks the Maxlength setting, if this is not zero we

add 2 to the Maxlength setting of each control. A similar routine would do the reverse when leaving Find mode.

Q+E Response: Enh Request # 58155.

Problem: Printing DBTextBoxes

Using the PrintForm method on a form with DBTextboxes results in a printout where the contents of the DBTextboxes are transparent (e.g. VB. or Program Manager Screens show through the DBText boxes.) I have tried this in a LaserJet III, IIIsi, IV and a postscript printer, also the problem only appears with DBTextboxes, so it is **not** a VB. / Windows problem as all other printing of bitmaps is fine.

I have written some code that takes a form as a parameter. This routine rasterises the forms & prints it scaled to the correct aspect ratio for the printer printable area, it works on both HP LaserJets and Postscript printers. This code can be used as a complete replacement for the PrintForm method. It is also not effected by the bugs in the Q+E DBTextboxes. Let me know if you need a copy of the code to add to your product / sample code.

Problem: Cursor Position after record Save

Q+E cursor positioning. Using DBText boxes, after completing the first save of a new record (fNew followed by fUpdate & fEndQuery), a subsequent fNew works except the cursor in each of the DBText fields is positioned at its last position from the previous edit. That is, in the middle of the box, even though the 2nd fNew has blanked out the previous data AND Ive even tried to explicitly set the field to NULL data. This means that the user has to hit extra keys or do an extra mouse click to position cursor at the beginning of the field.

Problem: Q+E DLLs & VBXs loading into Memory below 1 Mbyte

The problem being the use of FIXED and/or NONDISCARDABLE in the CODE and/or DATA statements in the DLL/VBX Definition (.DEF) file.

To prove this is a problem I have modified the MYOWNVBX.VBX custom control to be FIXED NONDISCARDABLE. I then ran the BELOW1MB memory mapper with the old and new versions. This custom control does no Global or Local Allocs, and only uses VB. properties for storage. So if any global memory appears allocated to the MYOWNVBX module below 1 MB, then this must be either the fixed DGROUP (data segment) or the fixed code segment or the TDB.

The following is the memory dump with the standard MYOWNVBX.VBX. Note that no entries appear for the control indicating that the it is not using any memory below 1MB.

BELOW1MB.EXE memory mapper output when using MYOWNVBX.VBX compiled with this Definition file :-

```
LIBRARY MYOWNVBX
EXETYPE WINDOWS
CODE PRELOAD MOVEABLE DISCARDABLE
DATA PRELOAD MOVEABLE SINGLE
HEAPSIZE 1024
EXPORTS
```

WEP PRIVATE

The following is the memory dump with the modified version of MYOWNVBX.VBX compiled with the following Definition file :-

```
LIBRARY MYOWNVBX
EXETYPE WINDOWS
CODE PRELOAD FIXED NONDISCARDABLE
DATA PRELOAD FIXED SINGLE
HEAPSIZE 1024
EXPORTS
    WEP PRIVATE
```

Free:

```
00030BA0 0020 00032 bytes
00055B00 0020 00032 bytes
00058E40 0060 00096 bytes
0008B720 0020 00032 bytes
0009FB00 0100 00256 bytes
```

Free: (Note not a single 512 byte block available)

```
00058E40 0060 00096 bytes
0008B6C0 0080 00128 bytes
0009FAE0 0120 00288 bytes
```

Used:

```
0007C420 2f60 12128 bytes (MYOWNVBX CODE Fixed)
0007F380 0fa0 04000 bytes (MYOWNVBX DGROUP Fixed)
```

Now the following two memory blocks are allocated to MYOWNVBX:

```
0007C420 2f60 12128 bytes (MYOWNVBX CODE Fixed)
0007F380 0fa0 04000 bytes (MYOWNVBX DGROUP Fixed)
```

Which are clearly the CODE and DATA segments for MYOWNVBX.

I have written code that forces Q+E multink VB VBX and DLLs to load in memory above one meg and the

Q+E multilink works fine. I have examined the behavior of the Q+E multi-link controls and DLLs are run time and do not they they use any real mode code at all, in which case they should not be load in low memory.

Let me know if you want a copy of the code to do this to include in your samples.

These are some of the offending DLLs/VBXs.

```
0002E760 0020 00032 bytes (QMUTL03 CODE Fixed)
0002FD00 0060 00096 bytes (QMBLB CODE Fixed)
00030B40 0060 00096 bytes (QMMDS03 CODE Fixed)
00054880 0500 01280 bytes (QMUTL03 DGROUPE Fixed)
00054D80 03c0 00960 bytes (QMMDS03 DGROUPE Fixed)
00055140 0060 00096 bytes (QMSQL03 CODE Fixed)
000551A0 02a0 00672 bytes (QMGUI03 DGROUPE Fixed)
00055440 0040 00064 bytes (W3DBLIB CODE Fixed)
00055480 0500 01280 bytes (QEUTL03 DGROUPE Fixed)
0007C420 9a40 39488 bytes (QESS03 CODE Fixed)
00085E60 1660 05728 bytes (QESS03 DGROUPE Fixed)
000874C0 2040 08256 bytes (QEBAS03 DGROUPE Fixed)
00089500 0a60 02656 bytes (QEMDS03 CODE Fixed)
00089F60 03c0 00960 bytes (QEMDS03 DGROUPE Fixed)
0008A320 0e40 03648 bytes (ODBCINST DGROUPE Fixed)
0008C900 f560 62816 bytes (QEBAS03 CODE Fixed)
```

Problem: Have you tested Multilink/VB. with Sybase Omni Server?

Q+E Response: At the current time our QA division has not tested Multilink against Sybase Omni Server, but this should not cause any foreseeable problems.

