

Sybase ASE Query Optimization

Embarcadero Technologies, Inc.

November 2010

Americas Headquarters
100 California Street, 12th Floor
San Francisco, California 94111

EMEA Headquarters
York House
18 York Road
Maidenhead, Berkshire
SL6 1SF, United Kingdom

Asia-Pacific Headquarters
L7. 313 La Trobe Street
Melbourne VIC 3000
Australia

TABLE OF CONTENTS

Table of Contents	- 1 -
Introduction.....	- 2 -
Q&A with Peter Dobler, Kevin Sherlock, Scott Walz and Harry Helal.....	- 3 -
Migration Issues	- 3 -
Performance Problems	- 4 -
Performance Solutions	- 6 -
Query Plans	- 7 -
Differences Between ASE Versions.....	- 7 -
Balancing Workloads	- 8 -
Avoiding Process Blocks.....	- 9 -
DBA Tools.....	- 9 -
Optimization Goals.....	- 11 -
Forcing Indexes and Plans.....	- 11 -
Abstract Query Plan	- 13 -
Determining Queue Size	- 13 -
DBA 1.0 v. DBA 2.0	- 14 -
Biographies	- 16 -
About Us	- 17 -
About Embarcadero	- 17 -
About ISUG	- 17 -
Additional Resources	- 18 -



INTRODUCTION

In October 2010, Embarcadero Technologies and the International Sybase User Group (ISUG) hosted a Webcast on the topic of Sybase ASE Query Optimization. As part of the program, the panel captured questions from registrants in advance and during the presentation, and shared valuable insights.

Among the panelists were Scott Walz, Senior Director of Product Management at Embarcadero; Harry Helal, Program Manager at Embarcadero; Peter Dobler, Founding Partner at Dobler Consulting; and Kevin Sherlock, Senior DBA at Children's Hospital and Medical Center.

What follows are highlights from the Webcast—in the form of a Q&A with all the panel participants—as well as supporting references designed to provide a complete set of resources to help DBAs learn more about Sybase ASE Query Optimization.

Q&A WITH PETER DOBLER, KEVIN SHERLOCK, SCOTT WALZ AND HARRY HELAL

In this Q&A from the Webcast, Dobler, Sherlock, Walz and Helal discuss a variety of topics including migration issues, performance problems, DBA tools, query plans and optimization goals.

ON MIGRATION ISSUES

Question: What are some of the issues involved with migrating from 12.5 to 15.X; What are some things to look out for?

Dobler: What I've seen with migration is mostly issues with query performance, especially when there were some workarounds involved where indexes were forced, forced plans were put in, or simply where queries that used to run quite efficiently pre-15 just simply do not perform anymore. There are several reasons for that. Number one is that Query Optimizer changed. The changes that came with it might surprise a few people. As an example, you now have three different optimization goals you can choose from, which introduces new join mechanisms. Before you only had one. What this does is it forces people into application changes. One thing that causes a lot of grief in the migration is the change of the cursor behavior. With the change of the cursor behavior, instead of using the work table to step through the cursor results it goes now directly back to the base table, and that causes a lot of problems. You need to know how to work with them.

Sherlock: Peter hit it on the head. There are a few more join methods and access methods now on the new optimizer. Before 15, a Loop join was generally the only join method. It was so much easier before 15 to reach a level of stability in terms of Optimizer predictability, without having to spend so much time in development writing more efficient SQL, or bothering with column level statistics and histogram data and density information. I think what's not really in the manuals is the fact that the default optimization goal now really leans heavily on more information being provided to the core processing engine than Optimizer, in order to make the correct query plan decision.

For more information about migration issues, including a description of “Compatibility Mode”, read this Sybase white paper:

http://www.sybase.com/files/White_Papers/Sybase-ASE-Compatibility-Mode-072309-WP.pdf

And the following white paper from Sybase describes changes that might need to be made to existing SQL code when migrating from ASE 12.5 to ASE 15 to ensure the SQL code continues to function correctly:

http://www.sybase.com/files/White_Papers/SY-Required-SQL-Changes-for-ASE15-v.1-073009-WP.pdf

For more on the issue of migration to ASE 15.x:

Mandatory Code Changes for Migration: <http://www.sybase.com/detail?id=1063534>

(http://www.sybase.com/files/White_Papers/SY-Required-SQL-Changes-for-ASE15-v.1-073009-WP.pdf)

Sybase ASE 15 Best Practices: Query Processing & Optimization:

<http://www.sybase.com/detail?id=1056243> (http://www.sybase.com/files/White_Papers/ASE15-Optimizer-Best-Practices-v1-051209-wp.pdf)

Planning your Migration to ASE 15: Tips & Tricks: <http://www.sybase.com/detail?id=1063027>

Using Compatibility Mode In ASE 15.0.3 ESD#1 For ASE 15 Migration:

<http://www.sybase.com/detail?id=1063556>

Optimizer Stability Changes in Adaptive Server 15.0.3 and 15.5: <http://www.sybase.com/detail?id=1080354>

ON PERFORMANCE PROBLEMS

Question: What is the best way to analyze a performance problem?

