(August 15th, 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,

We'll be taking a short summer break from Insider newsletters, as per usual every year. The next newsletter will be on Monday, September 25th.

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

SQLskills News

Glenn's writing a one-a-day blog series in August! Glenn's latest Pluralsight course is *SQL Server: Upgrading and Migrating to SQL Server 2016*; his new blog series is a companion to that along with his new 3-day IEUpgrade class debuting THIS October. Check out his blog post series here.

Our October classes in Chicago are fast approaching! We have the usual IE0, IEPTO1, IEPTO2, IEPS, and IESSIS1 classes running, plus THREE new classes: **Azure SQL Database and Azure VMs**, **Upgrading SQL Server**, and **Clustering and Availability Groups**. See here for the complete October Immersion Event class schedule.

We ran a competition last week to win a free seat in our IEPTO1 and IEPTO2 classes and ended up giving away FIVE seats! See here for the complete list of cool entries and the winners!

Our team is presenting a number of workshops at Fall conferences, we hope you'll join us:

- Paul will be 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 will be 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 will be 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 will be 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 will be 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). You can read more about our SQLintersection lineup in my blog post here.

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 85 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

The latest book I've read is Ed Yong's *I Contain Multitudes: The Microbes Within Us and a Grander View of Life*, that I was recommended by my buddy Buck Woody. This is a really excellent book that is hugely interesting! It details the history and current state of scientific investigations into microbes and explains how symbiotic microbes are basically necessary for all of life on the planet to exist and flourish. This book will change how you view microbes and the microbiome – the microbes that live inside all of us. 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.

I was helping someone out with a performance problem on twitter last week, and the scenario is worthy of inclusion here. The person was doing some *DBCC CHECKDB* testing and was running two of them back-to-back on the same database. The first one would take around 13s and the second one would take around 5s.

I racked my brains thinking what could cause that, like random concurrent workload on the server, or random concurrent workload on a shared I/O subsystem, but the server and storage were completely inactive apart from this test.

On further questioning, additional symptoms came to light: after the initial back-to-back test, subsequent runs of *DBCC CHECKDB* in quick succession would all take about 5s. But if the server was left alone for half an hour, the first *DBCC CHECKDB* after the pause would take 13s again, and then drop down to 5s for subsequent runs.

That pattern led me to the solution: balanced power plan. *DBCC CHECKDB* is a huge CPU user, so the CPUs started off at low speed; the OS noticed the high CPU usage and clocked the CPUs higher. So the second run was at normal speed. Waiting for half an hour allowed the OS to notice that the high clock speed was no longer needed and lowered it again. This neatly accounted for all the symptoms.

Bottom line: one of the simplest things to check when trying to determine the cause of weird performance slowdowns is the OS and BIOS settings around balanced power and power savings. In general, servers running SQL Server should not have any kind of power savings enabled.

Paul's Ponderings

Just like last time, the Curious Case topic this time gave me inspiration for the Ponderings. Although OS/BIOS power savings modes are outside of SQL Server, there's one thing *inside* SQL Server that can have a dramatic effect on the workload performance and is impossible to see using performance troubleshooting techniques: Extended Events sessions.

We just finished teaching our IEPTO2 class last week in Bellevue, and we explain and demonstrate how Extended Events can sometimes be what we call a 'silent killer', because there are no symptoms – no wait statistics, no perfmon counter, no DMV metrics – that indicate an Extended Event session is what's dragging down performance.

For instance, I created a simple test with 200 concurrent clients inserting into a table with a GUID cluster key and another 100 bytes of random columns. On my laptop, the workload averaged 22,000 transactions per second.

Next I created an Extended Event session that captured a SQL Server code call stack every time a wait occurred. Note that this absolutely isn't an event session I'd ever recommend as it's guaranteed to cause a performance problem – I'm just using it as an example.

As soon as the event session started, transactions per second dropped to 16,000 – roughly 25% reduction in throughput.