Problem Occasional GPFs Design Time Only

If an application is loaded from the network (Netware 3.11 /3.12) we sometimes get GPF [QELINK.VBX 0046:00D3] with our applications. We are using the latest ODI SPX/IPX stacks from WinUp09, on 3com etherlink III cards running normal VGA resolution on Compaq Laptop and Desktop 486 PCs.

This also happens setting the pQuery on a DBCombo (Style=ListBox) by double clicking to use the query builder. I have to type in the entry to get round this. [QELINK.VBX 0046:0306]

Problems with Message Passing in VB. with Active Q+E Controls

The Q+E controls seem to have problems with some Windows messages. For example:

Setting the Max length via a Windows message (when the Q+E Max Length is zero: Disabled) The same message is processed correctly for regular text boxes.

The right mouse click message is not processed correctly in Windows on a VB. form that has active Q+E controls on it. Everything is fine until the query to which the control are bound becomes active. After this the right mouse message is not always processed correctly. This means that it is not possible to develop applications that use the right mouse click messages while there are active Q+E controls on the form.

This looks like the unrecognised messages are not being passed along the message chain.

Record locking with Sybase using Q+E

The automatic record locking in Q+E is OK for ISAM work but for large client server work it is not ideal. What we have done is to use the techniques outlined in this section

Q+E locking of records can occasionally take several seconds. This happens even when the Lock on Keypress option is disabled.

What we want is explicit / transactional locking (which the fLock and fUnlock functions provide) and optimistic concurrency control which re-reads a modified record, and warns users that the data they tried to save has been changed. While the documentation is clear about how to lock and unlock records, implementing optimistic concurrency control is not clearly documented.

```
Connect.pOneHstmtOptions = Connect.pOneHstmtOptions Or  
QE_NO_COMPARE_AFTER_LOCK Or QE_NO_LOCK_ON_KEY_PRESS Or QE_HSTMT_NONLOCAL
```

[This sets the bit mask that Q+E uses to determining locking behaviour. Remember that in VB. if you get the variable name wrong and are not explicitly defining the variables, VB. will create a variant. If you set option explicit on and check that your connect name is correct this should not happen.]

OK that disables the client based locking that Q+E performs.

Next make sure that the pKey value of ALL your DBText controls are set to false. This will stop Q+E trying to use primary key information for locking, and speed up the app / stop Q+E getting this information it does not need!

Now the client software has been told not to do any locking. Sybase will perform all the locking etc. Finally you will have to trap the errors that can occur in your query, there are only two possible errors;

QE_RECORD_LOCKED

Record locked by another user

QE_RECORD_CHANGED

Record updated by another user since it was read

Use the **eError** event in your query to trap these errors gracefully. (Call functions in a module if you want to apply a standard response to these errors in VB.)

This will give you very fast locking (done by Sybase) while allowing you to trap lock related errors reported by Sybase to Q+E.

Enhancement Requests

pRecCount and use of dbCount

The pRecCount property does not seem to use the dbCount feature in dbLib. This returns the # records returned / updated/ deleted etc. This is MUCH faster than scanning all the records into a log file or even doing a `count (*)` as dbLib and Sybase give this information out with all successful calls to dbLib.

Q+E Response pRecCount does not take advantage of the dbCount feature. This as ENH # 58057.

Problem: dbLib requires all records to be retrieved to set dbCount!

Sub Queries in Find Mode

In find Mode, fields linked from another query are still live. While it is simple to turn off the second (linked) query during a find. Code must be written to display Not a search field and mark the control as read only. This takes time and could be done by ML VB.

Find Mode Mask for Numbers

using 99999 as a pMask means that Sybase / SQL server is handed a load of underscore characters `_`. This makes the pMask unsuitable for numbers. I have resorted to using `#####` as the format string. This is NOT what I want but at least it works without introducing spurious pad characters. If you must have a pad property then introduce a user definable property of setting it. Note the MS Pad has 2 properties one padded on not padded.

Setting the pMask Property

AA or larger causes QELINK to GPF during an edit of the custom control at design time, Even though the control operates correctly at run time. I have tried this with version 2.0.0, 2.0.2 and 2.0.3. Once in this state the control cannot be edited (that causes a GPF.) [Usual error codes are GPF at QELINK.VBX then (0002:19A9,0002:1A86)]

Non Existent Connect Control

Also running a program with queries that refer to a connect control that does not exist causes a GPF (in VB.!) using QELINK.VBX [version 2.0.2] switching to version 2.0.0 reports the error correctly.

Out of date Files on BBS

Several of the example files and demos on Q+Es BBS are out of date and actually wrong: They tell you how to work around problems that have been fixed by Multi-link VB. version 2.0.0 Having old files like this on the BBS is not very helpful. Are these old versions of Multi-link VB. still current supported products? I think you BBS is very useful and being able to download examples and new files is very useful, marred only by a few old (redundant?) .

Poor Quality / Old Samples

The sample programs that ship with Multi-link VB. are very poor and do not do justice to the product. The examples still use VB2.

ODBC Driver Test Program

ODBC Compliance include the source code for this program and elaborate on the required support for Q+E multi-link programs.

Q+E Tips

How do I disable record locking?

Add the statement

```
LOCKONKEYPRESS=0
```

to the **pLogonInfo** property of the Connect control eliminates many of these read ahead issues. Of course, it eliminates the automatic locking Q+E does to prevent concurrent updates of the record you are editing. This may cause concurrency problems, depending on the design of your application.

pMode property now returns 3, this is not documented?

Are the new value documented?

From the ML Help file qelink.hlp

- 0-Query or connection is not active
- 1-Query or connection is active
- 2-QBE mode is active (Query control only)
- 3-Find mode (Query control only)

Are there any read/write grids for VB. that Q+E have in the pipeline?