Dobler: You need to isolate the problem, to figure out where the problem is. The best case scenario is it happens right when you're on the system and you're able to in real time access the system tables. You're able to look at the logs as it happens, and you see all kinds of things you can do with the system tables, to look at what's happening right then. But again, you need to isolate it. If you can't do that, then you go back and try to figure out in the logs what's happening at a certain time. In either case, once you isolate it, either the SQL or an entire transaction—you need to be able to reproduce it. Lastly, the one really important thing is you need to take preventative action. Do what you can do, either from what Sybase supplies with the tools (monitor, server, MDA tables) or you go to a third-party vendor. But you need to track what's going on in the server and create your own repository so you can go back in time and know exactly if somebody says, “Hey, last night at 3 a.m. my job didn't finish on time. What happened?” You need to have a method to access this.

Sherlock: To me it's all about the toolbox you have, the toolsets you have available to you, whether it be products from Embarcadero or Sybase-provided tools. I count eight or nine different tools you can use. And of course, timing is important, as Peter mentioned. You have to institute some sort of methodology about using some of these tools to capture the necessary information. In the old days I used to use Historical Server and Monitor Server quite a bit. Those were useful tools to start collecting and aggregating data so that you can go back historically and look at long-running procedures and certain other bottlenecks in your system. I think the MDA tables should probably be heavily used, and there's sysquerymetrics, which is a feature in 15, which is good for analyzing statements. There are session-tracing capabilities in 15. One of my favorites is SP Monitor. SP Monitor has undergone quite a number of changes, and really good changes, in recent ASE releases. If you haven't used SP Monitor before I'd really encourage folks to look into that. You can really get a good dashboard from that procedure, with measuring response times. And then there's other tools [like] Sysprocesses, just your regular catalog tables that you can query to a temporary table and collect historical data there, and use that to isolate statements and store procedures that are taking a long time to run. Usually this comes down to isolating a single store procedure or statement, and as Peter mentioned, the secret is being able to reproduce that. Not very well documented is the fact that reproducing problems that you find can be challenging, and we're trading the flexibility of a really smart Optimizer these days for ease of reproduction. I would encourage everybody to [acquire a free version ASE], where you can import statistics from your production environment and start reproducing query plans without impacting your production servers.

For more on performance analysis, see:

ASE TechTalk Podcasts (Specifically "Monitoring Practices" and "Tips and Tricks")

<http://www.sybase.com/asetechtalk>

<http://www.sybase.com/products/databasemanagement/adaptiveserverenterprise?htab=DISCOVER&vtab=Webcasts/Podcasts>

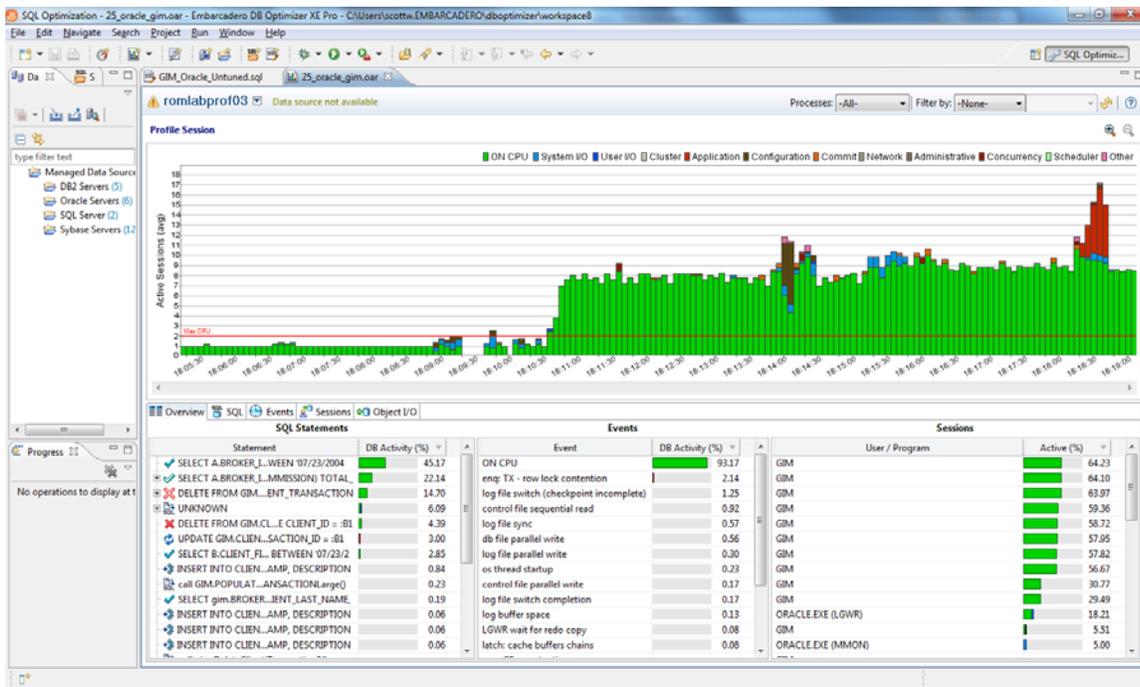
Techniques to Identify Poorly Performing SQL in Production – Techwave 2010 Paper:

<http://download.sybase.com/presentation/TW2010-Presentations/DATASE9.pdf>

ON PERFORMANCE SOLUTIONS

Question: That sounds like a lot of work. What kinds of tools from Embarcadero can help with performance issues?

Walz: We have a tool, DB Optimizer (DBO) that's able to identify problem areas [for example, in a SQL profiling session] really quickly and very graphically, so it really jumps out. The tool lets you slice and dice, hover over, and actually pinpoint areas that you want to look at. If you're interested in this spike or that spike, you can look at that. We can look at the actual SQL statements that are being run. We can look at the sessions that are taking up the most time. We can go into other areas of the product.



Helal: We actually collect information from Sysprocessors and a number of MDA tables. All this information is actually stored in a file that you can share and store for future analysis. It basically allows you to capture the SQL statements and sessions that have taken the most percentage of your processing power, and it also breaks down what it was doing, whether it's waiting for I/O or waiting for the network. And that basically gives you the ability to have SQL that you can start tuning. DBO will go through a number of steps. The first step is to look at SQL itself and see if there are any obvious issues. It will look at the index structures behind the tables, and it will put a collection of different plans that it will propose [to resolve any issues].

ON QUERY PLANS

Question: How do I interpret query plans in ASE 15 and why do query plans change from one run to the other?

Sherlock: Query plans change and are generated mostly because stored procedures or statements are being compiled, and you have a cost-based optimizer. In theory, if query plans are changing you'd like to think that 99% of the time they're changing for the better. The fact that you might have more compilation events happening could be a result of your procedure cache not being large enough, or you're not using statement cache. There's also something called deferred compilation in ASE 15. The age-old problem in ASE was after you created a temp table and a stored procedure and then you referenced that temp table and stored procedure and join it to a user table, you would like the ASE Optimizer to take advantage of any statistics on the temp table before it compiles a plan for that statement. Deferred compilation is an instance where that can be happening to you in 15. So, there are reasons for [changes]; they shouldn't be bad reasons. If you're finding that query times are changing and they're changing for the worse, you'd use these tools that we mentioned before and find out why that is. And there's probably a viable solution to resolve that issue. I think the idea here is customers have been asking for a smarter Optimizer, an Optimizer that is more flexible, and you might find that you get more regeneration of query plans or recompilations happening because of that, so you trade one for the other.

Dobler: When parameters change and the query plans are changing, it's probably most of the time for the better—the Optimizer found a better way to do [something]. Kevin mentioned something really interesting. Besides the statistics there are more and more important [factors] to be looking at. What is important is the procedure cache. With the change in the Query Optimizer in 15 is the query optimizer now uses more procedure cache, and you need to really take a look at that. I believe there was a new feature introduced where you can actually limit the size query plans to use the procedure cache, so that the entire procedure cache will not be filled up. Sometimes you really need all this procedure cache and you can turn this feature off. Think about it this way, the Query Optimizer now tries more and more combinations to find the best plan for the execution. That's really something to remember.

ON DIFFERENCES BETWEEN ASE VERSIONS

Question: What are some of the differences between ASE 15.0.3, ASE 15.5, and ASE 15.0 Cluster Edition?