I collected wait statistics before and after enabling the event session and the only difference was a slight increase in resource wait time for the *PAGELATCH_EX* waits when the session was enabled. Now, the first thing someone's going to think of with a slightly elevated resource wait time is that there's increased contention, not that there's extra CPU being used for each latch wait.

The reason Extended Events can cause this kind of performance issue is that the event processing is performed by the executing thread synchronously, so there's more code executed every time the event fires.

Now don't get me wrong – Extended Events are wonderful and they are the way forward for deep performance monitoring instead of Trace/Profiler. Just like with any performance monitoring tool, however, you need to be very careful that the monitoring itself is not causing a performance problem (called 'observer overhead').

Call to action: There are no metrics that you can reliably use to see Extended Events pulling down performance, so if you have an unexplained performance drop, check the active sessions using the DMV *sys.dm_xe_sessions*, and note that depending on your version of SQL Server, you're going to see one of more of these normal, system-defined sessions: *system_health*, *sp_server_diagnostics session*, *hkenginexesession*.

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.

New Enterprise Flash Storage from Toshiba

Toshiba has announced development of their newest flagship enterprise SSDs, both featuring Toshiba's latest generation of 64-layer 3D BiCS FLASH memory. The PM5 is a 12Gb/sec series of SAS drives, while the CM5 is their new enterprise NVMe drive series. Both are built with Toshiba's latest 64-layer 3D 3-bit-per-cell enterprise-class triple-level-cell (TLC) BiCS FLASH.

The PM5 SAS SSD comes in a 2.5: form factor, and will be offered in capacities between 400GB, and 30.72TB. The CM5 comes in U.2/2.5" and add-in card (AIC) form factors and utilizes an NVMe dual-port PCIe 3.0 x4 interface. The CM5 comes in capacities ranging from 800GB to 15.36TB. These drives are supposed to be available in Q4 of 2017.

SQL Server 2017 Benchmark Results

So far, there have been two official TPC-H benchmark submissions for SQL Server 2017 Enterprise Edition. The most recent result of 1,336,109 QphH@10000GB from July 9, 2017 is for a four-socket Lenovo ThinkSystem SR950 server with four Intel Xeon Platinum 8180M 28-core processors and 6TB of DDR4 2666 RAM, and eight ThinkSystem HHHL 3.2TB Performance NVMe PCIe 3.0 x4 flash storage cards (which are OEM Toshiba PX04P cards), running on Windows Server 2016 Standard Edition. In case you are wondering, this four socket server has 112 physical cores, and actually supports up to 12TB of RAM.

The other, somewhat older TPC-H submission on SQL Server 2017 Enterprise Edition from April 17, 2017 is a score of 717,101 QphH@1000GB on a two-socket HPE Proliant DL380 Gen 9 server with two Intel Xeon E5-2699 v4 22-core processors and 512GB of DDR4 2400 RAM. This system is using a combination of 400GB 12Gb SAS SSDs and 300GB 6Gb SAS HDDs, running on Red Hat Linux Server 7.3.

There has been one official TPC-E benchmark submission for SQL Server 2017 Enterprise Edition from June 27, 2017, with a score of 6598.36 on a two-socket Lenovo ThinkSystem SR650 server with two Intel Xeon Platinum 8180M 28-core processors and 1.5TB of DDR4 2666 RAM running on Windows Server 2016 Standard Edition.

What is interesting is that all three of these benchmarks were done before July 9, 2017, which was several weeks before even SQL Server 2017 RC1 was released. Normally, Microsoft does not allow companies to publish benchmark results on pre-release software, and they definitely have veto power over any TPC-H or TPC-E submissions. I am confident that the RTM version of SQL Server 2017 will have even better performance, especially since the release notes for SQL Server 2017 RC2 (which was released on August 2, 2017), states "This release contains bug fixes and performance improvements."