Q+E Response: No. Intersolv has not released this information to any third party developers, although it is being discussed.

Which Files are required for a Compiled VB. Application that uses Multi-link VB.?

This is an update to the Q+E Multilink/VB. Users Guide p. xvii. The following Multilink/VB. files are required for users to run your application:

```
QELINK.VBX
QEMLRUN.LIC
CTL3D.DLL and CTL3DV2.DLL
ODBC.DLL
ODBCINST.DLL
QMSQL03.DLL
QMGUI03.DLL
QMBLB.DLL
QMUTL03.DLL and QMUTL04.DLL
QMBAS03.DLL and QMBAS04.DLL
QMMDS03.DLL and QMMDS04.DLL
```

The executable form of your application (a file with an .EXE extension) Database drivers, named

QMxxx04.DLL, for the database systems to which the application connects. These drivers are listed in Guidelines for Distributing Applications in the on-line help file, QELINK.HLP QMFLT03.DLL and QMFLT04.DLL, only if the application you are distributing connects to Btrieve, dBase, Excel files, Paradox, or Text files. An ODBC.INI must either exist in the windows directory or must be added as part of the distribution. (If it already exists, it must be configured properly for the database drivers being installed.)

Debugging

This section contains tips for debugging ODBC / dbLib problems to see what is going on between the client and the server. It does not help with all problems but it is easier than breaking out a LAN Sniffer and decoding the packets on the WAN!! Which what I had to do to find some obscure bugs.

If your ODBC driver supports use the ODBC Admin. app to capture the ODBC stream to a LOG file to help you debug what is going on. This may help you figure out if the ODBC driver is the problem, or maybe the network connection, etc..

X-Ray SQL / X-Ray ODBC

These are two really useful products for find faults in code that uses dbLib and also for monitoring ODBC activity.

QEML2 Trace Option

To turn on qelib tracing, create a text file named QMBLB.INI and put the following in it:

```
[QELIB]
TraceFile=c:\qetrace.txt
TraceOptions=41
```

The Tracefile= can be set to any file name. That is the file that will contain the output. It will also be displayed on the screen.

Put the qmblb.ini file in the windows directory and run Multi-Link. You have to reload the Multi-link VB VBX or restart windows. Note that is technique also produces a dump to a debug window on the screen.

Trace option in ML/2

To turn of the qelib tracing, just move or rename the qmblb.ini file.

Common Questions

- [Q. Can I run NetLib, OpenClient, and DBLib concurrently on my workstation?](#)
- [Q. Can I use TCP/IP or SPX/IPX with Sybase / MS SQL Server?](#)
- [Q. Do I need ClientLib?](#)
- [Q. Do I need dbLib / Whose dbLib should I use, Microsoft or Sybase?](#)
- [Q. How can I restrict SELECTs, what do I do when I cant?](#)
- [Q. How can I stop using the Jet Engine?](#)
- [Q. How can I view the SQL being sent to the Sever / The servers response?](#)
- [Q. How can I use Combo Boxes in Client Server applications?](#)
- [Q. How do I clear locks?](#)
- [Q. Ideally, Id like to develop under SQLServer 4.21 and then just drop it into the System10 world and have it work can this be done?](#)
- [Q. ODBC Vs dbLib Performance When & What?](#)
- [Q. What are Dirty Reads?](#)
- [Q. What are Locks?](#)
- [Q. What are the benefits of Q+E ODBC drivers?](#)
- [Q. What does the Q+E ODBC pack offer me?](#)
- [Q. What is dbLib / NetLib do I need them?](#)
- [Q. What is NETAPI.DLL?](#)
- [Q. Why are all these different locks required?](#)

References

- [Q. Using the standard VB3 ODBC drivers seems to be working fine on the SQLServer, but can I use this software with a system-10 database?](#)
- [Q. Is there anyone out there who is using VB3 Pro with Sybase System 10? Any problems? What are the various pieces that you are using?](#)
- [Q. Can I replace the MS ODBC layer with Q+Es ? Can I use the VB3 Pro data access functionality through the Q+E driver?](#)

Q. Our application is a Windows application and can recognize other data files if an ODBC driver is installed. The problem is this: When our application attempts to access a Sybase data file, a message returns stating 'Context changed from...'

Q. How can I keep my application from reading through the entire result set when viewing data? This causes my workstation to temporarily lock up. I'd like to immediately close the query windows when I ask it to.

Q. Now, lets assume that I want to do something with the Q+E driver and that I have the latest and greatest Q+E Driver Pack (along with Multilink). Ive pretty well figured out that, while I might want to do some nifty stuff with some of the Multilink data-aware controls at some point, I really do not want to replace all of my Dynasets with Q+E data access functions?

referencesDo I need dbLib / Whose dbLib should I use, Microsoft or Sybase?

When you use an ODBC driver for Sybase it may or may not use dbLib on the client PC. However both the Q+E and Microsofts implementations MUST have dbLib (and an appropriate NetLib) on the client. This means that they all require a licensed copy of dbLib and NetLib. If you want to know what the Q+E ODBC drivers give you over MS when used with VB. apps, the answer depends on the SERVER PLATFORM in question. Microsofts implementation of ODBC for MS SQL Server is fine for NT Hosted MS SQL Server sessions. However if you are using Sybase 4.3 or later (e.g. 4.9.2) then you should consider the following;

dbLib version 4.25 is the current version supported by Sybase. This includes all the features in Microsofts version of dbLib that ships with SQLServer (MSs dbLib is currently version 4.21). The Sybase version of dbLib supports all the current version of MS SQLServer AND all current version of Sybase 4,x and System 10, SQLServer, Open server and Omni server over all transports that Sybase support. So the take home message is if you are a 100% Microsoft SQL sever shop, then use the MS version of dbLib, which they licence from Sybase. If you are a heterogeneous site: Sybase and MS SQLServer etc. then use the Sybase version of dbLib currently version 2.5.