Dobler: With 15.5 over 15.0.3, the major introduction is the in-memory database. Without the in-memory database there would be no real-time loading, because the in-memory database plays a pivotal role in the real-time loading environment. And then we have the Cluster Edition, which is available in 15.0.3 and 15.5. The official statement from Sybase is use Cluster Edition for high availability. It gives you access from several server nodes into the same master, same procedures, same databases. So it really lets you grow your application horizontally, and is more of a shared-everything cluster, compared to a shared-nothing cluster.

The white paper below has information on what's new with the ASE 15 Optimizer and Query Processing Engine, choosing the best optimization goal for your application, parallel query processing in ASE 15 and performance monitoring tips.

http://www.sybase.com/files/White_Papers/ASE15-Optimizer-Best-Practices-v1-051209-wp.pdf

ON BALANCING WORKLOADS

Question: How do I balance the optimization of mixed workloads; for example, mostly SELECTs and small updates/deletes during the day, with nightly batch processing involving massive deletes and inserts that are taking place at night?

Sherlock: I think that's the target audience for 15 especially, mixed workload. Keep in mind, as we progress over the last few releases to be a more dynamic server configuration capable server, there's lots of server-level buttons and switches to pull that don't require you to actually reboot ASE. In the context of this question, obviously it depends on what kinds of workloads we're talking about. Optimization goals can be changed at a session level or with log-in triggers or at the server level. Those are all available to make certain adjustments to what you think your workload's going to be. The usual dynamic configuration capabilities of ASE are always there: changing cache sizes; creating and dropping named caches; changing configuration of named caches; maybe moving a washmarker to lower in the chain so that you're not spending a whole lot of time processing dirty pages to disk; changing your cache partitions, configurations, optimization time-out limits, your statement cache settings. I think ASE is fairly well geared toward accommodating [mixed workloads]. Of course, it depends on what workloads we're talking about. The person who asked the question mentions massive deletes and inserts. Well, maybe it's more efficient to drop some indexes that aren't critical to that workload during that batch processing window, and recreate them giving those index recreations a higher priority with Resource Governor when that window is closing. Obviously testing is important in these kinds of scenarios.

Dobler: I want to talk about something that is a licensed feature of Sybase. What this allows you to do is really create great partitioning on your tables. Think about the concept of having local and global indexes. Local index means that your index is local to the partition, so you can have a hundred-million table that you divide up into a hundred partitions, and each partition has its own index range. So you can imagine that this goes much faster with inserts, and then you have global index on top of it.

For more on partitions in ASE 15.0, see this technical white paper from Sybase:

http://m.sybase.com/files/White_Papers/ASE-Partitions-040307-wp.pdf

This briefing by Dobler also addresses partitions:

http://www.sybase.com/files/Feature_Articles/Sybase_ASE15_SemanticPartitions_article.pdf

ON AVOIDING PROCESS BLOCKS

Question: What are some best practices to avoid process blocks?

Dobler: From my experience, design. It all comes down to A, shorten the transaction, and B, sequence your transaction in a way that you're able to keep the logs on the table as short as possible. And of course there's always the situation [to consider]. In a transaction with hundreds or even thousands of transactions in flight, one thing you can think about is using clustered index on non-key columns to create many lost pages for the inserts, so you're not fighting over one page at one time. That definitely is something to think about.

Sherlock: My ears perked up when I heard creating a clustered index on a non-key column. The only thing I'd add to that is I'd be careful with that too. It's a decent solution, but make sure you include enough columns in the key of that clustered index to make it unique. And with data-row locked tables as well make sure that you're careful with the space management settings for data-row locked tables, especially with regard to reserved row space and anticipated row space.

Question: Are clustered indexes useful in DOL tables with regard to ASE 15?

