ISUG TECHNICAL JOURNAL

JANUARY 2011

A PUBLICATION FOR USERS OF SYBASE, INC. PRODUCTS AND SERVICES

INSIDE

Information for Sybase Users Including Database Tools, Mobility and IT Strategy

PRESIDENT'S MESSAGE

Dear ISUG Member,

It's a new year and it might be a terminal cliché to say it, but this is an exciting time of immense change. Fresh economic horizons seem to be coming into view, new cloud computing delivery models are starting to gain significant traction and Sybase is now firmly on a path of dual-level integration with regard to its technology and that of its new parent company SAP.

ISUG itself is equally invigorated. We have seen the long-term benefits of prudent management by the ISUG board members to turn our organizational financial situation around. The last couple of years have been an arguably greater challenge for the organization that at any period in its past, but our recent Annual Member's meeting will pay testament to the fact that things are now looking dynamically positive.

To add to our strategic process and plans, we're also on the up with a new high-level technology partnership with HP. As our newest Exclusive ISUG Hardware Partner, HP brings a heritage in both systems management, ALM, laptop and server manufacturing and a not-inconsiderable software division that exists in its own right. We couldn't be more excited to be honest; this is a time of momentum, movement and dynamism.

But enough of my high spirits and ebullience. How about we take a look at the issue ahead to what we have inside? Well this issue we welcome Chris McCartney and Dave Putz's superb piece, "How To Take Full Advantage Of Deferred Compilation In ASE 15.x". In this feature Chris and Dave talk us through some of the finer points of ASE using their extensive combined experience with the technology.

Feeling hungry? Good – then you'll want to dive into, "Chomping On The Network With Wireshark [Part 2]" by Sybase's Paul Vero. In this second part of his feature Paul takes us on a further navigation into the Wireshark interface and sniffer output. Paul's colleague Wayne Smith is also featured this issue with a piece called "Using Data Access Application Blocks in Sybase ADO.NET Provider", which shows just how Wayne himself has progressed not only in terms of the breadth of his own technology experience – but also his competency when it comes to writing for us here at ISUG Technical Journal.

Consultant extraordinaire, Sybase guru and technology mover & shaker in his own right Jeff Garbus is with us again this issue – this time looking at a piece entitled "ASE Cluster Edition Connection Management". In this piece Jeff looks at all the connection information you will need in order to make the right decisions as well as some implementation tips.

Last, but oh so very much not least, is Eric Lai who holds a very special position at Sybase HQ as an ex-technology journalist now filling a role as a corporate communications and content specialist. Eric brings his skill to the written word to his piece entitled "Seven Semi-Fearless Predictions About Enterprise Mobility For 2011" – thanks Eric! Let's have more where that came from please.

Best Regards,
Bryan Enochs
President - ISUG

Oh Do Shut Up! The Cloud Is Finally Here!

Large companies are finally taking up the cloud. That was the headline on technology news website Computer Weekly (yes OK I do write for them too and this is a shameless plug!), on the evening of me sitting down to write my editor's letter for this issue of ISUG Technical Journal.

The reason I mention that we have some solid uptake of cloud computing delivery models among large organizations is that I was started to get a bit tired of the barrage of PR and marketing materials I was receiving; recently they have done nothing but talk around the so-called "challenges" associated with the cloud.

This latest iteration of on-demand IT delivery (and let's face it there have been others before) finally seems to have sorted much of its foundations out. The headline I was drawn to suggested that more than 92% of the largest enterprises have at least one cloud service in deployment; while 53% have more than six cloud services.

Of the 434 IT professionals surveyed, 68% said security was their biggest concern when moving to the cloud and 40% said they were worried about poor service quality. OK so there are still some big challenges ahead, but perhaps now we can start to discuss how the cloud model will impact Sybase's (and SAP's) technology stack. I certainly hope so.

Are you already working with cloud-driven technologies in your Sybase data environment? Please write to us and tell us what you're doing – the time for postulating is over and the time for cloud action is upon us now. Don't you agree?

Adrian Bridgwater
Editor

Table of Contents

FEATURES

DEPARTMENTS

ISUG Technical Journal

Managing Editor
Adrian Bridgwater
adrian@isug.com

Art Director
Teresa Foster

Senior Technical Editor
Anthony Mandic, Mandic Consulting, Pty Ltd.

Technical Editors
Mike Harrold, ISUG
Eric Miner, Client Services
Linda Morison, Lockheed Martin
Jerry Siegel, TeamSybase
Peter Thawley, Sybase

Chomping On The Network With Wireshark [Part 2]

By Paul Vero

Paul Vero has worked at Sybase for 15 years. He is a Principle Product Support Engineer in Broomfield, CO. supporting various Sybase Connectivity and Middleware products. He can be reached at paul.vero@sybase.com.

Last month I introduced you to Wireshark, an aptly named application used to capture network packets on a computer. In our case, we specifically want to capture the packets between a client, like an ODBC driver, and the ASE. Our goal is to analyze the TDS communication between the two end-points and to utilize this information for diagnostic purposes. Recall we can benefit by reading these packets for understanding some performance issues and voluminous result sets, seeking patterns and reading data unprintable in RIBO. Technical support might request a sniffer trace and you'll be totally prepared to provide one. Let's start with some further navigation of the Wireshark interface and sniffer output.

Mastering The Expanse Of Packets

Last month we learned how to set the capture where the Wireshark machine is one of the end points. We'd like to know what packets belong to the client and the server. I've provided a sample where Wireshark ran on the client machine, and the ASE churned away on my Solaris virtual machine. In **Figure 1**, I've pointed out the IP addresses and ports of the client and server. The ASE listener port is 1550, and IP address is 10.22.159.43. The ASE is typically the responder to the client requests so it's great to orientate yourself before diving into the mess and breaking things out. My client port is 1714 and IP address is 10.22.247.137.

The significance of knowing the source and destination is paramount to understanding the flow; realizing the requestor (client) and the responder (server). Referring again to **Figure 1**, look at the arrow to view the Src (source) Port and Dst (destination) Port. A sniffer trace tracking just one client is easy to follow. However, when there are multiple connections, things get a little crazy. We tend to follow one stream, one connection, but a sniffer trace involving multiple connections from the same client (like an IIS application, or an application server; both involve connection pooling) can be confusing. You can separate the different connections by selecting one of the packets belonging to a specific source port, right-click the entry, select "Follow TCP Stream" and all packets involved in that source stream are collected in the interface. You can save this file as a record of this port's history. Press the "Clear" button just below the menu and the full trace is restored. You can track another source port and repeat.

In this sample, let's look at the login sequence between the client and server. Before TDS comes into the

"The significance of knowing the source and destination is paramount to understanding the flow; realizing the requestor (client) and the responder (server)."

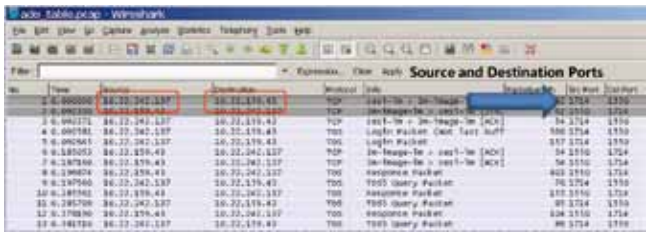

Figure 1: IP Addresses and Ports for Client/Server

picture the client and server need to establish the channel for communication. In the sample sniffer trace, we can look at packet 1 details by double clicking the line containing packet 1, expand the TCP/IP frame information like the frame in **Figure 2**:

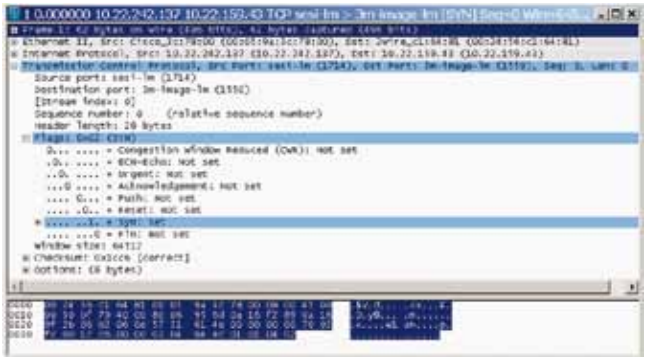

Figure 2: The Login Sequence

Starting with the window title bar, we see this is packet 1, at time 0.000000, and based on the first IP address, we know this is the client sending the packet to the ASE IP address. When a client talking the TDS lingo wants to establish a connection, it sets the SYN flag, meaning it wants to synchronize sequence numbers (where each successive packet contains a sequence number and this packet's relative sequence number is 0, it's next packet sent to the server is 1, then 2, etc…) with this destination. It's all part of the three way handshake. When the server side receives this packet it responds with its own SYN and ACK flag set. This means the server side will start its own set of relative sequence numbers, and acknowledges the client's SYN packet; which incidentally increments its acknowledgement number based on the received sequence number from the client. This is all seen in **Figure 3**. Now, it's the client's turn to finish the handshake.

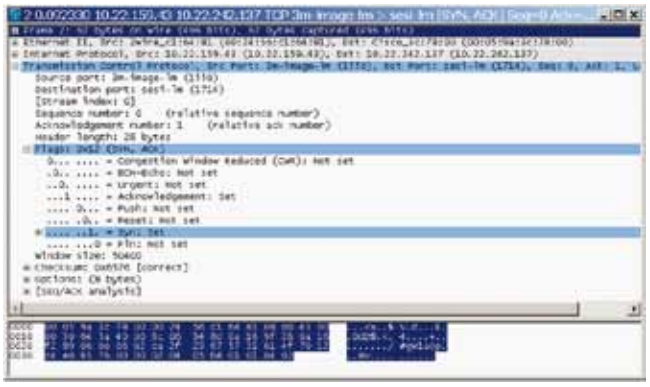

Figure 3: Server SYN-ACK Packet

And now for the moment you've been waiting for, the client ACK packet. After the client receives the server's SYN-ACK packet, it must oblige by sending an ACK packet (**Figure 4**) to the server, completing the three way handshake, establishing the connection so the TDS flow can start.