Remember to check all licencing agreements before installing any software.

See also;

[What is NetLib?](#)

[What is NETAPI?](#)

[What is ClientLib?](#)

[Problems with dbLib under Windows...](#)

What is NetLib do I need it?

Check with your network vendor and SQL Server Vendor about the issues relating to using this protocol for your application. The information provided here is only intended as background information, to show that the network may have a significant impact on client and server performance and scalability.

Connecting Client PC to SQL Server is an important and non trivial task. If you do not know about the networking issues (monolithic Vs ODI, routing, protocols, frame types, network numbers) it is well worth talking to your local network administrator or getting some books on the subject because the network is your applications link the SQL server. If there are problems with the network layer then your application and SQL Server may have problems connecting , intermittently, and it is common for network faults (e.g. not running with an old of date driver) to crash dbLib and or any part of the ODBC layer.

All versions of dbLib need a NetLib to work. NetLib is the part of dbLib that connects your apps to the ether, without this you cannot be connected! NetLib depends on the transport stack on your PC and the transport stack on the server, the 2 must be compatible.

Finally you can run multiple protocols concurrently on the same PC and connect to several servers using different network transports, so you may not have to chose between protocols but can adopt a mix and match approach, depending on your requires.

SPX/IPX

SPX/IPX is the network transport used by Novell Netware networks. Most Sybase and SQL server implementation now support this transport. While ideal for national networks SPX/IPX can be hard to manage internationally.

Microsoft pay a royalty to Sybase for each copy of SQL Server they ship. The SPX/IPX parts of NetLib that ships with MS SQL Server for SPX/IPX is called DBMSSPX3.DLL these were developed by Microsoft, as were the SPX/IPX transport stacks for Windows NT (This allows Microsoft to have there own SPX/IPX technology, which makes the software work over Novell Networks without using and code from Novell). Other than its origins DBMSSPX3.DLL is just a SPX/IPX NetLib for dbLib, that defaults to Urgent data (allows dbCancel to be processed on SPX/IPX even though this transport does not support Urgent data, it is implemented via Packet Type 20 [internal SPX/IPX type]) and uses Dynamic name resolutions (which means that the address of the Server is not stored on the client, it is kept by all Novell Netware servers as the SQL Server on NT issues Service Advertiser Packets just like a Netware Server [very neat trick]). DBMSSPX3.DLL will work with either the Sybase dbLib 2.5 or Microsofts dbLib version 4.21. It allows a PC to talk via SPX/IPX to a MS SQL Server. (Note if you have a Sybase Server running on HP with an SPX/IPX stack you will need another NetLib. NOVSP.DLL. This is because HP use portable Netware to allow their server to talk SPX/IPX where Microsoft have their own SPX/IPX stacks on their NT servers. Also on the client side Sybase does not have an agreement to use DBMSSPX3.DLL rather the licence code in NOVSP.DLL from Novell. This means that you may have to provide an address for the server and specify URGENT data.

TCP/IP

The NetLib you need for TCP/IP depends on the implementation of TCP/IP that you are using. Contact Sybase or Microsoft to see if there is a version of NetLib available for the TCP/IP transport stack you plan to use. TCP/IP is a robust and mature inter-networking protocol. While much more difficult to manage than SPX/IPX it is ideally suited to global deployment.

If you have WFW 3.11 you can use the stack that comes with that. If you have Windows 3.1 and do not want to move (like me!) then buy a good TCP/IP implementation that will support Sybase. Check with Sybase on this. If you need it, remember that the TCP/IP stack you buy might have to work with both MS SQL Server 4.x and Sybase system 10 SQL Server: buy one that will work with both.

NetBUI

This includes MS Lan Manager, DecNet and other network systems. Check with your network vendor and

SQL Server Vendor about the issues relating to using this protocol for your application.

Common Network Problems

This is a quick guide to the most common problems to cope with multiple connection to SQL server it is not intended as a complete guide to networking issues with SQL Server, it makes assumptions about your network configuration. If you are unsure about your network contact your network administrator or supplier. As SQL Server is a single statement database multiple request must be sent through different connections. If you are using named pipes to communicate to the SQL Server is might be worth checking these setting. Make a note of the setting before making changes so that you can always go back to your original settings.

Novell

The Shell.CFG or NET.CFG file on the DOS workstations is used to configure Named Pipes options. Check these settings **prior** to loading NETX.

```
NP MAX OPEN NAMED PIPES = n
```

Where n is the number of pipes for the client workstation. To connect to 4 servers n would equal 4 or more.

```
NP MAX COMM BUFFERS = m
```

Where m is 2n.

LanMan

To configure a LAN MAN client to connect to multiple SQL Server connections ensure your LANMAN.INI file contains the following;

```
[Workstation]      Section Heading  
NumResources = n
```

Where n = 1 for each Named Pipe. To connect to 4 servers , n would be 4 (or higher)

Problems with dbLib under Windows

These problems relate to the 16 bit Windows 3.1 implementation of dbLib. The file from either vendor is called W3DBLIB.DLL.

The problems outlined occur in dbLib 4.21a (stamped (C) Copyright Microsoft) and dbLib 4.25 from Sybase, not stamped with any version or vendor information.

Calling dbWinExit generates an Invalid hWnd message in the debug window of MSVC version 1.5.

The structure of dbProcess and dbLogin still use near pointers. This means writing dbLib code is made much more difficult and error prone.

