(October 23rd, 2017)

If you know someone who would benefit from being an Insider, feel free to forward this PDF to them so they can sign up here.

Quick Tips for our Insider friends!

Hey Insiders,

It's all go this month! After a couple of weeks packed with *8* Immersion Events in Chicago, we're home for two weeks working hard on Pluralsight courses and preparing for our Fall SQLintersection conference in Las Vegas next week. You can read more about our SQLintersection lineup in my blog post here. There are a bunch of incredible workshops and sessions by industry leaders, if you need a week to dive into SQL troubleshooting, performance tuning, and SQL Server futures – you still have time to join us. And, I've been practicing my craps strategy so hopefully I'll at least break even this year…

Note: you can get all the prior Insider newsletters here.

SQLskills News

Erin's latest Pluralsight course has been published! It's called ***SQL Server: Analyzing Query Performance for Developers***. Check out the details here.

Kimberly's latest Pluralsight course has been published! It's called ***SQL Server: Indexing for Performance*** and is a 7+ hour behemoth packed with her indexing wisdom. Check out the details here.

The first batch of 2018 classes is open for registration! We have all our usual classes and we'll be adding some more new classes over the coming weeks around security, data mining and AI, data due diligence/BI. See here for the 2018 Immersion Event class schedule. **Note: register in 2017 for discount prices!**

If not in Chicago, our team is presenting a number of workshops at Fall conferences, we hope you'll join us:

- Paul is presenting a workshop at the Fall SQLintersection conference in Las Vegas on Sunday, October 29th, titled *Performance Troubleshooting using Waits and Latches*. Check out the details here.
- Erin is presenting a workshop at the Fall PASS Summit in Seattle on Monday, October 30th, titled *Solving Common Performance Problems Using Query Store*. Check out the details here.

- Jonathan is presenting a workshop at the Fall SQLintersection conference in Las Vegas on Monday, October 30th, titled *Extended Events: WTF or FTW!* Check out the details here.
- Kimberly is presenting a workshop at the Fall SQLintersection conference in Las Vegas on Friday, November 3rd, titled *Very Large Tables: Optimizing Performance and Availability through Partitioning*. Check out the details here.
- Tim is presenting a workshop at the Fall SQLintersection conference in Las Vegas on Friday, November 3rd, titled *Common SQL Server Mistakes and How to Correct Them*. Check out the details here.

And all of us are presenting sessions at SQLintersection or the PASS Summit (Erin and Glenn).

Don't forget to check out our SQL101 posts… hopefully they'll help refresh or reinforce topics for even the more seasoned DBAs in the community. The blog posts are automatically collected here.

Finally, even if you can't join us in person, we've renewed our call for remote user group sessions for the second half of this year. We have almost 100 scheduled and completed so far; if you'd like one of us to present for your user group, check out my blog post here.

Book Review

One of the recent books I've read is Michael Riordan and Lillian Hoddeson's Crystal Fire: The Invention of the Transistor and the Birth of the Information Age. This book recounts the history of the development of the transistor, from the beginnings of solid-state physics in the late 1800s and early 1900s through to the invention and fabrication of the first microchips in the early 1960s. It's also something of a biography of William Shockley, who was arguably the central character in the transistor's creation, along with John Bardeen and Walter Brattain, fellow Bell Labs employees and co-winners of the 1956 Nobel Prize for Physics. Hugely interesting and very readable – highly recommended!

The Curious Case of…

This section of the newsletter explains problems we've found on client systems; they might be something you're experiencing too.

In our recent IEPTO2 class, I was talking to one of the students about a problematic process they had that seemed to be taking longer and longer to complete. I asked if they knew what the wait statistics were for the time covering the process and they didn't.

I set them up with two scripts: one to track all the waits over the half hour of the process (using the code in my post here that uses *sys.dm_os_wait_stats*) and another to look at the waits as they

occurred (based on my script here that uses *sys.dm_os_waiting_tasks*), running every 15 seconds and persisting into a table.

The second script was to let me see the resources being waited for, not just what waits were occurring. This is really important, because a wait type might show up and you can't tell whether it's normal or really bad.

As it happened, this was one of the cases where the second set of information was completely necessary. Their workload was generating a lot of *CXPACKET* waits, which we could see from the *sys.dm_os_wait_stats* output, but only with the output from *sys.dm_os_waiting_tasks* could we see that they had a skewed parallelism problem.