Figure 4: Client ACK, the Acknowledgement to the Three-way Handshake

Next in line is the TDS_LOGIN packet. After the handshake, it's time for the TDS channel to be established with a login packet which is sent to the ASE so the TDS can flow. This packet is identified under the Info column as "Login Packet (Not last buffer)". In a full-on sniffer trace, you can use the "Find Packet…" function under the *Edit* option from the menu, or use the *Control-F* key sequence, and a small window will appear as shown in **Figure 5**.

Figure 5: The Find Packet Function, Useful to Search Through a Large Trace

Since you want to search the packet list, select "Packet list", type in "Login" in the filter field, and *Find by String*. Click the *Find* button and the first "Login packet" is highlighted as shown in **Figure 6**.

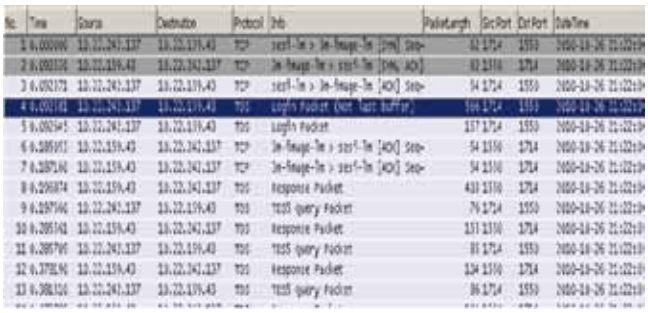

Figure 6: Login Packet is Found in the List.

In this sample, double click the login packet and you'll see the details. The TCP/IP protocol information is found along with the data. Important information such as sequence numbers, TDS type, flags, and length of data is displayed in readable format, seen in **Figure 7**. From the flags, you can see the ACK and PUSH flags are set; the ACK is a continuation of the client and server keeping in touch and the PUSH indicates the packet is ready to push on out.

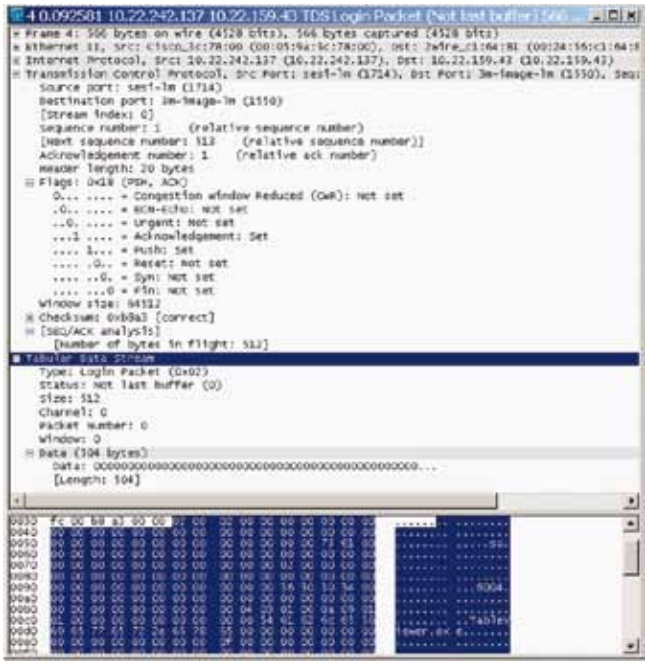

Figure 7: Login Packet Details

Looking back to **Figure 5** notice there is a second login packet. TDS clients always send the login packet in two packages, a 512 and a remaining amount, 103 bytes in this example.

Timing Is Everything

In the Time column, you can display this item in different formats depending on what you want to do. To check on the available options, from the menu select *View | Time Display Format*. The options are shown in **Figure 8**. As you can see, there is a rich variety of settings. In the current traces, I've selected seconds since the beginning of the capture, using the automatic precision: microseconds.

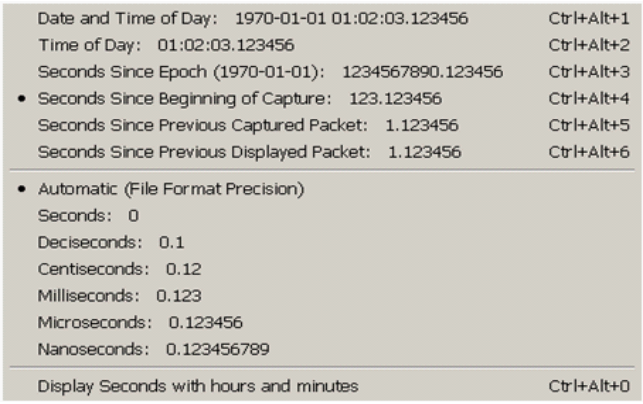

Figure 8: Time Display Format Options

Depending on what you're looking for, you can change these options to assist in timing issues. For example, you probably won't need microsecond precision, millisecond might be more practical. Also, maybe you're looking for packet response time from the ASE so you can select "Seconds Since Previous Captured Packet", a great way to observe how quickly ASE is responding.

Suppose you have some locking issues and you want to find out how long it's taking for ASE to respond to your query. In my example shown in **Figure 9**, I simulate a lock by locking a table manually and catching a trace of the waiting client. The TDS5 Query Packet was sent 7.605 seconds since the last packet, due to me setting up the lock and taking my time in executing this query. The ASE sends an ACK packet in 0.261 seconds, letting the client know it received the response, and yes I'll get back to you as soon as I can. In this case, it took 9.913 seconds after the ACK to send results, the "Response Packet", once the lock on the table released. Maybe the locks you encounter won't be as dramatic, but you get the picture. Notice in our locking sample, if the ASE isn't sending a response back in less than 0.20 seconds, it sends the ACK packet #21. If you see lots of these packets, you might want to investigate your ASE and perhaps give it a tune-up. Excessive ACKs from ASE immediately following the request (TDS5 Query) can indicate a performance issue.

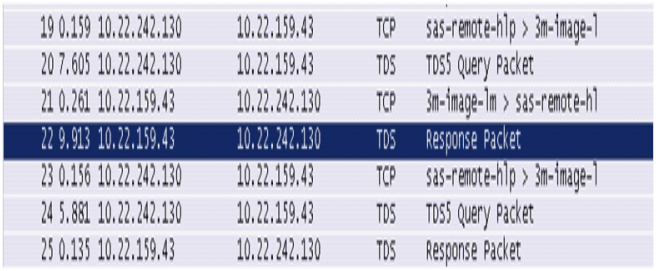

Figure 9: Long Response due To Lock

Tracking The Command

Recalling the "Finding Packet…" function, now is a good time to search for your command. If you have this humongous trace and know the precise table in a select statement, the search comes in handy. You simply enter the table name in the filter field, selecting "String" in the "By" section. Look for something unique. Keep in mind if there are schema calls preceding your select command they might match so control "F" again to continue your search.

In TDS, a select statement is commonly sent as a TDS_LANGUAGE token, displayed in **Figure 10**. The token is 0x21, and selecting the *Language* text highlights the data section. This technique is handy when you look at the response to identify the ROW_FORMAT data. **Figure 11** shows such an example.

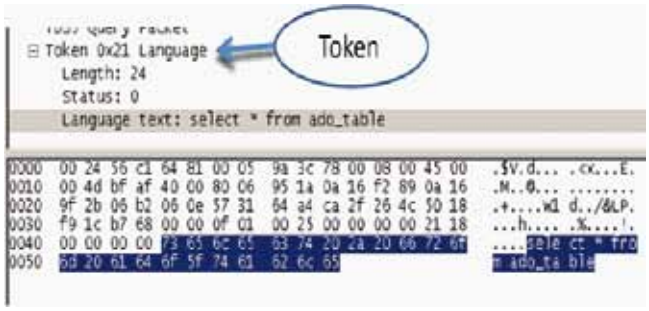

Figure 10: . Highlighting the Language Data

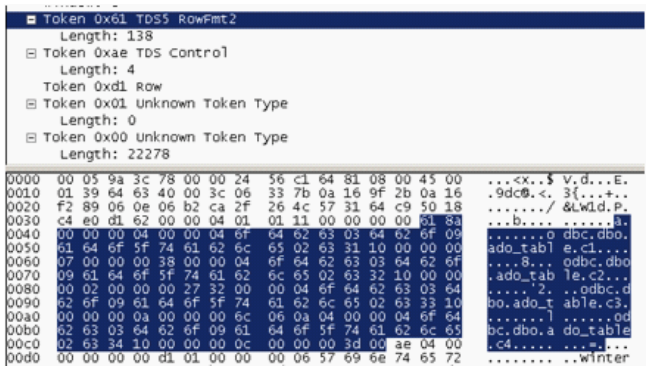

Figure 11: Identifying a Row Format Section

From the example we can see the individual columns, complete with schema information. If you're familiar with the TDS specification, you can identify the data types and other pertinent fields.

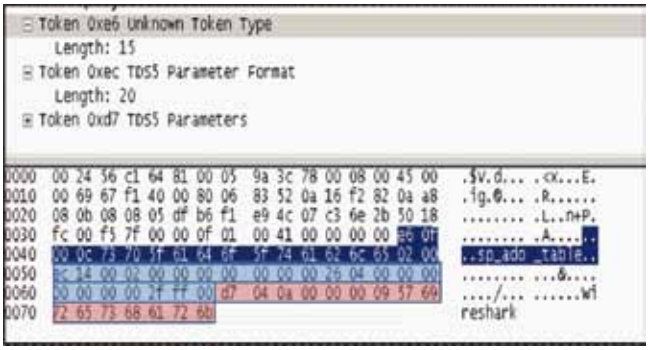

Figure 12: Stored Procedure Breakout

You can apply the same technique to stored procedures. Normally, a stored procedure is sent as TDS_DBRPC, the 0xEC token. Wireshark isn't quite sure on what this is, so it marks the token as *Unknown*. In **Figure 12**, we can examine the breakout: The dark blue is the TDS_DBRPC token, the name of the stored procedure, *sp_ado_table*. In light blue is the parameter format token, in this case defining the data types, integer (0x26) and varchar (0x2f, with size 0xff or 255). Last is the light red, the parameter values: 0x0a, or 10 and "Wireshark".

With a little practice, this sort of breakout will become easier. Of course, using the RIBO utility is easier, but if you can't redirect your client's ASE port number, such as when you're trying to track down some production problem, the sniffer trace is best. It's not as intrusive in terms of hampering performance since the capture occurs at a low level and keeps everything in one space, whereas RIBO creates separate files for each connection.