Sybases version for dbLib does *not* contain any version resources. This means that standard Windows version stamping and version control cannot be used in a production.

Do I need ClientLib?

This depends on what you want to do. If you want to access all the new features that Sybase system 10 can provide ClientLib is the way to go if you need multiple read cursors (useful for maintaining editable grids on large tables) asynchronous requests etc.. If you must have all these then you need ClientLib.

If you'd like to develop under SQLServer 4.21 and then just drop it into the System 10 world and have it work this can be done...

As you are developing for 2 platforms (4.x and System 10) you may have to accept the lowest common denominator approach. MS SQL Server 4.21 is Sybase 4.2 and System 10 is fully backward compatible with 4.2 and 4.9.2. All you have to do is use the appropriate drivers;

If you want to talk to either then get dbLib 4.25 from Sybase.

Can I run NetLib, Open Client, and dbLib concurrently on my workstation?

Yes, once configured everything will work fine.

MS SQL Server

A Low cost option for developing with Sybase 4.x is to but MS SQL Server. This is Sybase 4.2 it may not have all the new System 10 features but it still packs a punch.

Also MS SQL server supports SPX/IPX and TCP/IP out of the box.

MS SQL Server Version 4.2.1a and NTAS operating system 3.5 support SPX/IPX and TCP/IP.

A new version of SQL Server, SQL Sever 95 is currently in beta. This will offers many of the Sybase System 10 features at a fraction of the cost.

Sybase is aimed more suitable for high end applications. MS SQLServer is more suitable for smaller developments (<500 users, <400 TPS etc..)

Using the standard VB.3 ODBC drivers seems to be working fine on the SQLServer, but can I use this software with a System 10 database?

There is no problem using the existing ODBC drivers (i.e. the ones that ship with VB. Pro 3.0) and ODBC drivers with Sybase System 10, as System 10 supports dbLib. You will not benefit from the new features in System 10 (multiple read cursors etc.) but your existing software (the drivers that ship with VB. use dbLib) will work.

Is there anyone out there who is using VB.3 Pro with Sybase System 10? Any problems? What are the various pieces that you are using?

Use System 10 on a development system. I have used;

VB. with MS ODBC and MS data aware controls (all shipped with VB. 3.0 Pro).

VB. with Q+E ODBC pack and MS data aware controls.

VB. with Q+E Multi-link VB. and Multi-link ODBC drivers for Sybase SQL Server.

All of the above with Sybase System 10 drivers from the Q+E.

There are currently a few problems with ODBC to System 10 these are due to various problems not least of which are bugs in system 10.

Now, lets assume that I want to do something with the Q+E driver and that I have the latest and greatest Q+E Driver Pack (along with Multilink). Ive pretty well figured out that, while I might want to do some nifty stuff with some of the Multilink data-aware controls at some point, I really do not want to replace all of my Dynasets with Q+E data access functions?

You do not have to change anything. All your Dynaset and MS data aware controls will work with Q+E ODBC drivers. All you are doing is replacing the ODBC drivers form Microsoft with those from Q+E. You have a choice 4.x drivers or System 10. Either way you do NOT have to re-code (or recompile). You may have to modify any installation software to include the new drivers though. You can then add the Q+E controls as your program develops. This way the old and the new can co-exist. If you want you can standardise on one or the other or keep both its up to you.

What does the Q+E ODBC pack offer me?

The Q+E ODBC drivers (or any ODBC driver) offer an alternative to the ODBC drivers for Sybase that ship with VB. Pro 3.0. Notice I said offer an alternative because you can have the drivers co-existing without any problems. (E.g. App A use driver X and App B uses driver Y, even at the same time) If you want to undertake client server development take a look at Q+E Multilink VB., this is a VBX and a set of ODBC drivers (that carry a per seat licence / royalty free) that allow you to develop VB. apps. These controls do not use the same data controls as MS. Rather Q+E controls use a different model, which seems far better thought out.

Can I replace the MS ODBC layer with Q+Es ? Can I use the VB.3 Pro data access functionality through the Q+E driver?

Yes you can replace the MS ODBC driver with the Q+E driver you will not have to change or recompile any code.

What are the benefits of Q+E ODBC drivers?

Try the MS ODBC drivers and Q+E I have found the Q+E driver faster and more reliable than the equivalent MS offerings.

Getting around the Jet Engine ✈

The functionality and easy of use of SQL Server can be severely limited if you connect via Jet. A way around this is to code around the jet engine by keeping the SQL extremely simplistic (time consuming and not guaranteed to work with later releases of the jet engine. I found I was coding around the Jet engine more time than developing apps!)



Jet is an optimiser for ISAM databases. It can make modify SQL on the way to the server. Unlike the ODBC layer and dbLib that only translate APIs the Jet can modify the SQL. SQL Server has its own optimiser (which does not support tree pruning) this make is more sensitive to changes in the SQL than another DBMS (say Oracle).

Alternatively use Jet Independent bound controls ODBC implementation (such as Q+E ODBC drivers and Multilink VB or the ADE ODBC Tools) that allow call the ODBC API direct. This eliminates all of the compromises introduced by Jet layer.

How can I stop using the Jet Engine?

If you...

Do not use any MS data aware controls

Replace bound controls with the bound controls with [ADE ODBC Tools](#)

Use the [Q+E Multi-link](#) controls instead,

Call the ODBC API directly

Call dbLib via an [intermediate VBX](#)

Then the software you produce will not use the Jet Engine. These are not alternatives they are options! You can also use them in conjunction with bound control.

Why would you want to?

Performance, bypassing Jet makes things run much faster (e.g. ODBC connecting to a server less than 3 seconds, and data retrieval etc. is significantly faster.) I know that there are work around for the Jet Layer like SQL_PASSTHROUGH. But some of the work-arounds are specific to the version of Jet and generally make the SQL / Code less readable.