AMD Ryzen Threadripper Released on August 10, 2017

AMD has finally released their new super high-end desktop (SHED) processor family, the AMD Ryzen Threadripper. There are two initial SKUs in the lineup, the 1950X and the 1920X. The $999.00 Threadripper 1950X has 16 cores plus SMT (which is AMD's version of hyper-threading), so you get 32 logical cores total. The $799.00 Threadripper 1920X has 12 cores plus SMT, for a total of 24 logical cores. Other SKUs with lower core counts will be added over the next few months. These first two models are actually available for sale today, so this is not a "soft launch". Multiple motherboards are also available to support this new processor family.

These processors use the X399 chipset, and they will all support 64 PCIe 3.0 lanes (less 4 for the chipset), which means you will have a lot of flexibility and capacity for things like multiple discrete video cards, multiple M.2 NVMe storage devices, and 10GbE networking.

They also have eight DDR4 DIMM slots that support ECC RAM, so you can have up to 128GB of quad-channel RAM with 16GB DIMMs. The processor itself will support up to 1TB of RAM as larger DIMMs become available and affordable.

Another interesting fact about these processors is that they have two NUMA nodes at the hardware level. NUMA mode can be disabled in the BIOS (since some desktop/consumer software is not NUMA-aware).

SQL Server 2016 has a new feature called automatic soft NUMA which is enabled by default when you have more than eight logical processors in a NUMA node. You can disable automatic soft NUMA with an *sp_configure* setting. Both of these Threadripper models have enough logical cores to let you experiment with hardware NUMA and with automatic soft NUMA.

The bottom line here is that you can buy/build a very powerful desktop machine for virtualization or heavy duty SQL Server development and testing usage for a lot less money than if you use an Intel Skylake-X HEDT platform. Here are a few initial reviews:

- The AMD Ryzen Threadripper 1950X and 1920X Review: CPUs on Steroids
- Ryzen Threadripper review: AMD's monster 1950X stomps on other CPUs
- AMD Threadripper 1950X review: Better than Intel in almost every way
- The AMD Ryzen Threadripper 1950X and 1920X Review

#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.

Glenn's half-way through his upgrade blog-a-day series, so that's the theme for this TBT:

- Glenn's course: *SQL Server: Upgrading and Migrating to SQL Server 2016*
- Glenn's course: *SQL Server: Installing and Configuring SQL Server 2016*
- Tim's course: *SQL Server: Consolidation Tactics and Best Practices*
- Glenn's course: *SQL Server 2012: Evaluating and Sizing Hardware*
- Glenn's upgrade blog post series

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

- Paul: *SQLskills SQL101: REBUILD vs. REORGANIZE*
- Paul: *So why do you want to come to our training? And the winners are…*
- Glenn: *SQL Server Diagnostic Information Queries for August 2017*

I hope you find these useful and interesting!

Video Demo

In the demo video this time, Tim demonstrates how you can expose certain host level information when using Azure virtual machines. Having this information can be very useful when troubleshooting any performance related issues.

The video is around 6.5 minutes long and you can get it:

- In WMV format here
- In MOV format here

Enjoy!

Upcoming Immersion Events

All 2017 classes are available for registration!

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, October 2017

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - October 2-6
- **IE0**: Immersion Event for the Junior/Accidental DBA
  - October 2-4
- **IECAG**: Immersion Event on Clustering and Availability Groups
  - October 5-6 **** New class! ****
- **IESSIS1**: Immersion Event on Learning SQL Server Integration Services
  - October 2-6
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - October 9-13
- **IEPS**: Immersion Event on PowerShell for SQL Server DBAs
  - October 9-11
- **IEAzure**: Immersion Event on Azure SQL Database and Azure VMs
  - October 9-10 **** New class! ****
- **IEUpgrade**: Immersion Event on Upgrading SQL Server
  - October 11-13 **** New class! ****

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