Conclusion

We've just scratched the surface of the capability of Wireshark and if you're network savvy you can track down other protocols such as HTTP, LDAP, SSL and more. Just select *Edit | Preferences* and look at the protocol list. Don't be intimidated by the complexity. Practice makes perfect and occasional use makes it easier and becomes another valuable item to add to your toolkit. As for the TDS flow, use the TDS specification as a guide to translate the tokens and you'll become the resident expert in analyzing and diagnosing client application and database server situations. Remember, this shark doesn't bite! ■

ASE Cluster Edition Connection Management

By Jeff Garbus

Jeff Garbus has 20 years of expertise in architecture, tuning and administration of Sybase ASE, Oracle, and Microsoft SQL Server databases with an emphasis on assisting clients in migrating from existing systems to pilot and enterprise projects. He has co-authored 15 books and has published dozens of articles on the subject. Jeff and his team can be reached at sales@soaringeagle.biz

When you implement Sybase Adaptive Server Enterprise Cluster Edition, you have access to a single-system presentation of multiple servers. In other words, you point to the cluster, which to the client looks like a single server, and any connection to specific servers happens without user interaction. The logical connection allows ASE CE to redirect the query to another physical server behind the scenes.

This enables the architecture group to perform some pretty cool machinations behind the scenes from a high availability perspective (lookup cluster architecture for more detail). But, it also means that there are some additional possibilities (read: Decisions to make) in your connection / connection string(s). In fact, while you can connect to a cluster, this is not a requirement; you may choose to connect to a specific instance.

This article covers information you need in order to make the right decisions, as well as some implementation tips.

Client Connection Failover

ASE CE ships with Open Client, jConnect, ODBC, OLE DB, and ADO.NET which support Login Redirection, Connection Migration, and High Availability Failover.

Login Redirection

If you do not have the capability of spreading the logins across the clustered servers, you can get single-threading through a single connection at login time. Login redirection is the ability of an instance to redirect the incoming client connection to another instance prior to the handshake. The client has no notification of the redirection. You do not need to do anything special to make this work. This is done automatically by the CE workload manager based on the logical cluster configuration and the cluster's current workload.

Connection Migration

ASE CE has the ability to transfer connections from one server within the cluster to another. Again, this happens transparently to the client. We'll discuss how this happens shortly.

High Availability (HA) Failover

At connection time, in a high-availability configuration, the server provides a list of failover addresses to clients that are paying attention (i.e. are "high availability aware") when they connect. This allows the client to failover multiple times if the instance to which they are connected fails. The client does this with or without a "HAFAILOVER" entry in their interfaces file or directory services, as the list comes from the cluster.

Client Application Changes

Existing applications can connect to CE, however to take advantage of HA features, you may need to change some application code. If you have existing HA applications based upon existing HA functionality, no code changes are required. In fact, the failover lists are automatically updated by CE even if you have already coded a HAFAILOVER entry.

Existing non-HA applications may benefit from some aspects of the HA capabilities of CE with little or no code changes. The only catch is, failover is not transparent. Your application gets an error message when failover is detected, and has to resubmit the batch or transaction to initiate failover. To enable failover for isql, specify the -Q option when connecting to Adaptive Server. For applications linked with Client-Library, set a corresponding connection property (we'll do these shortly) that enables failover.

To make failover transparent to users, check for failover error status, and automatically resubmit the batch or transaction. To do this, you'll have to update the Client-Library version used by the application to use the cluster-related HA capabilities. To enable failover in Client-Library applications, set the CS_HAFAILOVER property at either the context or the connection level using the ct_config or ct_con_props Client-Library API calls.

If you attempt to connect to a failed instance, behavior is the same as with a non-clustered Adaptive Server. The Client-Library tries all the query entries for the instance name in the interfaces file until one of them works, or it has none left to try. You should include query lines in the client-side interfaces file for all instances. Applications can connect to the cluster, which is represented by a series of interfaces file query entries.

When a successful failover occurs, the Client-Library issues a return value named CS_RET_HAFAILOVER, which is specific to several Client-Library API calls during a synchronous connection. In an asynchronous connection (a routine that requires a server response returns CS_PENDING immediately), these APIs issue CS_PENDING, and the callback function returns CS_RET_HAFAILOVER.

"If you have existing HA applications based upon existing HA functionality, no code changes are required. In fact, the failover lists are automatically updated by CE even if you have already coded a HAFAILOVER entry."

Depending on the return code, you can perform the required processing, set up the context, and send the next command to be executed.

You should rebuild your applications using the Open Client SDK with a version that is at least equal to the version of the Open Client SDK shipped with the Cluster Edition.

More On Connection Migration

To implement login redirection and connection migration, make sure the application uses a current copy of the client libraries; or, more specifically, a set of client libraries created for ASE CE. Connection migration occurs when an existing client is transferred from one instance of a cluster to another. For example, it may migrate because the instance it is currently connected to is brought down for maintenance, or it may migrate for load balancing. The migration is transparent to the client application, and allows the cluster to balance the load throughout the logical cluster.

Connection migration allows the workload manager to gracefully move existing connections between instances during administrative failover, failback, or logical clusters going offline. The workload manager can use migration for dynamic load distribution, during which some existing connections are migrated between instances to more evenly distribute load. Connection migration is enabled automatically when an instance uses the Open Client 15.0 client libraries. You do not need to perform any additional configuration for connection migration. This differs from failover because it is planned, rather than incidental. You do not need to code for migration, but you will need to code for failover. The workload manager can initiate a migration when a request can successfully be sent to a client. Specifically, migration can occur:

- To connections that have completed their login, though the migration may occur prior to the initialization of the login;
- After an instance receives a new batch from the client, but before the batch is parsed and executed;
- After an instance completes the processing of a client batch but before it sends a final completion to the client;
- When an instance is not executing any batch on behalf of the client;
- According to the workload manager's algorithm (you'll help define this);
- When a connection's context accommodates migration.

Note: No migration can occur inside a transaction.

At The Operating System (OS) Level

The number of user connections used by CE can't exceed the number of file descriptors available to the CE at the OS level. Note that some file descriptors are used by CE for opening files and devices, so be sure to over configure a bit (i.e. give yourself some room).

CE uses the "KEEPALIVE" option of the TCP/IP protocol to figure out when clients are no longer active. When a connection to a client is inactive for the timeout period, the OS sends KEEPALIVE packets at regular intervals. If the client machine doesn't respond for any of these packets, the OS notifies CE that the client is no longer responding, and it terminates the client's connection. Note that on many platforms, the (configurable!) KEEPALIVE default timeout period is 2 hours (7,200,000 ms).

Configuring Client Connections

There are a variety of new connection properties available to HA to make all of this work as shown in **Table 1**.

Changes To Connection Strings

The below examples are used only at connection time; the workload manager will update the client with the list of available instances that the client can connect to should that list change. This is all transparent to the application itself, it doesn't have to make any decisions about which servers to connect to, or identify which ones are up, which ones are down, or which ones are busy. When a failover occurs, all temporary tables, intermediary result sets, and cursors are lost. The application will need to take this into account when failing over from instance to instance. You need to change the connection to the server as follows.

Interface File

Using the traditional interfaces file, change the "servername" to the name of the cluster. The interfaces file entry would look something like this:

```
# Interface File

  MYCLUSTER
  query tcp ether host1 5000
  query tcp ether host2 5000
  query tcp ether host3 5000
   query tcp ether host4 5000
```

CS_PROP_REDIRECT	Users connect to the cluster itself, and not a specific instance of ASE. Thus, it's not possible to guarantee that a connection will always connect to a specific instance of ASE. Login redirection uses CS_PROP_REDIRECT at the connectivity layer to implement this. This property, which is enabled by default, allows an incoming connection to be re-directed to a specific instance within the cluster, as defined by routing rules and load, no matter which instance they initially connect to. Note that this property is only used at connect (login) time.
CS_DS_RAND_OFFSET	NOT enabled by default, this property randomizes the list of servers that a client will connect to. With this property set, the connections will pick a random server instance to connect to, preventing the overload of a particular server with login requests. While this won't prevent connections coming in should the instance be down, it will minimize the chance that all users of a particular application will be delayed when trying to connect to the server.
CS_PROP_EXTENDED-FAILOVER	This property, enabled by default, allows a connection to failover multiple times within the cluster, as long as there are instances available to service the application (according to routing and failover rules). The number of times that a connection is able to failover is only limited by the number of instances available and the routing rules that you have defined within the cluster.
CS_PROP_MIGRATABLE	This property, enabled by default, determines whether or not a particular connection is eligible for migration to another instance within the cluster, depending on the system load. A connection is eligible for migration if it meets the following criteria: • It is not executing a query batch. • Has no open transactions. • Has no session-level temporary tables. • Has no declared cursors. • Has not changed its password since its initial connection. • Has not run set user or set proxy. • Is not connected to the cluster using SSL. • Is not bound to an engine using the logical process manager. • Is not using real-time data services (RTDS). • Is not a logical connection associated with an inbound site handler.
CS_HAFAILOVER	As mentioned earlier, this property is something that must be coded for in order to insure that, when a failover occurs, the batch that was in-flight gets resubmitted.

Table 1: Connection Properties

Sql.ini

Similarly, on windows platforms, the sql.ini would also change to reflect the differences in the interface file entry. The user no longer connects to an individual server name; in this case they are also connecting to the name of the cluster.

Jconnect / JDBC

If you are using JDBC or jConnect to connect to the cluster, the connection string needs to be modified. Since it's not possible to specify a server alias, multiple host / port combina-

tions must be specified in order to have the user both connect to the cluster and try additional instances if the primary one isn't available. A sample connection string is as follows:

```
URL="jdbc.sybase.Tds:host1:5000,host2:5000/
mydb?REQUEST_HA_SESSION="true"
```

In the above example, "host1 and host2" are the hostnames that the instance(s) of ASE are running on, and 5000 is the port number. Mydb is the database name that should be connected to, and REQUEST_HA_SESSION="true" refers to the ability to failover from instance to instance, should the one we're connected to fail.

ODBC

Using ODBC to connect to the cluster, you can use the ODBC connection manager to specify the appropriate connection to the cluster. To use the connection manager, one needs to simply do the following:

1.) Open the Adaptive Server Enterprise dialog box.