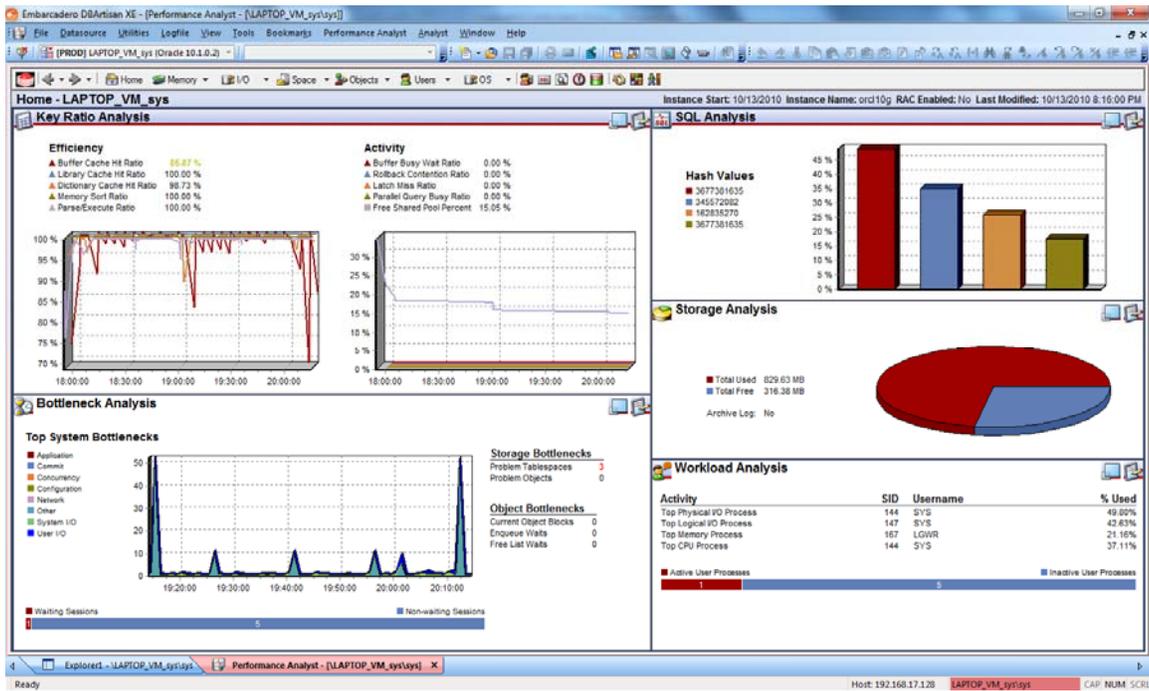
Sherlock: Absolutely, they act just like clustered indexes on APL tables as far as I'm concerned. Now they're called placement indexes because the basic table structure is essentially ordered by the key. Now it's not guaranteed to be ordered that way because of the way the placement index works. But you don't have a heap table like you do if it were a non-clustered index. There's no heap on the data page side that's just being added to at the end. It's still attempted to be kept in clustered order, not guaranteed but still attempted. So they act just like they do in APL tables and certainly are perfectly viable. [As for] clustered indexes on data-only locked tables, there's nothing unusual about them in comparison to the APL side.

ON DBA TOOLS

Question: When I hear process blocks, I think database administration, and when I hear database administration, I think DBArtisan. How might some of these things be handled in a tool like DBArtisan?

Helal: We do have some features in Artisan that helps simplify that process for you. One of the features is called the Database Monitor. I have it running here, and essentially it's a real-time feature that gives you a list of processes and related information that you can get out of system tables, and you can control things such as the refresh rate, whether you want to see the SQL or not. And essentially in this screen you'll have to see the processes themselves. If you have that SQL toggle on you'll be able to see the running SQL. You can see some stats coming out of sysprocesses about how much memory this session has and what's the current command that it's running. If it's green than that means it's currently running. This SQL Server I have is only a single engine so you'll only get one green. Under the All Locks area it basically goes out to syslocks, and retrieves the list of the locks. Again, if you have many locks you'll see them here, so you'll know which sys they're related to, you'll know if they are blocking or not. You'll see red coded if these locks are blocking any other sessions. Under the blocking logs you'll actually get the list of all the blocking logs. We don't have any in this server, which is not doing anything. Once you have lots of sessions, you're able to see the SQL and

you're able to go and run the infamous kill command. You can also pull that SQL and do a query plan on it, or you can take it into your SQL Editor and see what it does.



Earlier on we were talking about collecting information from SP sysmon. I just want to take the opportunity to talk about another feature in Artisan called Performance Analyst. This is basically an add-on analyst that runs SP sysmon and queries some MDA tables, collects information periodically. You can set up frequency and a number of other aspects. You can store that information for later. You could also put a list and thresholds on it. Instead of writing your own scripts and collecting that information, you could have Artisan do that for you. We have a server version of it called Performance Center that does the same thing except it's a 24x7 with notification alerts and action, so you can actually have the server take an action for you.

ON OPTIMIZATION GOALS

Question: How do I go about determining the right optimization goal for my environment? What are the differences among optimizer levels and what level of testing is required?