It turns out that once you open a connection through the database objects, that VB may not release that session even if you "close" the DB object and set it to "NULL". This may be a "feature" of VB and the Jet Engine. Perhaps the thinking was that you might wish to open the same connection latter in the program. Maybe this is why it is kept open for you in the background. This may save needless time of re-connecting!

If you have examples of code that uses Jet quickly please let me know.

What is NETAPI.DLL?



NETAPI.DLL comes with the Netware windows kit although in the updates it comes in a directory called 'OLD' which suggests it may not be around long or at least not supported for long. The documentation (in the OLD directory) says it is an API to the SQL Server front-end, looking at the functions in it they seem to implement the named pipes interface. From what I remember to use it there is also a DOS TSR which implements these named pipe functions over SPX.

The functions in there are as follows..

Exports:

ord seg offset name

1	1	11f9	WEP exported, shared data
532	1	0058	DOSDISCONNECTNMPPIPE exported, shared data
533	1	007c	DOSSETNMPHANDSTATE exported, shared data
534	1	01fe	DOSWAITNMPPIPE exported, shared data
535	1	033e	DOSTRANSACTNMPPIPE exported, shared data
527	1	0f8a	NETSERVERENUM exported, shared data
536	1	05f6	DOSQNMPHANDSTATE exported, shared data
530	1	0010	DOSMAKENMPPIPE exported, shared data
537	1	078e	DOSQNMPIPEINFO exported, shared data
531	1	0034	DOSCONNECTNMPPIPE exported, shared data
538	1	0992	DOSPEEKNMPPIPE exported, shared data
539	1	0bd2	DOSCALLNMPPIPE exported, shared data

ODBC Vs dbLib Performance When & What



The most common questions I am asked are...

I understand that ODBC doesn't have all the features found in the DB-Library, but what about performance issues on the functions that are the same?

Are there any major gotchas when using ODBC on a client (like the cost for the ODBC client side library is X times the cost of the Sybase library, or something of that nature). I would like to hear any pros/cons?

Many people seem to be confused about ODBC and dbLib. Using ODBC does not mean you can't use dbLib! The two are not mutually exclusive. In fact ODBC is only a wrapper for dbLib. I use ODBC for all my form based work and tabular data, and dbLib for quick lookups (e.g. For filling combos, call SP, general purpose SQL, etc. where the benefits of ODBC are not required. Most people try the ODBC drivers and controls that ship for free with VB. / Access and base their opinion of ODBC on this. Unfortunately this set-up comes with a layer of software called the jet engine. This ISAM optimizer generates unneeded / slow SQL. Also the controls that MS ship are not ideal for client server work. Look at some of the other commercial ODBC vendors and tools.

There are a number of issues that need to be clarified about Accessing Sybase using ODBC. Firstly ODBC actually uses dbLib to access Sybase. (If you do not believe this look at the entries in ODBC DLLs or just delete the W3DBLIB.DLL from your system and try connecting to a SQLServer.) If you have System10 then you can use client library. dbLib and Client lib are the only ways to communicate with SQLServer. By the way the only ODBC driver for System10 is produced by Q+E. If you use any other it will be talking to Sybase using dbLib: 4.x compatibility mode, so you won't see the benefits of System10, multiple read /write cursors etc..

ODBC Pros

Wide support for applications, e.g. Access, Excel, Word, etc.

There are several vendors of ODBC, ODBC is a specification NOT a product.

If you do not like the restrictions of a manufacturer then you can switch vendors or mix and match. We use MS for our Office users (Word, Excel, Access etc.) and we use the Q+E multi-link ODBC drivers for application we develop. The later are much faster (under 2 seconds to connect to a server) and provide real performance benefits especially on large tables and optimistic locking, full low level (Sybase) error handling, multiple connections etc...

While dbLib is great for simple SQL selects and updates it does not really lend itself to form or grid based applications. Such features have to be implemented by the programmer. This is time consuming and is part of the ODBC specification. Using ODBC means that this sort of ground work has been done for you. ODBC is not a complete answer, because dbLib is still really good for very simple things like populating a list box or quick lookups etc. It really depends on the type of application that you are developing.

ODBC also makes your application relatively database independent.

ODBC Cons

The MS implementation via Access or Visual Basic using a layer known as the Jet engine. This is designed to optimize SQL for Access (ISAM) this is a real problem when running against an SQL server. The reason for this is that the SQL for SQLServer tends to be much more complex than the SQL sent to Access. The result is that only very simple SQL can be sent through the Jet Engine. Further the Optimization in Jet actually slows down the execution. The result is that developers spend time trying to bypass the jet engine. The optimizer is jet version specific. This means that when subsequent versions of Jet come out you may have to re-code your SQL. Some people say that calling stored procedures ensures that the Jet engine won't get in the way, while this is true, I think it speaks volumes about the problems using MS ODBC. If you use an alternative ODBC layer this non-sense disappears. I have written code using MS ODBC and

Q+E ODBC, the latter does not use the jet engine, so works without any silly optimisation problems and seems to much faster.

ODBC Driver for Sybase 4.x

Much of the functionality of dbLib is NOT used by the current Q+E ODBC drivers. For example dbLib returns the number of records retrieved (e.g. matches in a select) updated etc. This information is returned to the client but never used. In fact using Q+E pRecCount function retrieves all the records to the client, and then counts them!! This and other anomalies are outlined in the problems with Q+E Multilink VB & Sybase section.