2.) Go to the Connection tab.

3.) Select "Enable High Availability".

4.) Optionally (but strongly advised), enter the alternate servers and ports in the Alternate Servers field using the format: (you may also specify a service name parameter) server1:port1,server2:port2,...,serverN:portN.

In establishing a connection, Sybase's ODBC Driver first tries to connect to the primary host and port defined in the General tab of the Adaptive Server Enterprise dialog box. If ODBC fails to establish a connection, ODBC will go through the list of hosts and ports specified in the Alternate Servers field.

C# / ADO.NET

"HA Aware" applications are coded to handle failover. The connection string is modified to:

```
Data Source='linuxhost03';Port='6000';HASession
=1;Secondary Data Source='linuxhost03';Secondary
Server Port=5000;UID='sa';PWD='';Database='pubs2';
```

Seven Semi-Fearless Predictions About Enterprise Mobility For 2011

Rather more based on empirical evidence than audacious prognostication, Eric Lai gazes into the crystal ball of enterprise mobility to provide a little informed conjecture on the months ahead.

By Eric Lai

New Year Predictions stories have always been one of my favorite genres to read (along with, um, relationship advice columns (but I digress) – and I now I get to write one!

Just So You Know…

When I was a kid, I hated falling off my bike and other stuff that made me look foolish. So I don't pretend to be the most audacious of prognosticators. On the other hand, do expect at least some empirical evidence to back up my mental dart throws. So with that, here goes:

Eric Lai is an ex-Computerworld reporter, now blogging on the evolution of enterprise mobility at Sybase.com and ZDNet.com.

1 35.5 Million Tablets Will Be Sold Globally In 2011

How did I arrive at that ultra-precise number? By simply averaging the public forecasts of six researchers, ranging from a high of 55 million (Gartner) to a low of 22 million from ABI Research. Hip social media types would call that 'crowdsourcing the experts'. Whatever you call it, that's strong growth. (See the chart that follows, based on pre-Thanksgiving forecasts.)

2 Enterprise iPad Deployments Will Accelerate To Become Unnewsworthy

Frustrated by the lack of evidence presented by both sides of the 'Is the iPad enterprise-worthy?' debate, I began collecting a list of companies and schools going big on the iPad last fall.

Check out the list, which I maintain with an educational technology expert, Jim Siegl, here at *http://ipadpilots.k12cloudlearning.com/* As of late November, we had documented 250 organizations that had publicly announced trials and deployments of the iPad. The largest include the Long Island University system, which has given 6,000 iPads to its students, Medtronic, which has deployed more than 4,500 iPads to employees, and our parent company, SAP, which has deployed 2,000 iPads, with its CIO Oliver Bussmann vowing to arm up to 17,000 employees with the Apple tablets. Siegl and I are already planning to start maintaining lists of enterprise deployments of Android tablets and BlackBerry PlayBooks, too. So check us out if you're interested.

3 Nevertheless, Netbooks Will Still Outship Tablets For Many Years

It's amazing how quickly the iPad has rendered netbooks as last year's flavor. The facts are, though, that netbooks still outsell tablets by a factor of 4 or so, according to ABI statistics. "Netbooks continue to grow at a very respectable rate," Orr says. "It's not the same as the last two years, but there's not decline." Even by 2015, when ABI expects 80 million tablets to be sold, he doesn't

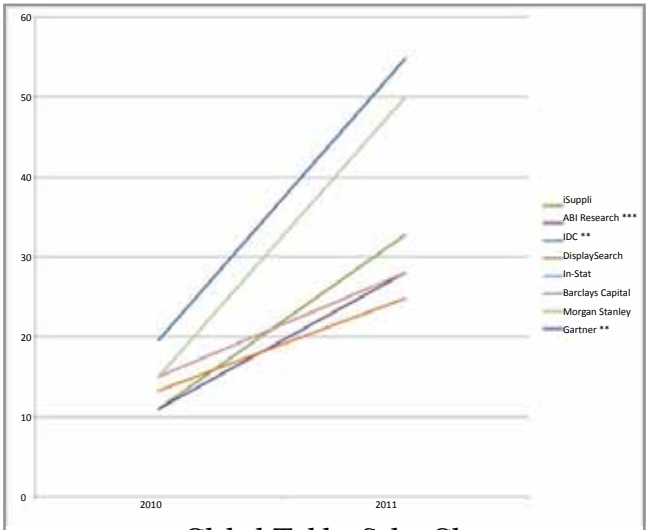

Global Tablet Sales Chart

*'Moderate' forecast, from interview with ABI analyst Jeff Orr, Nov. 12, 2010
** Upwardly revised figures*

see them surpassing netbooks. There's still life in the Wintel platform yet.

4 Tablets Will Start To Deliver Mobile Unified Communication

It's taken many years for UC to become semi-established in the office. The next stage, mobile UC, may arrive much more quickly, courtesy of tablets such as Cisco's Cius or the Avaya Flare. Docked on the desktop, they will be able to enhance the normal office desk phone with features like video conferencing, real-time presence and instant messaging.

When taken on-the-go, they will provide more convenient and powerful collaboration capabilities than available today. "We'll soon be able to talk to colleagues while editing the same document – something impractical on smaller smartphones," predicts longtime mobile industry analyst Philippe Winthrop, in an article in Sybase's coming 2011 Guide to Enterprise Mobility (available in print and e-book – look for it on Sybase.com in early February).

5 The Cius And The Flare Will Still Be Shoved To The Side In The Enterprise By The iPad And Android

Pre-iPad, these tablets would've felt like revelations compared to the dowdy and pricey Windows tablets available at the time. Now we're fully in the iPad/Android/Bring-Your-Own-Device era, the Cius and Flare look like the kind of gear that only an IT manager could really love. ABI Research analyst Jeff Orr agrees, adding that it's even a "bit of a distraction" to lump them with the iPad or PlayBook since it will be such an unfair fight.

6 Google Will Need To Rein In Android, Unless It Creates A Vista-like Bomb

Innovation, openness and free – sounds great in theory, but can be a pain in practice. For mobile developers used to Steve Jobs', how best to put it, iron fist-in-velvet glove, Android initially felt like a breath of fresh air. That is, until the never-ending dessert-themed updates combined with the buffet of different hardware, each with their own idiosyncrasies, started leaving them pining for Apple's autocratic ways.

I think Google will have to do more – much more – to rein in developers and partners. A vibrant ecosystem is one thing – a chaotic one is another. Microsoft's failure to force developers to properly rewrite applications and drivers for Windows Vista caused numerous blue screens and misfiring peripherals, and was one of the reasons for that operating systems' massive failure. Google needs to tone down its talk about innovation and openness and find a middle way that works.

7 The Enterprise App Market Will Start Its Shift From Packaged To Custom - And Begin To Explode

For mobile developers thus far, it's all been about casual-games-this or social-media-app-that. Something like Super Twario, a Twitter-reader/80s-era-scrolling-game, was the unholy but predestined union of the two genres. There's been a serious dearth of serious business apps. The market just wasn't there. With Android smartphones and iPads infiltrating the enterprise and Bring Your Own Device policies becoming the standard, the market now IS there. That means the big ISVs, including SAP/Sybase, will not only get serious about bringing out mobile extensions to their big business apps, they'll even start to bundle them into their core license. That will lower the entry barrier to enterprises, who have thus far had to build their own since there was nothing on the shelf for them.

While much of the development pie will be taken away from small developers, it also means that the remaining slices – customization for industry verticals, for instance, will grow substantially in size. Or to mangle a third metaphor – the rising mobile tide should lift all developer boats. ■

Rock your development skills with PowerBuilder® 12.

You don't have to choose between being a .NET developer, a WebForm developer or a WPF developer. Using the skills you already have today, you can leverage .NET inside PowerBuilder 12 to deploy a wide range of applications and easily satisfy your customers' demands. In fact, PowerBuilder's new functionality and features give you the freedom to be the programmer you've always known you could be – writing exactly the applications you want, and doing it faster and easier than your .NET peers.

Only PowerBuilder 12 offers powerful productivity, the ability to deliver rich user experiences and unparalleled ease of migration. So you can evolve your applications and focus on core functionality, saving significant person hours and development costs along the way.

SYBASE®

PowerBuilder® 12

GET SET
TO ROCK .NET

PowerBuilder® 12. The hottest thing yet to rock .NET

When PowerBuilder was first released, it shifted the development paradigm with speed, simplicity and power. The latest version of this award-winning development tool takes everything you love about PowerBuilder and builds it for the future, providing more power and .NET all the way.

Here's another way to look at it: if the first PowerBuilder release was a race car, PowerBuilder 12 is like a private jet. It has never been easier or faster to get where you need to be, ahead of your competition and in true rock star style.

Blazing speed. Smokin' features. Let the performance begin.

Five lines of code. That's all PowerBuilder 12 needs to access powerful business logic using the patented DataWindow, which is completely rewritten to support Microsoft's WPF technology. Built on top of the Visual Studio Shell, PowerBuilder 12 combines productivity, flexibility and security with the .NET capabilities you need today.

An incredibly fast and cost-effective migration path evolves your Win32 applications to .NET. There's no need to toss out your PowerBuilder code to get your application to .NET – PowerBuilder 12 gets you there, while protecting your investment.

Unleash your inner programming rock star. Download your FREE trial of PowerBuilder 12 by visiting sybase.com/powerbuilder today.

SYBASE®

Using Data Access Application Blocks With Sybase ADO.NET Provider

By Wayne Smith

In this article, I will introduce you to the Data Access Application Block (DAAB) and explain how to use it with respect to the Sybase ADO.NET Provider. Additionally, we'll examine how to reference the DAAB in a Visual C# .NET Project.

MS Patterns And Practices

The Patterns and Practices Team at Microsoft have provided a collection of patterns and practices as well as code blocks that we can use in our applications. These patterns and practices are based on proven architectures, guaranteed production quality code, and known lifecycle best practices. Some of these include:

a.) DAAB for .NET
b.) Caching DAAB for .NET
c.) DAAB Exceptions for .NET
d.) DAAB Logging for .NET

To learn more visit: *http://www.microsoft.com/resources/practices/code.mspx*

What Is A DAAB?