Dobler: With ASE 15 there are essentially three optimizer goals that you can set from a pre-set environment. There are all rows (OLTP?), all rows mixed and all rows DSF. What they mean is they influence the way that joins are being treated. The all rows OLTP tends to narrow the joins more towards the nested loop join where you have the large table looking at the smaller table for each value. The all rows mixed setting is more a pre-15 type of setting, and the all rows DSF tends to change the joins more towards a merge or a hash join. How do you determine which one is the best? First, always run your test against a cold server. You want to take out any skewing from caching or other memory enhancements. You can do this several ways. For single SQLs you turn off the statement cache or set cache size to zero. Recompile stored procedures before every run. It's actually running your statistics time into a file, with different optimizer goals, and you simply grab your output for nested loop join, merge join or has join, and with the timing you can pretty much figure out which one works best. If you know that there's a lot of nested loop joins and if these queries work the fastest, than your optimization goal would be all rows OLTP.

Sherlock: This is a very well-documented, very well-addressed issue in several white papers. Peter mentioned the different optimization goals. It's not only just about join methods, it's about parallel query processing. Parallel query processing is treated a little bit differently in 15. One thing I might suggest that's not necessarily mentioned in any manual anywhere is that if you were a parallel query processing shop before going to 15, try turning that off while you're exploring the different optimizer goals in 15. You may find that some of the other join methods and support for Star Schema and other methodologies that are in that query optimizer eliminate the need for resource-intensive requirements for parallel processing. QPTune is a tool that Sybase I think has released with 15.0.3. The main objective of QPTune is to automatically put your workload through various scenarios with different optimization goals and give you an idea of where your best fit would be. I'd recommend looking into QPTune if at all possible.

ON FORCING INDEXES AND PLANS

Question: Why do I need to force indexes in my query? Why doesn't the Optimizer pick it up on its own from the index?

Dobler: Probably most of the time the reason it doesn't pick up the index by itself is outdated statistics, or the query optimizer determined that using the index is not worth the effort and a full table scan would be better served at this time. Not using the index is not always a bad thing.

Sherlock: I think shops that are forcing indexes in their development are probably running into some other issues that are forcing them to take that path. Very common issues for why an optimizer won't choose a particular index have to do with data-type mismatches. With 15, the whole concept of data mismatches is a little more relaxed than it was pre-15, but they're still there and it's still possible. Index design is very important in terms of your most critical queries, choosing the right indexes. I've seen shops where they forced use of an index, but it's not key positioning on that index. Indexes are only really valuable if they're being key positioned on. In other words, if they're being used like you use a phone book. If you turn right to the Sh page for Sherlock, that's key positioning. If you're just using index from A to Z to find Sherlock, your using the index, but you're not using it efficiently. If you were forcing indexes before, you'll probably be forcing them in 15. But it's probably much more worth your time and efforts to find out what else underneath the covers is causing the Optimizer to choose an alternate access path.

Question: Should I force the plan or let the Optimizer pick the plan?

Dobler: My choice always is to let the Query Optimizer pick your plan, because once you start forcing the plan, you're locking yourself into a situation where you may not be able to recover from this if you're going to the next upgrade. So forcing the plan for me is the absolute last method. If everything else fails, then you do the forced plan. But forced plan as a way of development, I stay away from that. With a forced plan you essentially take a query plan that works in a certain environment and lock this query plan in. If this environment changes, you are locked into a forced plan that doesn't represent reality anymore.

Sherlock: I think situations where shops have to force plans have to do more with somebody being put in a position that says make what we have work, we're not going to change it. Maybe it's a vendor, a written piece of code that they don't control. But obviously if you're in a development environment and you're creating procedures and designing indexes and so forth you'll let the Optimized do the hard work. I think forced plans and those types of scenarios are last resort.

Question: I have many instances of CPU usage going to 100% because a query was using the wrong index. Doesn't ASE become aware of this situation and do something about it?

Sherlock: No, ASE itself isn't designed to "self heal" in any way like this. If you think about it, the ASE query processing engine is doing only what you asked it to do. If you think it is using a "wrong" index, then you need to determine why that is. Most likely, it's because there isn't enough information provided to the optimizer to determine the most efficient query plan. In that case, there isn't anything ASE could do about it other than guess at other access paths. Something I probably wouldn't want as a "feature". Having said this, there is nothing to keep you from using the diagnostic information ASE provides you via MDA tables to detect this condition of 100% CPU utilization by a single task, and act accordingly either manually, or by some automated process that you can easily develop. The ASE kernel's purpose is to serve query tasks, and it leaves the monitoring and process management up to the DBA or other tools which serve that purpose better.

In most cases, bad query plan decisions can be traced to a definitive cause which typically point to a lack of necessary and accurate table/index/column level statistics, inefficient SQL design, inefficient database design.

If you have a specific query that is giving you problems, I would invite you to post a very specific question (providing ASE version, actual SQL code, table/index DDL, and optdiag statistics information

for each table) to the newsgroups at <http://www.sybase.com/newsgroups> where myself and many others can help you pinpoint where there might be something you can do to avoid this.