This is where the amount of work for each of the parallel threads for an operator is not distributed evenly. One thread (or a few threads) does most of the work and the other threads are stuck waiting. You can read more about skewed parallelism, with a simple example here.

The way to detect skewed parallelism is very simple: if you see any *thread ID* in my script (or *exec_context_id* in the raw *sys.dm_os_waiting_tasks* DMV output) other than 0 waiting for *CXPACKET*, you have skewed parallelism. My script will tell you the query plan node ID where it's happening and then you can figure out why the cardinality estimate is incorrect.

Bottom line: *CXPACKET* waits, in their simplest form, mean you've got parallelism happening, which you can prevent for a query in many ways (of which, increasing the *sp_configure* option *cost threshold for parallelism* is probably the easiest and best). But you should always make sure to check for skewed parallelism, as that can *really* cause performance problems.

Paul's Ponderings

For this newsletter, I invited Kimberly to write about a consistent and all too-common problem that's sometimes difficult to troubleshooting – application session settings causing execution problems. Enjoy!

You've probably seen this problem before… a stored procedure that works well in Management Studio but performs slowly or "differently" in the application.

There are a variety of possible causes for this but one of the most frustrating can be the wide variety of differences around session settings. ODBC DSNs (Data Source Names) used to be a significant cause of this many years ago, but with the resurgence in ODBC DSNs over the last few years, that could be something you've run into. And, it doesn't have to be an ODBC DSN that causes the problem; almost all applications have this potential for session settings affecting query performance – especially around one particular session setting: arithabort.

So, what *is* the problem? Stored procedure query plans can only be re-used by sessions with essentially the same session settings. And, some performance-related features can only be used

when a certain number of session settings are set consistently throughout an application (from creation to usage). Furthermore, because your applications can have different session settings, you might get inconsistent behaviors and performance problems across query executions.

So, if you have wildly varying performance of a procedure when used in an application (especially if it's different in SSMS), you're probably suffering from inconsistencies in session settings.

And, where and how these session settings are set can be equally frustrating. It's true that most are set by the application, but some can even be changed by the user – for example, for the ODBC DSN settings. So, if you have any applications that use ODBC DSNs and you're having inconsistencies in performance – the first thing to check is each of the application's session settings.

The easiest way to do this is with *sys.dm_exec_sessions*. About half-way through the output column list is a column called *is_user_process* and the columns that follow it describe many of the more critical session settings.

At a minimum, you can see if your applications (or a specific connection) have any settings that are not set consistently. Consistency is very important for stored procedure plan reuse.

But, I'd go further than that with some performance-related features (such as filtered indexes, indexed views, indexes on computed columns, and spatial indexes – to name just a few). Without the *right* session settings set, you might not be able to leverage those features from sessions that don't match. And, in the worst cases, applications that don't have the correct session settings (where you are using one of those features) will actually FAIL when modifying data in tables that leverage any of those features.

Call to action: Session settings are often overlooked as a cause of stored procedure performance problems, so it's good to understand them and make sure they're set correctly. I cover this in depth in modules 2 and 3 of my Pluralsight course on stored procedure performance – see here.

Glenn's Tech Insights

This section of the newsletter highlights recent news and views from the hardware and Windows worlds that we think will be interesting to SQL Server community members.

Windows 10 Fall Creators Update is Available

Microsoft released the Windows 10 Fall Creators Update on October 17. This update (which is similar to a Service Pack) has a number of interesting new features that may be useful to you, depending on your interests and usage patterns.

For example, Windows Defender has new ransomware protection, there is a slider that lets you more easily manage your power options, you can use Cortana voice control to shut down your machine, and there is a new GPU tab in Task Manager.

You can get more details about the all of the new features here. Microsoft will be pushing this update out in phases, or you can go here to download an Upgrade Assistant that will let you pull down the update.

New KRACK Attack Exposes Vulnerability in WPA2 WiFi Protocol

Mathy Vanhoef, a researcher from the University of Leuven (KU Leuven), has discovered a severe flaw in the Wi-Fi Protected Access II (WPA2) protocol that secures most modern protected Wi-Fi networks. The flaw affects the WPA2 protocol itself and is not specific to any software or hardware product.