The Microsoft DAAB for .NET consists of a single .NET-based assembly, which contains all of the functionality necessary to perform the most common data access tasks against a Sybase ASE database. Specifically, the DAAB helps you:

- Call stored procedures or SQL text commands.
- Specify parameter details.
- Return *AseDataReader*, *DataSet*, *XmlReader* objects, or single values.
- Use strongly typed table and field names.
- Support parameter caching.
- Allow additional tables to be added by passing in pre-existing datasets.
- Update datasets with user-specified update commands.
- Create *AseCommand* objects.
- Allow strongly typed data rows to be passed in place of parameters.

Where Can I Get DAAB?

You can download the "Data Access Application Block" from the Microsoft website: *http:/www.microsoft.com/downloads/details.aspx?FamilyId=F63D1F0A-9877-4A7B-88EC-0426B48DF275&displaylang=en* and *http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnbda/html/daab-rm.asp*.

For more information on the Sybase ADO.Net Provider 2.0, visit: *http://infocenter.sybase.com/help/topic/com.sybase.infocenter.dc20066.0215/pdf/adonet2155.pdf [Chapter 5 Supported Microsoft ADO.NET features]*

There are newer versions than DAAB 4.1. I have tested the 5.0 DAAB and there is no real difference in usage so you can use the newer version with the Sybase ADO.Net Provider 2.0. To keep everything from getting confusing, we will use the 4.1

Wayne has worked at Sybase for over 13 years. He is a Principle Technical Support Engineer supporting various Sybase Connectivity and Middleware products. He can be reached at wayne.smith@sybase.com.

in this article. There have been several updates to the Sybase drivers that impact developers in a positive way. First, starting with the 15.5 Sybase 2.0 Provider, support is available for the Unity framework as well as LINQ. The second part is with 15.5ESD#7 which now supports integration into Visual Studio 2005 and 2008. Visual Studio 2010 integration should be available by the end of 2010. This will make developing in Visual Studio much easier for Sybase developers.

How Do I Use The DAAB?

By default, the DAAB DLLs are all signed by Microsoft. This causes an issue for all vendors that want to access the DAAB. All vendors such as Sybase need to use their own signed respective DLL. So in order to incorporate the Sybase Provider, we need to add some references and code into the DAAB and build the DLL in combination with the ADO.NET Provider. After installing, you'll need to add the Microsoft Enterprise Library DAAB for Adaptive Server. The Adaptive Server ADO.NET Data Provider 2.155 extends the Enterprise Library 4.1 DAAB to support Adaptive Server.

The DAAB is a collection of classes that simplifies common database functions such as creating database instances and updating database records. The DAAB also encapsulates database specific features, which allows for a database agnostic application design. The DAAB for Adaptive Server classes are supported in Microsoft .NET 2.0 and up to 4.1. Here are the steps for using the DAAB:

Step 1: Step 1: Install the DAAB and system requirements. Download and install the EntLib41Src source. Keep in mind it has to be the source for the Enterprise Library. You need Microsoft Visual Studio 2008. To use DAAB for Adaptive Server, you need to update the ADO.NET on your developer's system to ADO.NET 3.5 and the 4.0.

Step 2: Add Sybase to the reference of DAAB. CD to the directory where you installed the EntLib41Src (for example: C:\EntLib41Src\Blocks\Data). Once in the Data directory, start the project Data.sln. Next, find the DbProviderMapping.cs source code. After the line: *public const string DefaultOracleProviderName = "System.Data.OracleClient"*, add our ADO .NET Provider as follows:

```
/// <summary>
/// Default name for the Sybase managed provider.
/// </summary>
public const string DefaultSybaseProviderName =
"Sybase.Data.AseClient";
```

Next, find the DatabaseConfigurationView.cs file and after the line: *using Microsoft.Practices.EnterpriseLibrary.Data.Sql;* we need to add the following using statements:

```
using Microsoft.Practices.EnterpriseLibrary.Data.
Sybase;
using Sybase.Data.AseClient;
```

Additionally, after the line: *private static readonly DbProviderMapping defaultGenericMapping = new DbProviderMapping(DbProviderMapping.DefaultGenericProviderName,typeof(GenericDatabase));* we need to add:

```
private static readonly DbProviderMapping de-
faultSybaseMapping = new DbProviderMapping(
DbProviderMapping.DefaultSybaseProviderName,
typeof(AseDatabase));
```

Then, after the line: *(OracleClientFactory.Instance == providerFactory) return defaultOracleMapping;* we must add the following:

```
(AseClientFactory.Instance == providerFactory)
return defaultSybaseMapping;
```

Lastly, add the following assembly references to the Enterprise Library DataAccessApplicationBlock\Data solution. Note that this will be the only time you will need to reference the EnterpriseLibrary.AseClient.dll.

- AdoNet2.AseClient.dll
- EnterpriseLibrary.AseClient.dll

At this point, we need to perform a complete rebuild of the Enterprise Library. Once the DLLs are built, we will be able to add them to our project.

Step 3: Simply use the DAAB methods for the ASE calls.

Referencing The DAAB In A Visual C# .NET Project

When you create a data access client that uses the DAAB, you should perform the following steps:

Step 1: Set a reference to the DAAB assembly in your application:

- Microsoft.Practices.EnterpriseLibrary.Data
- Microsoft.Practices.EnterpriseLibrary.Common

Note: There is no need to add the Sybase.Data.AseClient as it will be pulled from the DAAB.

Step 2: Add a using statement at the top of your data access client source file(s) to reference the Microsoft Application Blocks Data namespace. All DAAB types are located within this namespace.

Make sure the DLLs that you reference are being pulled from C:\EntLib41Src\Blocks\Data otherwise the framework will give you errors about a mismatch if you happen to have the signed DAAB blocks installed on the system. I'm using Visual Studio 2008. You need to start by creating a C# form application project. Once you have created your project we can now start coding. On the form, we want to add a Data Grid. Add a pulldown menu called File, then add an Exit label as shown in **Figure 1**. Double click on *Exit* and add the following code:

Figure 1: Exit Label

```
this.Close();
```

Figure 2: Count and Load

Next, add a pulldown menu called Authors, add a Count label, then add a Load to the pulldown menu as shown in **Figure 2**. Double click on Count and add the following code :

```
Database db = null;
db = DatabaseFactory.CreateDatabase("Sybase");
int count = (int)db.ExecuteScalar( CommandType.
Text,"select count(*) from authors");
string message = string.Format("There are {0} au-
thors in the database",count.ToString());
MessageBox.Show(message);
```

Double click on Load and add the following code:

```
Database db = null;
db = DatabaseFactory.CreateDatabase("Sybase");
DataSet ds = db.ExecuteDataSet( CommandType.
Text,"select * from authors");
dataGrid1.DataSource = ds.Tables[0];
```

We also need to add two using statements:

```
using Microsoft.Practices.EnterpriseLibrary.Data;
using Microsoft.Practices.EnterpriseLibrary.Common;
```

Once you're finished, your Form should look like **Figure 3**. Double click on the Data Grid to bring up the source code. Your code should look like this:

```
using System;
using System.Drawing;
using System.Collections;
using System.ComponentModel;
using System.Windows.Forms;
using System.Data;
using Microsoft.Practices.EnterpriseLibrary.Data;
using Microsoft.Practices.EnterpriseLibrary.Common;

namespace CustomerManagement {
  /// <summary>
  /// Summary description for MainForm.
  /// </summary>
  public class MainForm : System.Windows.Forms.
  Form {
     # region Uninteresting Windows Stuff

     // The  first item in the Authors pull down menu
     private void mnuCount_Click(object sender,
     System.EventArgs e) {
        Database db = null;
        db = DatabaseFactory.CreateDatabase(
              "Sybase");
        int count = (int)db.ExecuteScalar(
              CommandType.Text,
              "select count(*) from authors");
        String message = string.Format(
              "There are {0} authors in the database"
               count.ToString());
        MessageBox.Show(message);
     }
     private void mnuLoad_Click(object sender,
     System.EventArgs e) {
        Database db = null;
        db = DatabaseFactory.CreateDatabase (
              "Sybase");
        DataSet ds = db.ExecuteDataSet(
              CommandType.Text,
              "select * from authors");
        dataGrid1.DataSource = ds.Tables[0];
     }
     private void dataGrid1_Navigate(object sender,
     NavigateEventArgs ne) { }
  }
}
```

Very simple right? Oh WAIT! Something is missing? Where is the connection information? In the DAAB, the connections aren't stored as a connection string in the application though you could do it that way. The idea is to create

Figure 3: Final Form

a connection string and store it in a App.Config file so that all of the application can access it. In that way, when the application is distributed, it will access only one source file. This method is important for a number of reasons. For example, when building applications for the cloud or directory services your connection information is easily accessed and separate from the application itself. If you need to change the Server Address, all you would have to do is edit the App.Config file to make the change. There is no need to rebuild the application.

Let's take a look at how we create an App.config file. There are several methods that can be used: XML editor, Visual Studio and the Enterprise Configuration Utility. Remember, when we downloaded the source code for EntLib41Src, there was the source code for the Enterprise Configuration Utility as well. We can use that to create our connection to ASE.

In Visual Studio 2008, go to your project and right click on Add | New Item as in **Figure 4**. The new item file will be Application Configuration File, the name will be App.Config.

Figure 4: Add New Item

Now that we have our App.Config file, we can add our connection information. Again, as I have indicated before, you have options on how to edit the file. In Visual Studio, the App.Config is just an XML file, nothing special. This is the view from Visual Studio. As you can see, the file has no connection information to start out with.

```
<?xml version="1.0" encoding="utf-8" ?>
<configuration>
</configuration>
```

Here is a App.Config that is configured to connect to my ASE:

```
<?xml version="1.0" encoding="utf-8"?>
<configuration>
  <configSections>
    <section name="dataConfiguration"
type="Microsoft.Practices.EnterpriseLibrary.Data.
Configuration.DatabaseSettings, Microsoft.Prac-
tices.EnterpriseLibrary.Data, Version=4.1.0.0,
Culture=neutral, PublicKeyToken=neutral" />
  </configSections>
  <connectionStrings>
    <add name="QuickStarts Instance" connect
ionString="Server=wsmithxp3;Port=5000;User
ID=sa;Password=;Database=pubs2"
      providerName="Sybase.Data.AseClient" />
  </connectionStrings>
</configuration>
```

There is a tricky part that can get you into trouble. On the Visual Studio 2005 and 2008 Menu, you may see the Enterprise Library Configuration. The location of the executable should match the location of the DAAB dll that you have configured and compiled: C:\EntLib41Src\bin\EntLibConfig.exe. The best way to do that is to right click on App.Config and select "Open With". Then select "Add", then browse to the location where you have the EntLib41Src stored and go in to the bin folder and select EntLibConfig.exe. as shown in **Figure 5**.