ON ABSTRACT QUERY PLAN

Question: The Abstract Query Plan feature is a great feature, but I've never had any success implementing it. For me, it's very cumbersome to use. Any advice?

Sherlock: I think it is a difficult feature to use. It's very obtuse language. When you're at this level of having to use abstract plans and so forth, this is last resort. I haven't actually used abstract plans a whole lot. The cases where I would use abstract plans are in reproduction, trying to reproduce an issue in production, and I can't seem to get the Optimizer to behave the way it was behaving when the problem was occurring. I might force a plan that way through abstract plans so that I can start seeing how the query behaves.

Question: Would you recommend and ASE 15 to run a cursor from the table that gets updated, or copy the data into a temp table and run the cursor from there?

Dobler: There is an option you can turn on called Semi-Sensitive that institutes the old, pre-15 behavior, which then means you can update the cursor. But it is a code change. You have to add it at the cursor open statement. The reason why cursors were changing goes back to the Query Optimizer changes. I wouldn't necessarily create a work table, just use the option that you can add with Semi Sensitive or Sensitive.

ON DETERMINING QUEUE SIZE

Question: If we are using MDA tables, since they depend on queue size how do we arrive at the best queue size?

Sherlock: By "queue size", I'll assume you are referring to the "pipe max messages" settings that apply to some of the MDA tables. These are a small subset of the MDA tables that have data which will "age out" of the table depending on the configuration settings for the "pipe max messages" for these tables. If you are collecting data for these MDA tables, you need to capture these rows before they "age out".

deadlock pipe max messages - monDeadLock: How many deadlock messages do you want to retain?

errorlog pipe max messages - monErrorLog: How many lines of the current errorlog do you need access to?

plan text pipe max messages - monSysPlanText: How many historical query plans do you need to be able to capture?

sql text pipe max messages - monSysSQLText: How many historical SQL queries do you need to capture?

statement pipe max messages – monSysStatement: How many historical Statements do you need to capture?

Of the 5 “pipe” tables, the last 3 are the most common for performance analysis. These are historical versions of their “monProcess%” counterparts. You only really need to activate these pipe tables when collecting historical information for these events. Also, you shouldn’t just leave these pipes active all of the time since they do consume ASE resources to collect. So, given your “sp_configure” settings for “pipe max messages”, ie:

“plan text pipe max messages”
“sql text pipe max messages”
“statement pipe max messages”

You need to collect the contents of these tables and store them in a repository somewhere to analyze later. How do you know if the settings are big enough? When you select from any of the mon% tables above, you should record how many records were retrieved from the tables. If the number of records is close to, or equal to your “pipe max messages” setting, then you either have to reduce the time between sampling/collecting the data, or increase your “pipe max messages” settings.

ON DBA 1.0 v. DBA 2.0

Question: You hear the term DBA 1.0 versus a DBA 2.0, where the DBA 1.0 is in there at the command line writing his own scripts, and the DBA 2.0 is using more modern tools. Which one are you?

Sherlock: I’m a DBA .17. This summer I painted my house. My wife wanted me to hire somebody [to do it], and I said no I want to paint the house myself because I want to have the control. I wanted to but the materials. I was a homeowner 1.0 this summer. Next time I need to paint my house I’m going to be a homeowner 2.0. It’s all about what you’re willing to trade, time and convenience, knowledge. If you’re using tools and GUI products, such as those from Embarcadero, you’re leveraging someone else’s expertise, you’re saving resources and time on your end. Folks who are scripters, you’re calling the DBA 1.0s, those are probably people who’ve been around the command line a lot. I will say if you’re using the tools, it’s very important to know what that tool is doing underneath the covers. Because if you move on to a different position at another company and they may not have a tool that you’re used to using, you’re going to have to know how to go back on the command line and perform your job.

Dobler: I'm probably a 1.5, somewhere in between. There are situations where I need to have full control in command line and I don't want to be on a PC connected through a client-server connection to a server and execute something. For whatever reason the network goes down than I have to clean up the mess after that. That's when I like to have command line on the server where I know it's still executed even when I lose connectivity. Then there are the time savings, as Kevin said, where it just doesn't make a whole lot of sense to use command line when you can do it with a few clicks. But again, you need to know what happens on the tool when you click a button. If you don't have an understanding of the consequences, what's happening when you click your way through, then you go on to another position, as Kevin said, with another company where they don't have this tool, what do you do then? So you need to understand what's underneath. Use the tools to your advantage but know what they're doing.