KRACK stands for **K**ey **R**einstallation Att**ack**, and an attacker must be within the victim's WiFi range in order to use it. The only way to fix this vulnerability is to get updated firmware from the manufacturer of your WiFi capable device. You can get much more detail about this issue here. The academic paper behind this is here.

So far, this exploit is not known to be out in the wild. The proper course of action is to check to see if your device manufacturer has any information about their products that could be affected and whether they have a firmware update available. CERT has a Vulnerability Notes database here. Microsoft released a patch for this issue on October 10.

Microsoft Surface Book 2

Microsoft has announced the new Surface Book 2. They will have seven different configurations, ranging from a 13.5" Intel Core i5-7300U with 8GB of RAM, a 256GB PCIe SSD, and Intel HD 620 integrated graphics for $1,499 to a 15" Intel Core i7-8650U with 16GB of RAM, a 1TB PCIe SSD, and an Nvidia GeFore GTX 1060 discrete GPU for $3,299.

The Intel Core i7-8650U is a quad-core (plus HT) Coffee Lake processor that is significantly more powerful than the dual-core (plus HT) Intel Core i7-6600U Skylake processor used in the original high-end Surface Book models.

It would have been nice to have a 32GB RAM option, but the mainstream market probably would not support that.

#TBT

(Turn Back Time…) This section of the newsletter highlights some older resources we've referred to recently that you may find useful, plus select blog posts we've published since the previous newsletter.

Parallelism and CXPACKET waits is the theme for this TBT:

- My post: *Knee-Jerk Wait Statistics: CXPACKET*
- My post: *More on CXPACKET Waits: Skewed Parallelism*
- The waits library entry for *CXPACKET*
- Jonathan's post: *Tuning 'cost threshold for parallelism' from the Plan Cache*
- Jonathan's post: *SQL 101: Parallelism Inhibitors – Scalar User Defined Functions*
- My post: *Updated sys.dm_os_waiting_tasks script to add query DOP*
- My post: *Who is overriding MAXDOP 1 on the instance?*

Here are a few of the blog posts we've published since the last newsletter:

- Paul: *SQLskills SQL101: Should you kill that long-running transaction?*
- Paul: *SQLskills SQL101: How can corruptions disappear?*
- Glenn: *Windows Server Servicing Model Changes*

I hope you find these useful and interesting!

Video Demo

The demo video this time is taken from Kimberly's new Pluralsight course *SQL Server: Indexing for Performance*. In the demo, Kimberly explains how a nonclustered index that covers a query can be used to seek the results of a query.

The video is around 5 minutes long and you can get it in MP4/MOV format here.

The demo code is here.

Enjoy!

Upcoming Immersion Events

We've released the first set of 2018 classes for registration. Over the coming weeks we'll be adding a few more classes to the Spring line-up in Chicago, including classes on data mining and AI, security, and data due diligence/BI. We'll also be adding one or more classes in Europe, in London and/or Dublin, in the second half of the 2018.

To help your boss understand the importance of focused, technical training, we've also added a few items to help you justify spending your training dollars with us:

- Letter to your boss explaining why SQLskills training is worthwhile
- So why do you want to come to our training? And the winners are…
- Community blog posts about our classes
- Immersion Event FAQ

Chicago, IL, April/May 2018 **(all classes have discounts for 2017 registrations!)**

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - April 23-27
- **IE0**: Immersion Event for the Junior/Accidental DBA
  - April 23-25
- **IEUpgrade**: Immersion Event on Upgrading SQL Server
  - April 23-25
- **IECAG**: Immersion Event on Clustering and Availability Groups
  - April 26-27
- **IEAzure**: Immersion Event on Azure SQL Database and Azure VMs
  - April 26-27
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - April 30-May 4
- **IEPS**: Immersion Event on PowerShell for SQL Server DBAs
  - April 30-May 2
- **IESSIS1**: Immersion Event on Learning SQL Server Integration Services
  - May 7-11

Bellevue, WA, June 2018 **(all classes have discounts for 2017 registrations!)**

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - June 18-22

Click here for the main Immersion Event Calendar page that allows you to drill through to each class for more details and registration links.

Summary

We hope you've enjoyed this issue - we really enjoy putting these together.

If there is anything else you're interested in, we'd love to hear from you - drop us a line.

Thanks,
Paul and Kimberly

Paul@SQLskills.com and Kimberly@SQLskills.com