Figure 5: Selecting the Path to the Enterprise Manager

Figure 6: Open Dialog

Select "OK" and make sure you highlight the line and set it as a default as shown in **Figure 6**. Once you're done, you have access to everything from the EntLib41Src.

One warning: If you do use the wrong EntLibConfig.exe , the connection object in

the App.Config will be signed. Because the original DAAB DLLs are the ones we are using and are not signed, you will get a mismatch error about using the wrong Enterprise Library.

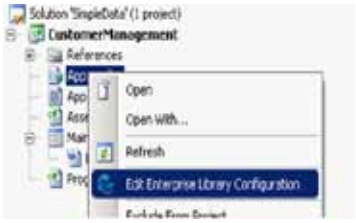

Figure 7: Edit Enterprise Library Configuration

Now, if you have the 4.1 or 5.0 installed, it is important to build the Enterprise Library Configuration utility yourself **(Figure 7)** and make your own App.Config. Make sure at run time that Visual Studio is pointing to the C:\EntLib41Src\bin\EntLibConfig.exe. The main reason is for version control. You really want to make sure you're using 4.1 for everything. If you want to use 5.0 that is fine, just make sure you don't mix versions of the framework. The nice part about the App.Config is that you can have as many connections as you want. In this case, we have one called Sybase. You could add another one called SAP and have the application make a connection for each command to two different servers. Here's a nice example:

```
<?xml version="1.0" encoding="utf-8"?>
<configuration>
 <configSections>
  <section name="dataConfiguration" type="Microsoft.Prac-
tices.EnterpriseLibrary.Data.Configuration.DatabaseSettings,
Microsoft.Practices.EnterpriseLibrary.Data, Version=4.1.0.0,
Culture=neutral, PublicKeyToken=null" />
 </configSections>
 <connectionStrings>
  <add name="SAP" connectionString="Server=NUTMEG;Po
rt=5909;User ID=sa;Password=sybase;Database=pubs2;"
   providerName="Sybase.Data.AseClient" />
  <add name="WSMITH-WIN7" connectionString="Dat
abase=Database;Server=(local)\SQLEXPRESS;Integrated
Security=SSPI"
   providerName="System.Data.SqlClient" />
  <add name="Sybase" connectionString="Server=wsmith
xp3;Port=5000;User ID=test;Password=testing;Database=p
ubs2"
   providerName="Sybase.Data.AseClient" />
 </connectionStrings>
</configuration>
```

As you can see, you can even mix vendors providers in the App.Config. In our code *db = DatabaseFactory.CreateDatabase("Sybase")*; we are calling the App.Config and sending our server name via the DatabaseFactory class which is very nice and very simple. The next piece, that is a little tricky, is building the application. By default, the Visual Studio will default to the Systems path pointing to Microsoft Signed DLL's as in **Figure 8**.

Figure 8: Path Pointing to Microsoft Signed DLL's

We need to change that to point to the Enterprise Library that we built which isn't signed. On the pulldown menu, select EntLib41Src. It will prompt you about thechange; just click "OK".

Now that we have our code and connection information configured, let's run our sample application and see how we did. The Count is shown in **Figure 9** and the Result Set is shown in **Figure 10**.

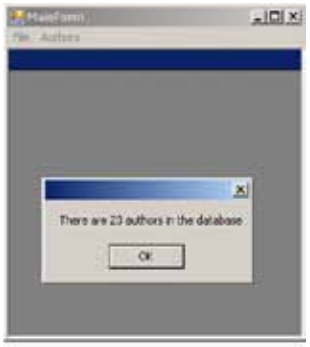

Figure 9: Selecting Count from the Menu

Figure 10: The Results Set from the Load Menu

Conclusion

Using the Sybase ADO.Net Provider works very well with the business libraries. Now that the Sybase Provider is integrated into Visual Studio, it will now make development that much easier. The key to all of this is to make sure the environment is setup correctly: don't mix signed and unsigned DLLs. If you need signed DLLs, sign them through Visual Studio. The DAAB does not have a long history but it is having an impact in coding circles and has Sybase recognized the benefits of using the DAAB with the Sybase ADO.Net Provider for ASE. That's it for now. Happy coding! ■

LearnIT Step Into The Online Classroom For PowerBuilder

Increase your value with web-based training courses for Sybase professionals everywhere...

Presented by one of the most well known PowerBuilder educators in the world - Mr Yakov Werde - eLearnIT's PowerBuilder Core Skills training course enables software developers at all levels to rapidly develop a comprehensive PowerBuilder skill set.

Each module is presented in an intuitive format that is sure to make the learning experience fun, compelling and rewarding.

Working closely with ISUG, eLearnIT has produced a clearly presented product that is tuned to PowerBuilder's most recent 11.5 release.

This course is focused on PowerBuilder deloyments in real world scenarios. It will give you the skills to build your own PowerBuilder workgroup scale application from scratch, or perform maintenance tasks on an enterprise scale application. In short, it will increase your value in the workplace, so visit the URL below today!

Modules covered:

- IDE Skills and Concepts
- PowerScript Fundamentals
- GUI Basics
- Object-Oriented Programming
- DB Connectivity and Embedded SQL
- DataWindow Technology
- Deploying An Application

Essential PowerBuilder Series

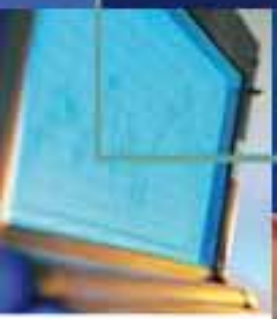

>> **For More Information, visit:**
www.isug.com/pbtraining

eLearnIT
Helping People Know and Do

How To Take Full Advantage Of Deferred Compilation In ASE 15.x

By Chris McCartney and Dave Putz

Chris McCartney joined Sybase in 1990 and has worked on every version of SQL Server/Adaptive Server since SQL Server 4.0.1. Chris has developed and presented extensive internal training for Sybase Engineering and Support on ASE. Chris can be reached at cmccartn@sybase.com

Dave Putz has been working with computers for 36 years, doing application programming, networking, operating systems support, hardware support, and, for the last 11 years, has worked in support for Sybase. His primary focus is ASE; but as part of Worldwide Product Support Engineering he deals with issues realted to many areas.

Adaptive Server Enterprise version 15.0.2 introduced a new and powerful feature called Deferred Compilation (DC). The purpose of this document is to help Sybase DBAs and Developers better understand:

1.) Why Deferred Compilation was needed.
2.) How Deferred Compilation works.
3.) How to enable / disable Deferred Compilation.
4.) How to determine if Deferred Compilation was used.
5.) What queries are and are not supported for Deferred Compilation.
6.) How to write ASE-based applications with Deferred Compilation in mind.
7.) How to handle ASE upgrades with Deferred Compilation.

Why Deferred Compilation?

Before discussing what the Deferred Compilation (DC) feature does, it is important to understand the reason why it was needed. There are two categories of coding structures that traditionally had problems prior to the DC feature. They are:

- **Category 1** – queries within stored procedures that use local variables in the where clause, and
- **Category 2** – queries against temporary tables that were created in the stored procedure.

Category 1: Local Variables

Consider the following stored procedure, as processed prior to ASE 15.0.2:

```
create proc myProc1  @p1 int,@
p2 int as
  select col_x from table_1 where
  col_1 = @p1 and col_2 = @p2
```

For the purpose of this example, assume that:

- The table being queried, ***table_1***, has two indices, as shown below:
  - Non-Clustered on (col_1) named ***idx1***.
  - Non-Clustered on (col_2) named ***idx2***.
- The DBA has done the correct thing and run the following command not too long ago:

```
update index statistics table_1
```

As a result of the update statistics command, the table has the following Optimizer Statistics:

- Histogram on col_1, and
- Histogram on col_2

Given that stored procedures are optimized without knowing the value

of any local variables, the Optimizer, in this case, will be asked to determine the best access method **without** knowing the values of either of the two Search Arguments (SARGs) in the query. Because the values of ***@p1*** and ***@p2*** are not known until execution time, it is not possible for the Optimizer to know what they are when it optimizes the query. As a result, the Optimizer has to guess how many rows will match and therefore, whether to do an index scan on ***idx1***, an index scan on ***idx2***, or a complete table scan.

Category 2: Temporary Tables Created In The Procedure

Consider the following stored procedure, as processed prior to ASE 15.0.2:

```
create proc myProc2 as
  select id, type, crdate into #temp from sysobjects
  select m.id from #temp t, mytab m
  where m.id = t.id and t.type = 'P'
```

In this case, the Optimizer is tasked with determining the optimal query plan, **BUT** it does not know at optimization time how many rows will be in the ***#temp*** table, nor how large this table will be. Therefore, it simply estimates that the ***#temp*** table is 10 data pages in size and that 100 rows will qualify. This estimate will rarely result in the ideal query plan.

Clearly, there are many different stored procedure coding structures that will fall into this issue. For this reason, the Deferred Compilation feature was implemented.

How Does Deferred Compilation Work?

Now that we understand the problem, let's discuss how the DC feature solves the problem. To do this, let's examine the following two examples:

```
-- Example 1: Category 1- Local Variable Case

create proc myProc1 @p1 int, @p2 int, declare @count
int as
  select @count=count(*) from table_1 where col_x > 0
  select col_x from table_1 where col_1 = @p1 and
  col_2 = @p2
```

When the above procedure is first run with the DC feature enabled, the Optimizer will:

1. Completely optimize the first query in the procedure since all the SARGs are known.
2. Skip the second query which was flagged for deferred compilation during the normalization phase when the procedure was created.
3. Copy the actual text of the query into the query plan for later optimization.