Walz: I'd be in trouble if I didn't say 2.0. But I feel I'm at my best when I know what the tool's doing, and that's one of the great things about DBArtisan.

Helal: I'm going to say I'm 1.95. I'm holding that .05 from 2.0 because I think the tools have a responsibility to be clear to their users what they're doing for them. We do that in Artisan and all the Embarcadero tools where we actually allow you to preview any scripts if you would like to before we keep them on your behalf. We allow you to log the SQL so you can go back and utilize that SQL in your own scripting. And I think there are cases where you do preferably have to get off the tool and go back to the server. We will strive in the future to build tools that actually will fill the need to run things server side.

BIOGRAPHIES



Peter Dobler is an accomplished IT database professional who "makes a difference" by improving efficiencies and reducing costs for small and medium sized businesses. He founded Dobler Consulting, a Tampa, Florida consulting firm that delivers implementation expertise for Oracle, Sybase, and MS SQL Server in 2000. His 25 plus years in technology started in Basle, Switzerland in 1985. He currently engages in strategic alliances and special projects with Sybase, the enterprise software and services company. He's the author of two popular blogs including his Database Trends Blog (a behind the scenes look at the evolution of database technology) and his Technology Tips Blog (step-by-step instructions on today's challenging technology).



Kevin Sherlock is a Database Administrator for Children's Hospital in Omaha Nebraska and a member of TeamSybase. Kevin has been a frequent contributor to the ASE newsgroup community for many years and has extensive experience with all aspects of ASE including performance and tuning the server, and Transact-SQL. Kevin has been working with ASE since version 4.9.2, and also has experience working with Replication Server, and Sybase IQ. He was the recipient of the ISUG Innovation and Achievement Award in 2005, and the ISUG Technical Journal Award in 2008.



Harry Helal is Program Manager for Embarcadero's Database flag ship products DBArtisan and Rapid SQL. He has twenty years of database experience including ASE, IQ and Replication Server. Prior to joining Embarcadero, He was the Database Manager at Lavalife for more than six years. Prior to that, he was a Senior Sybase DBA consultant, providing services to many leading DBA teams including Pepsi Bottling Group, Enbridge Gas distribution, and Royal Bank Dominion Securities.



Scott Walz has more than 15 years of experience in database development and administration and currently serves as the senior director of product management for Embarcadero Technologies. In this position, Scott oversees the direction of the company's database product family, while focusing on database development and administration products. Prior to joining Embarcadero eight years ago, Scott served as a development lead for Louisville Gas & Electric. He holds a bachelor's degree in computer information systems from Western Kentucky University.

ABOUT US



ABOUT EMBARCADERO

Since 1993, Embarcadero has been the trusted tooling provider to thousands of Sybase ASE customers. Explore our [Sybase ASE solutions page](#) or [download a free trial](#). You will quickly discover why experienced data management professionals place their confidence in the depth of Embarcadero's Sybase-specific capabilities to get their jobs done quickly.

Embarcadero Technologies, Inc. is the leading provider of software tools that empower application developers and data management professionals to design, build, and run applications and databases more efficiently in heterogeneous IT environments. Over 90 of the Fortune 100 and an active community of more than three million users worldwide rely on Embarcadero's award-winning products to optimize costs, streamline compliance, and accelerate development and innovation. Founded in 1993, Embarcadero is headquartered in San Francisco with offices located around the world. Embarcadero is online at www.embarcadero.com.



ABOUT ISUG

Founded more than 20 years ago, the International Sybase User Group is an independent association that represents thousands of users of Sybase products in more than 60 countries around the world. Members network together through local, regional and virtual meetings to share ideas, solutions and insights into Sybase technologies, all of which results in increased learning, and thus increased earning potential. Membership in ISUG is obtained by simply joining one of ISUG's regional, local or virtual user groups. For more information on ISUG, or to become a member, please visit www.isug.com.

ADDITIONAL RESOURCES

Learn all about ASE, including features and product editions, on this Sybase site:
<http://www.sybase.com/products/databasemanagement/adaptiveserverenterprise>

This link goes to a number of Sybase presentations (you will need to log in):
<http://www.sybase.com/techwave/TWpresentationarchive>

Information about Sybase products is available at this site operated by the International Sybase User Group:
<http://www.isug.com/common/JournalArchives.html>

Communicate with product engineers and other users about specific issues related to Sybase products, on community forums/newsgroups:
<http://www.sybase.com/support/community-forums>

More information is available on the Sybase ASE database blog site:
<http://blogs.sybase.com/database/>