ODBC Drivers for System 10.0.2

The Q+E ODBC drivers for System still not work reliably (despite several updates and patches) so we are forced to use the 4.x ODBC drivers against a System server. This means that all the benefits that system 10 confers e.g. Notification, multiple cursors etc. Are not available at ODBC level. This sort of functionality could solve many of the problems we have using Q+E against Sybase 4.x. The ODBC drivers make a huge difference to the performance of a client applications. For example the logging problem that we have with Q+E against Sybase does NOT exist if we use the same client software with the Q+E ODBC driver for Oracle. In fact the Oracle ODBC services seem much faster even on large tables. This is not because of the back-end being faster, but because features like multiple cursors are used by the Oracle ODBC layer for editable grids. If we had a working ODBC layer for system then we could see these sort of improvements using Sybase.

To be fair some of these problems are due to defects in CT-Lib.

ODBC to OMNI Server

This is supported Q+E do not seem have tested this. As there is no SP support on an OMNI server and some system tables are different there might be problems if your code or your ODBC driver used them.

Local Storage of Configuration & Drivers

The network address and login details for the ODBC data sources are stored locally. This might be fine when you have one or two servers. However managing a large number of server requires special software that centralises the management of ODBC information and drivers. Updating the drivers for ODBC is an on going process as new version of the drivers are shipped. It is not possible to have the configuration stored on a network. This makes managing the ODBC software, its installation and maintenance and important task. There are packages that allow you to maintain a central set-up for all your ODBC configuration and allow using to install and update the server definition and drivers automatically.

Learning Curve

ODBC is seen by many as a simple to use connection mechanism. It is. Many mistakenly assume that it is simple to implement and use. At a simple level this is true but there are a great deal of internal functions that need to be considered especially where performance is important.

I use both ODBC and dbLib at the same time within application this allows me to use the best system for the task at hand. (E.g. Complex grid view populated in seconds by ODBC, and drop down list boxes filled in under a second using dbLib.) This gives me the best of both worlds.

References



The ODBC 2.0 Programmer Reference and SDK Guide ISBN 1-55615-658-8 from Microsoft Press.

Sybase InterNet FAQ: Command questions and answers relating to all aspects of Sybase use. Useful for MS SQL Server users. (Location:)

Displaying information About Locking

Locking is the tem used to mediate contension for database resources, this may be between users or even a single use accessing several databases or tables.

To get a report on the locks currently being held on SQL Server, use the ***splock*** system procedure:

Type ***sp_lock*** in ***iSQL***

The results might look like this;

SPID	LOCKTYPE	TABLE_ID	PAGE	DBNAME
4	Sh_intent	1987902549	0	static
4	Sh_page	1987902549	3325	static
9	Sh_intent	2044534317	0	master
9	Ex_extent	0	296	tempdb
13	Sh_intent	1037962774	0	hacks
13	Sh_intent	310044636	0	static
13	Sh_intent	1827901979	0	static
13	Sh_intent	1987902549	0	static
13	Ex_extent	0	40	tempdb
13	Ex_extent	0	272	tempdb
13	Ex_extent	0	280	tempdb
13	Ex_extent	0	288	tempdb
18	Sh_intent	1987902549	0	static
18	Sh_page	1987902549	3322	static

The *locktype* column indicates not only whether the lock is a shared lock (*Sh*) an exclusive lock (*Ex*), or an *Update* lock, it also tells whether the lock is held on a table, on a page, or on an extent, and whether the lock is blocking (*blk*) another process or is attempting to acquire an exclusive lock (*demand*).

See Also: [Clearing Locks](#)

What are all these lock types and what do they mean?

In general, read operations acquire shared locks. A read lock allows other user to read the same data so read locks are often refered to as shared read locks.

Why does SQL Server place locks on reads being Read?

When you select records SQL server places a read lock on the records being selected. This is to make sure that no one can change these records while they are being read.

This simple read consistency system means that a read will never bring back data that has not been committed: No dirty reads.

What is a Dirty Read?

A dirty read is used to describe a database where no read consistency is enforced. In such a system one user may read a record that is being modified by another user. In some application that do not require read consistency this behaviours maybe acceptable and even desirable. MS SQL Server 95 and new releases of Sybase System 10 may allow the *option* of dirty reads. The ability to use dirty reads depends on the type of application you are developing. Dirty reads mean less contention caused by read operations. Dirty read also mean that SELECTs may fetch information has not been committed (and may never be!): Things that are not in the table.



Write operations acquire exclusive locks: Only one person can update a record at a time. Update locks are created at the page level. There is no row level locking in SQL Server, this *may* be an option on System 11. This means that unrelated information may be locked. In practice this is not usually a problem. The reason that row level locking is not used is for performance. Calculating the row boundaries within a page takes time. This strategy allows fast locks but can lead to the situation where one users exclusive lock acquired during an update locks a page containing another record from another completely unrelated table! In such cases the second users process will block until the unrelated update has completed and release the exclusive lock. This is not a bug a part of the design of SQL Server. Usually the speed conferred by not having to calculate record boundaries within a page mean that locks are acquired, updates applied and locks release, very quickly.

Mediating Contention



Update locks are acquired during the early part of an update operation while the pages are being read. The update locks are compliant with shared locks. Later, if pages are changed, the update locks are promoted to exclusive locks. This prevent unnecessary exclusive lock being placed and reduce the contention that this would generate.

An intent lock indicates the intention to acquire a shared or exclusive lock on a data page. Setting an intent lock prevents another transaction from subsequently acquiring an exclusive lock on the table that contains that page.

An extent lock is a lock held on a group of eight database pages while they are being allocated or freed. Extent locks are set during a CREATE or DROP statement, or during an INSERT operation that requires new data or index pages.

The *blk* suffix indicates that one process is blocking another process that needs to acquire a lock. As soon as the blocking process completes, the other processes proceeds.