Then, as it executes, the ASE Execution Engine will first execute the initial portion of the Query Plan until it gets to the portion marked as a DC query. At this point, it will send both the query text **and** the values of the previously unknown SARG to the Optimizer for processing. The Optimizer will use the SARG value(s) along with the existing Optimizer Statistics to determine the optimal path, which will immediately be executed. All subsequent executions of this procedure or at least all of those that find this Query Plan available in the Procedure Cache, will execute the Query Plan as is. In other words, with Deferred Compilation, **the SARG values in the very first execution are the ones on which the query plan is based.**

Now let's examine the stored procedure for Category 2:

```
-- Example 2: Category 2 - #Temp Table Case

create proc myProc2 as
    select * into #temp from table_1 where col_1 = 0
    select m.col_x from myTab m, #temp t
    where m.col_1 = t.col_1 and m.col_2 = t.col_2
```

When the above procedure is first run with the DC feature enabled, the Optimizer will:

1. Completely optimize the first query in the procedure since all the SARGs are known.
2. Skip the second query which was flagged for deferred compilation during the normalization phase when the procedure was created.
3. Copy the actual text of the query into the query plan for later optimization.

Then, as it executes, the ASE Execution Engine will first execute the initial portion of the Query Plan until it gets to the portion marked as a DC query. At this point, the Optimizer will know both the size of the ***#temp*** table and the estimated number of qualifying rows. Therefore, the Optimizer will be able to select the most optimal Query Plan.

If, however, the Query Plans for either of these stored procedures get aged out of the Procedure Cache – as a result of the Procedure Cache becoming full – or nearly full, then the next time the procedure is executed will be analogous to the first execution. This situation could result in getting a different Query Plan.

How to Enable / Disable Deferred Compilation

The Deferred Compilation feature was introduced in ASE 15.0.2. In this initial release, the feature was enabled by default and could only be disabled by turning on trace flag ***7730***.

As of ASE 15.0.3 ESD #1, however, the above trace flag is no longer needed. As of this release, DBAs can turn this feature off and on via the new ***sp_configure*** parameter ***procedure deferred compilation*** – which is set to 1 (on) by default. This parameter is ASE wide and there is no session level equivalent. Enabling this parameter does not require any additional ASE memory nor is an ASE restart needed to change it. If you are running on an ASE release between 15.0.2 **and** 15.0.3 ESD #1 and you do not want to have DC enabled, then you must use the ***7730*** trace flag. You may do so either by including this trace flag in your ***RUNSERVER*** file via the ***–T*** parameter, or by running the following command each time your ASE is restarted:

```
dbcc traceon(7730)
```

Note also that, because the DC feature is first applied when the stored procedure is created, it is at creation time that ASE evaluates the value of both the ***procedure deferred compilation sp_configure*** parameter and the ***7730*** trace flag – not at stored procedure execution time!

Determining If Deferred Compilation Was Used

All messages that ASE uses to communicate that DC was used are included in ***showplan*** output. The first time a stored procedure with a qualifying DC query it is executed, ***showplan*** will include the following messages:

To be Optimized at runtime using Deferred Compilation.

Optimized at runtime using Deferred Compilation.

On subsequent procedure executions, *showplan* includes the following message in its output:

Optimized using Deferred Compilation.

What Queries Are And Are Not Supported?

It is very important to note that, as of the date this article was published, not all query types are supported by the DC feature! More specifically, the query types that are **not supported** by DC are:

- Queries using extended data types.
- XML queries.
- CIS pass-through queries.
- Internal SQL queries.
- Non-user connection queries , like the threshold procedure or upgrade.
- Execute immediate queries .
- Conditional Queries ("if..." "if exists...", "if not exists").
- Select into queries.
- Cursor queries .
- All queries if *compatibility* mode is on at procedure creation time.

For example, the following procedures are **not supported** by DC and, therefore, have a higher chance of not using the optimal query plan:

```
-- Example 3

create proc myProc3  @p1 int as
  if exists (select 1 from table_1 where col_1 = @p1)
  begin
     print "Above query is not supported by DC!"
  end
```

"if exists" queries are not supported

```
-- Example 4

create proc myProc4  @p1 int as
  if (select count(*) from table_1 where col_1=@p1)=0
  begin
     print "Above query is not supported by DC!"
  end
```

"if" queries are not supported

```
-- Example 5

create proc myProc5  @p1 int as
  select * into #temp from table_1 where col_1= @p1
  print "Above query is not supported by DC!"
```

Select into queries are not supported

```
-- Example 6

create proc myProc6 declare @cmd varchar(200) as
  set @cmd="select col_1 from table_1 wherecol_2=0"
  exec ( @cmd )
```

Execute immediate queries are not supported

Note that the following example shows that only the queries that are part of the conditional statement are not supported. Queries **inside are supported by DC**.

```
-- Example 7

create proc myProc7  @p1 int as
  if(select count(*)from table_1 where col_1=@p1) > 0
       select col_2 from table_1 where col_1 = @p2

  else
     select col_3 from table_1 where col_1 = @p2
```

This query is supported by DC

BUT, each or these queries are supported by DC

Note that these examples do not represent an exhaustive list.

How To Write ASE-based Applications Using Deferred Compilation

By now, the concepts provided in this section should be pretty straightforward. In other words, where it matters, do your best not to use the query types that are not supported by DC when either Category 1 or 2 from above apply. One example of alternative coding structure which **is supported by DC** is shown below:

```
-- Example 8: Alternative Coding Structure to Replace
-- the 'if exists .." Case

create proc dc1  @p1 int as
  declare @count int
  select @count = count(*) from table_1 where col_1 =
  @p1
  if ( @count <> 0 )
  begin
     print "Above 'select' query IS supported by DC!"
  end
```

What About ASE Upgrades

The only reason to consider ASE upgrades in this discussion is that, as mentioned above, stored procedures created with the DC feature disabled – including those created on an ASE release prior to the existence of the DC feature will be processed as if the DC feature had been purposely disabled by the user.

"Note that, if the DC feature is enabled before a stored procedure is executed for the very first time following an upgrade, the automatic recompilation that occurs due to the upgrade will use DC for the stored procedure."

Note that, if the DC feature is enabled before a stored procedure is executed for the very first time following an upgrade, the automatic recompilation that occurs due to the upgrade will use DC for the stored procedure.

Otherwise, ASE users who have recently migrated their ASE to a post ASE 15.0.3 ESD #1 release from a release prior to 15.0.2, should consider performing the following steps to ensure that the DC feature will be used where applicable:

Step 1: Verify that the DC feature is enabled via the following command:

```
1>   sp_configure 'procedure deferred'
2>   go
```

This should indicate a Run Value of 1, indicating that the feature is enabled. If the Run Value = 0, then run the following command to enable it:

```
1>   sp_configure 'procedure deferred', 1
2>   go
```

Step 2: Perform ONE of the following 3 actions:

- Run ***sp_recompile*** on all tables that are accessed by any user created stored procedures, OR
- Execute all user created stored procedures one time only using the ***"with recompile"*** option, OR
- Drop and recreate all user created stored procedures

Conclusion

Clearly, the Deferred Compilation feature is a powerful and positive new feature for ASE. It is important, however, to fully understand some of the ins and outs of the feature – as presented in this paper, to ensure that it is used to its fullest potential. ■

ISUG's goal is to provide the most valuable information possible to our users, which is best achieved by connecting the vendors and the users. The Partner Program provides a vehicle for vendors and service providers to participate in user group events and communicate to the Sybase user group community. Amongst the other significant benefits, our Partners are able to send product information to our world-wide membership, plus have the ability to conduct user surveys or polls providing invaluable feedback that can't be obtained by any other means.

Each year our members renew their memberships, so our contact lists remain current. The same is true for our Partners since we want to assure that we are providing the most current product information and introduce those who are setting the trends. Many Partners offer their own benefits directly to the Sybase users, including discounts on software, consulting, training, conferences, and seminars. Experts who work for our Partners often share their extensive Sybase expertise by participating in ISUG chat groups, forums, and Special Interest Groups (SIGs).

APPEON

Appeon is the industry leader in providing PowerBuilder-based development products and IT outsourcing services to global customers.

Principal Products

Appeon for PowerBuilder – "the most productive platform" to deploy PowerBuilder applications to J2EE and .NET on the web with over 1,200 companies in 60 countries worldwide choosing Appeon for PowerBuilder for web conversion and development of their mission critical web applications.

Appeon for PowerBuilder deploys existing PowerBuilder client/server applications to the web, while maintaining the look, feel and functionality of the original application. No major rearchitecting is required. Customers can interact with their applications in the way to which they're accustomed and businesses leverage their existing software and developer resources.

Appeon Professional Services – takes a partnership approach with its clients, forging long-lasting relationships to not only develop but also maintain its software applications, thereby empowering its customers to better focus on improving and expanding their core business.

Corporate Contact

Email: sales@appeon.com
Phone: 1-877-3APPEON
Website: www.appeon.com

BRADMARK TECHNOLOGIES

Bradmark Technologies is a global database solutions provider that develops, markets, sells and supports tools that provide anytime — anywhere data access across all the major RDBMS installations and beyond. For over 20 years, Bradmark has helped Fortune 500 companies manage their data, increase system reliability, and drive down IT costs.

Principal Products

Surveillance™ for Sybase – Surveillance is the first and only performance monitoring solution that supports all the major Sybase environments (Sybase ASE, Replication Server, Sybase IQ and Sybase Mirror Activator) running on UNIX, Linux and Windows operating systems.

Surveillance is a comprehensive, cross-functional toolset that satisfies all of your data management requirements by detecting and preventing problems before they adversely affect productivity, ensuring that IT systems are up and operating at peak performance, and is the preferred choice for Sybase monitoring.

Corporate Contact

4265 San Felipe, Suite 700
Houston, TX 77027
Phone: 1-800-621-2808
Fax: 1-713-621-1639
E-Mail: sales@bradmark.com
Website: www.bradmark.com

ELEARNIT

eLearnIT is an independent source for cost effective PowerBuilder developer training using rich media eLearning solutions.

eLearnIT Courses Combine:

- Content based on the knowledge of industry-recognized subject matter experts, organized and structured according to the principles of adult learning.
- Presenters who are both enthusiastic and engaging to help keep learners engaged.
- Exercises that are thoughtfully designed to support self-directed discovery oriented learning.
- Courses that are packaged into a modern rich media presentation format that is delivered via a globally 24 x 7 available web site.

Principal Products

Essential PowerBuilder Series – a 7 unit self-paced rich media learning experience providing PowerBuilder core skills training. *See a preview of the PowerBuilder core skills course at http://www.elearnitonline.com/PBCourseOverview/player.html*

MDA Tables Training – a 5 unit self-paced rich media experience providing comprehensive performance and tuning training. *See a preview of the MDA Tables course at http://www.elearnitonline.com/MDACourseOverview/player.html*

Corporate Contact

Email: PB@eLearnITonline.com
Phone: 1-973-896-7241
Website: http://www.eLearnITonline.com

eLearnIT
Helping People Know and Do

EMBARCADERO

Embarcadero's family of database tools make it easier for you to design, build, manage, monitor, and tune all your Sybase ASE databases. Since 1993, Embarcadero has been the trusted tooling provider to thousands of Sybase customers. In fact, in a recent survey of International Sybase User Group (ISUG) members, Embarcadero was selected as their number one database tooling provider.

Discover why experienced Sybase data management professionals place their confidence in the depth of Embarcadero's Sybase-specific capabilities to get their jobs done quickly.

Principal Products

DBArtisan – the industry's premier database administration tool.

DB Optimizer – prevent, find, and fix performance issues via visual SQL profiling and tuning.

Rapid SQL – full-featured SQL IDE helps you write quality SQL faster.

DB Change Manager – simplify, automate, and report on database changes.

ER/Studio – multi-level data modeling for logical and physical database design.

While Embarcadero's products are known for providing deep, native support for Sybase ASE, they also provide in-depth support for Oracle, Microsoft SQL Server, DB2, MySQL (and more). With Embarcadero's new XE line of tools, support for all major database platforms is included in a single, affordable license, so you can extend your Sybase skills into less familiar databases while enabling your entire organization to standardize on a common, powerful solution.

Corporate Contact:

100 California Street, 12th Floor
San Francisco, CA 94111
1 415 834 3131 x232
Email: elias.terman@embarcadero.com
Corporate Site: www.embarcadero.com
Sybase Tools: http://www.embarcadero.com/solutions/sybase

embarcadero

Kickstart Your Career With A Membership In Your Nearest Sybase User Group

BASIC PACKAGE

Free Copy of Appeon for PowerBuilder Developer Edition, 20% off Werysoft's QweryBuilder, Discounts on O'Reilly Products and Training, Networking Opportunities and Member Access to the ISUG Website:

- Access to user communities and Special Interest Groups (SIGs)
- Technical Documents Repository
- Announcements of upcoming events
- Ability to vote in the enhancement process to communicate to Sybase what users want for the future of Sybase products
- Online access to selected *ISUG Technical Journal* content

GOLD / GREEN PACKAGE

Benefits of the Basic package, plus…

Monthly issues of the *ISUG Technical Journal*, including articles such as:

- Discovering Real-Time, Performance Monitoring and Event Management
- Encrypted Columns in ASE 15
- PowerBuilder 11.5 Top Feature Picks
- Understanding & Managing Optimizer Statistics in ASE

Copies of current versions of software with personal use licenses of:

- DataWindow.NET
- EAServer Developer Edition
- PocketBuilder
- PowerBuilder Enterprise (with InfoMaker)
- SQL Anywhere Studio
- Workspace Studio

Note: Benefits requiring a physical shipment, including the quarterly print version of the ISUG Technical Journal, are only available with the Gold package.

Magazine Discounts

- 40% off on a subscription to the *PowerBuilder Developer's Journal*

Documentation

- Adaptive Server Enterprise System Tables Diagram fold-out poster
- Adaptive Server Enterprise Technical Library CD
- PowerBuilder Technical Library CD
- Replication Server System Tables Diagram

Software and Training Discounts

- 25% Discount on PowerDesigner Developer
- 10% off list price on all individual registrations for Sybase Learning Center and SyberLearning LIVE classes
- 20% on CSP Certification Exams (DBA & IFD) Sylvan Prometric
- 20% on a computer-based training from Branick Consulting

All benefits are subject to change. Please refer to www.isug.com for the current list.

Join today at my.isug.com/join

Corporate Memberships

Workplace Green
$495 for 6 members
+ $99 for additional members

Enterprise Green
$895 for 12 members
+ $99 for additional members

Workplace Gold
Starting at $595* for 6 members
+ $119* for additional members

Enterprise Gold
Starting at $995* for 12 members
+ $119* for additional members

* Price varies by region

For more information:
www.isug.com

For online sign-up:
my.isug.com/join

*According to a leading **global hedge fund firm**, prior to utilizing Surveillance, they experienced x# of unscheduled outages a week. Since implementation, they reported **ZERO** outages!*

Surveillance *DB*™

Proactive Management Solutions for Sybase

Sybase ASE | Sybase ASE For Cluster Edition | Sybase IQ | Sybase Replication Server | Sybase Mirror Activator

The Recommended Choice for Reliable System Availability

Surveillance *DB*™ for Sybase is the only **end-to-end** solution for database management that provides essential real-time monitoring, issue identification and resolution, rules-based event management, data repository and reporting, enabling IT professional to ensure reliable system quality, productivity and performance.

A strategic alliance Sybase partner, Bradmark is the recommended choice for monitoring and managing Sybase ASE, Replication Server, Sybase IQ, and Mirror Activator investments that run on UNIX, Linux and Windows operating systems. And with support for Oracle, SQL Server and DB2-UDB databases as well, Surveillance provides a single, integrated source for DBAs to manage multi-database, multi-OS system enterprises from a single source.

Surveillance delivers tangible business value:

- **Reduce time to identify and resolve database, application or OS issues**
- **Mitigate risk of unscheduled downtime avoiding lost business costs**
- **Determine database usage by department for internal cost allocation**
- **Deploy existing labor and resources more efficiently**
- **Provide a structured process for monitoring and managing databases under scrutiny creating an enterprise standard**
- **Improve capacity planning**

*Go to **www.bradmark.com/sybase** today and discover Bradmark's proactive approach to database management —*

Or, contact Bradmark directly for a free Surveillance product evaluation:

For Bradmark U.S.: **(713) 621-2808**, or toll-free **(800) 621-2808**

For Bradmark EMEA: **+31 (0) 251-268 248**

bradmark technologies

The Choice Provider for Sybase Management

BOARD OF DIRECTORS

President
Bryan Enochs
AT&T
Orlando, FL
Phone: 407-826-6171
bryan@isug.com

Secretary
Martin Clark
Enable Development LLP
martin@isug.com

Vice President, Business Relations
Cynthia Gill
CSG Technology, Inc.
Houston, TX
cynthia@isug.com

Vice President, Finance
Eric Van Patten
ING
Des Moines, IA
eric@isug.com

Vice President, Membership
Anne Walker
St. George Group Treasury and Capital Markets
Sydney, Australia
anne@isug.com

Benefits Director
Sheri Page
Accumen Industries, LLC
Cheverly, MD
sheri@isug.com

Communications Director
Michael Bendorf
First Pacific Corporation
Salem, OR
michael.bendorf@isug.com

User Group Director, European
Jeremy Chubb
Maidenhead, UK
jeremy@isug.com

User Group Director, Asia-Pacific
Joe Woodhouse
Melbourne, Australia
joe@isug.com

User Group Director, North America
Cindy Bean
Bradmark Technologies, Inc.
Houston, TX
cindy@isug.com

User Group Director, South America
Jean-Pierre Dareys
Mexico City, Mexico
jeanpierre@isug.com

Enhancements Director
Aryeh Silverberg
Hennepin County
Minneapolis, MN
aryeh@isug.com

Partner Director
Marcelle Von Wendland
London, UK
marcelle@isug.com

TeamSybase Liaison
Terry Voth
Toronto, ON
Canada
terry@isug.com

Sybase Liaison
Jonathan Baker
Sybase, Inc.
Bethesda, MD
jonathan.baker@sybase.com

ISUG Executive Director
Mike Harrold
ISUG
Columbus, OH
mike@isug.com

ISUG Editor
Adrian Bridgwater
ISUG
London, England
adrian@isug.com

ISUG WEBSITE LINKS

User Group Directory
Asia Pacific -- http://www.isug.com/ap
Europe, the Middle East and Africa – http://www.isug.com/emea
North America – http://www.isug.com/na
South/Latin America – http://www.isug.com/sla
Virtual User Groups – http://www.isug.com/sigs

Membership Information
Joining ISUG – http://www.isug.com/join
Benefits – http://www.isug.com/membership_benefits
Academic Memberships – http://www.isug.com/academic_memberships
Corporate Memberships – http://www.isug.com/corporate_memberships

Resources
e-Journal – http://www.isug.com/e-journal
Enhancement Process – http://www.isug.com/enhancements
Newsletter – http://www.isug.com/current_newsletter
Partner Program – http://www.isug.com/partner_program
Sybase FAQ – http://www.isug.com/Sybase_FAQ

*Be sure to visit us at **www.isug.com** for all the very latest information!*

Step Into The Online MDA Tables Training Classroom For ASE DBAs

Anywhere on the Globe, Anytime of Day or Night, Cost Effective, Comprehensive Performance and Tuning Training for ASE 15

MDA Tables Training

Master ASE Instructor and DBA, Brian LeFevre, guides you through acquiring the knowledge and skills you need to install the MDA Tables and effectively use the information they contain to help you identify: system choke points, production tables killing performance due to table scans, overworked devices, under-used caches or users that hinder performance due to the stored procedures they run.

Armed with this valuable information, you can rapidly take remedial action to speed up production response times and reduce down time.

In five in-depth modules you will cover these performance and tuning topics: Installing and configuring MDA Tables, Access Methods, Process Tuning, Server Tuning, and Database Tuning.

eLearnIT courseware; designed and delivered by professionals, for busy professionals. eLearnIT/ISUG courses are comprehensive, up to date, rich media, web delivered, self-paced and available for unlimited play, 24/7 via subscription to its learning portal. Courses include text, audio descriptions, video demonstrations, lab guides and lab installation scripts as well as learner support from the course instructor.

Ready to view the course overview and experience a sample topic?
Visit www.elearnitonline.com/MDACourseOverview/player.html

eLearnIT
Helping People Know and Do

http://www.eLearnITonline.com