A Demand lock prevents additional shared locks from being set. It indicates that a transaction is next in line to lock a table or page, because shared locks can overlap so that read transactions keep monopolising a table or page, forcing a write transaction to wait indefinitely.

An example of this is an update (in a transaction) being blocked by several SELECTs say an update is waiting for change the address of a customer in a table that is constantly being queried by other users looking up customer addresses. After waiting on four different read transactions, a write transaction is given a demand lock. This mechanism allows write operations to pulling rank over read operations. As soon as the existing read transactions finish, the write transaction is allowed to proceed. Any new read transactions have to wait for the write transaction to finish and release its exclusive lock.

Clearing Locks

In a well design system locks and contention for records are kept to a minimum. This is a very important part of the database design. This depends on good logical and physical database design. Many of the standard approaches (cookbook style database design) do not contention into consideration. This is often due to their theoretical background. I think this sort of approach produces simplistic database design which can lead to unnecessary contention caused by unthinking dogma e.g. Use 3rd normal form, primary keys etc.. and you cant go wrong attitude. Good database design is based on simple (not simplistic) design, using a variety of techniques to design out (or greatly reduce) the incidence of any contention. This allows large high turnover system to be designed, built and operated with a high degrees of reliability, performance and uptime. these are real world demands that are often overlooked by general purpose database design books. This is not a comment on such books, indeed SQL Server is not a typical DBMS. If you are design and building using SQL Server is worth getting books that are specific to the version of SQL Server that you are using.



Despite all the best planning and careful database design there may be cases where you have to detect locks and remove the process that is causing problems. This can be caused by high volumes or simple a rouge process (e.g. a user issuing a nonsensical JOIN)

I you want to find out which user owns an SPID use the **splock** system procedure: Type `sp_lock` in **iSQL**

The results might look like this;

1	sleeping	sa	0	master	MIRRORHANDLER
2	sleeping	sa	0	master	CHECKPOINT SLEEP
3	sleeping	sa	0	master	LAZY WRITER

To clear rouge locks (once users have logged out) you may try to KILL the SPID (use this with care). Also remember that some processes (MIRROR HANDLER / CHECKPOINT / etc..) may not be killed. Even if you issue a KILL command these processes may not die. In some cases the only sure way to release locks it to take the server down.

How can I keep my application from reading through the entire result set when viewing data? This causes my workstation to temporarily lock up. I'd like to immediately close the query windows when I ask it to.

Sounds like you have not enabled URGENT (out of bound data). This is a transport level switch. Remember that the client and the server OS need to be able to process urgent data sent by dbLib / CTLib. Also the transport stack on the server has to pass this information to the server process via the server OS to the SQL Server process.

In practical terms you need to try adding a ,URGENT to the end of your SQL Server section of WIN.INI (or SERVER.INI).

If you are using ODBC as well then open ODBC Admin. and make sure that the Cancel behavior is set correctly.

There are a few server OS that cannot support dbCancel (e.g. VMS) however more OS support this feature than do not (e.g. Unix, Windows NT). The use of dbCancel makes canceling large requests **very** fast. It sounds like you do not have this facility enabled, so the entire results set is being processed: A loop calling dbNextRow until the entire result set is pulled back. Not only is this bad for the client (Machine locks up under 16-bit Windows, you are using up network bandwidth)

You can use the ODBC Log SQL or run a spy program, such as the excellent X-Ray SQL / X-Ray ODBC, [dbLib / ODBC spy programs](#) to see what is happening.

Our application is a Windows application and can recognize other data files if an ODBC driver is installed. The problem is this: When our application attempts to access a Sybase data file, a message returns stating 'Context changed from...'

When you connect to SQL Server via dbLib (which what your ODBC driver is doing) if a database is not specified (dbUse) you are defaulted to master. This is not an error as the severity returned by dbLib is zero.

You can specify the database in the ODBC driver. Alternatively use the database in the tablename (e.g. hack.dbo.credit_card_numbers) this will prevent the problem.

Combo Replacements

What we decided to do is not to use combo boxes because of

1. The large amount of data involved and pulling down the data for each combo was too high an overhead
2. Fetching the data in advance was no good as we have several hundred forms and anyway the data would be out of date
3. 64k limit combined with the fact that >64k of data is not really useful as it may be difficult to use. (e.g. For code expansion eg, component code to component description)
4. Multiline tabbed list boxes does not allow searching either component code or component description.

What we did was use the idea in some mail programs for address lists (where the user has to select a few names from a list of thousands).

We also added the ability to search on either component code or component description.

Finally we needed some that looked like a Windows control.

The results was two edit boxes separated by a button which looks like a dropdown button. The component code was coloured black on white text and we put description in blue. The user can type in either. Clicking the button generates SQL to search using the input from either component code or component description. Now were to put the results? We can't use a drop down because the information may be multi-columnar. The solution we adapted was a pop up Windows with a tabbed list box and an OK Cancel show all matching records.

Restricting SELECTs

While it is true that you should restrict the results of an SQL SELECT whenever possible. Having said this this is not always possible.

If you are using SQL Server then there is a dbCancel function. This allows you to get say 100,000 records and populate say a list box or grid control. Once you have enough to populate the visible part of the screen you can paint the client region. You can fetch the pending results in background to maintain an off screen buffer. If the user scrolls down and you run out of off screen buffer then issue an hourglass while the dbNetrow loop catches up.

If the user decides to view other details or closes the app / window then issue a dbCancel.

We do this with a large customer database (100,000 records plus) accessed from all over Europe. This seems to work well. Remember that SQL server has contention mediation to cope with updates on tables with read locks (intent lock escalation et al.). This means that updates etc. should not be a problem.

