(June 11th, 2018)

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.

Note: As an Insider, you can read all prior Insider newsletters here.

Quick Tips for our Insider friends!

Hey Insiders,

It's a very special day for us today as our eldest daughter graduates from high school and speaks as her class salutatorian this afternoon; we're both incredibly proud of her and also amazed to think how fast time flies!

In SQLskills news, we're excited to announce **three brand-new live, online classes with special newsletter discounts** – see below!

And coming up next week, our Bellevue IEPTO1 class has a couple of seats left; it's not too late to register and spend some of that fiscal budget that might be lost at the end of June! Additionally, we still have a few seats remaining in our London classes in September; they're our last, in-person Immersion Events to be delivered this calendar year!

SQLskills News

Live, online classes: there are three new classes available for registration:

- August 28-30: ***IEPUM2017****: Immersion Event on Planning and Implementing an Upgrade/Migration to SQL Server 2017* (taught by Glenn)
- October 9-11: ***IETLB****: Immersion Event on Transactions, Locking, Blocking, Isolation, and Versioning* (taught by Kimberly)
- October 23-25: ***IEQUERY****: Immersion Event on Fixing Slow Queries, Inefficient Code, and Caching/Statistics Problems* (taught by Erin, Jonathan, and Kimberly)

These classes will be delivered live via WebEx 10am-3pm PST, Tuesday-Thursday (roughly 12-13 hours of content including open Q&As; similar to two full workshop days without leaving the comfort of your home/office!) and you also get access to the class recordings. By dedicating only 3 half-days of your time you still have time to get some work done during the day and with lifetime access to the recordings, you get amazing ROI!

Each of these classes are priced at US$699 and we're offering a combo package of all three for US$1,749, saving US$350. And as a special discount ONLY for newsletter subscribers, **use the**

discount code 'newsletter' before July 1st (this offer will NOT be extended) to save US$100 off each class (US$599 each) or an additional US$150 off the combo price (US$1,599 for the three-course combination for a total saving of US$500!). Click on the class links above for all the details.

In-person US classes: Our single, remaining, in-person class in the US for 2018 is IEPTO1 in Bellevue, WA next week.

In-person London classes: We're bringing four of our Immersion Events to London in September: IEPTO1 and IEPTO2, plus our new classes: **IEAzure** on Azure, Azure VMs, and azure Managed Instance and **IECAG** on clustering and availability groups **(register for one of the new classes and get the other one half-price!)** See here for all the details.

Finally, even if you can't join us in person, I've put out a call for 2018 remote user group sessions. In 2017, we did more than 100 of these around the world and we have set up more than 70 for 2018 already! If you'd like one of us to present for your user group, check out my blog post here. **Note: Tim has a new user group session on Azure Managed Instance** that he's happy to present to your group – see here for details.

Book Review

The latest book I've read is Colm Toibin's *The Blackwater Lightship*. It turns out I've read this before, back in 2011, and I realized half-way through but I finished it because it's so good. It's an excellent story about a set of estranged Irish women (daughter, mother, grandmother) who are forced back together as they deal with the decline of the daughter's brother who is gay and very close to dying of AIDS (which they didn't know until one of his friends turns up out of the blue). Toibin does a great job of exploring their emotions as daughter and mother slowly reconcile, and the daughter recollects the events leading to all the estrangement. More interesting reading it this time around as I've been to Dublin and around in Ireland a lot since 2011 and know all the places described. Highly recommended!

The Curious Case of…

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

I had a question from someone last week who was doing wait statistics analysis and was seeing a huge number of *SOS_SCHEDULER_YIELD* waits but couldn't see anything in *sys.dm_os_waiting_tasks*. I explained that they wouldn't show up there as that wait type means a thread has yielded but isn't waiting for a resource (see this post for a deeper explanation) and so they'd have to use *sys.dm_exec_requests* instead (see here for the script to use).

They still couldn't see anything, which was really weird.

I asked if they are using Azure SQL Database, and they are. I then asked how they were seeing all the *SOS_SCHEDULER_YIELD* waits, and they were using my regular waits script that uses *sys.dm_os_wait_stats*.

Aha! I explained to them that on Azure SQL Database, you must use *sys.dm_db_wait_stats* instead; otherwise you're seeing all the waits for the host instance of SQL Server, which is the wrong set of metrics. I pointed them to Tim's blog post that explains how to look at wait stats and I/O metrics when using Azure SQL Database, and that solved their analysis problem.

Bottom line: If you're using Azure SQL Database, make sure you're using the correct wait stats DMV otherwise you won't be able to properly interpret the results.

Paul's Ponderings

Last week I presented to two user groups in the UK, Manchester and Leeds, which we usually do on consecutive days and on the same topic and this time it was my favorite – wait statistics. I had two follow-up questions in email about preemptive wait types that I thought would be useful to explain here, and about preemptive waits in general.

SQL Server uses non-preemptive scheduling, where it controls what happens with thread scheduling rather than allowing the OS to do the scheduling (obviously, while running within the OS schedulers too, but it's best to just ignore that when thinking about how SQL Server performs scheduling.)

When a thread has to call an API outside of SQL Server, SQL Server relinquishes control of the thread – i.e. the thread goes preemptive and incurs a *PREEMPTIVE_XX_YY* wait type, depending on what the thread is doing. These wait types were added to SQL Server in SQL Server 2008 and are pretty much undocumented (well, nearly all wait types are only documented to the point of a single sentence in the documentation for *sys.dm_os_wait_stats*, but most of the preemptive ones aren't even mentioned!). I've documented some of them in my waits library, and for the *PREEMPTIVE_OS_XX* ones where I don't have in-depth guidance, I've at least made clear what OS API the thread is calling.

Now for the questions I was asked…

Firstly, someone said they're seeing increasing *PREEMPTIVE_OS_CREATEFILE* waits on one instance. I asked if they're using *FILESTREAM* data and they are. I then asked if they'd configured the NTFS volume hosting the *FILESTREAM* data container correctly and they didn't know they had to, although they suspected the waits were something to do with NTFS.

When a *FILESTREAM* file is created, a *PREEMPTIVE_OS_CREATEFILE* wait is incurred while the thread calls out to the OS to create the file. If the NTFS volume is not prepared properly (specifically, disabling 8.3 name generation), the file creation will start to take longer and longer

as more and more files are created and the wait time will increase. You can read about correctly preparing for *FILESTREAM* use in the *FILESTREAM* whitepaper I wrote for Microsoft back in 2008.

Secondly, someone said they see lots of *PREEMPTIVE_OS_WRITEFILEGATHER* waits after every full backup that they perform. That didn't make sense to me as that wait occurs when a new file (or new portion of an existing file) is zero-initialized, and nothing a data backup does would cause that to happen. I investigated a bit more and found that they were shrinking the log before every full backup, and so a lot of log growth was happening. I explained that they should stop shrinking the log, and the issue would go away, which it did.

Preemptive waits shouldn't be ignored, as any wait type that becomes prevalent on your system may indicate a performance problem that should be addressed.

Call to action: Any time you see a wait type that you don't understand fully or you've never seen before, check it out in my waits library (and if there's no information there, shoot me an email using the link on the waits library page for the wait type you're interested in). And if you've never looked at wait statistics before, check out my seminal post: *Wait statistics, or please tell me where it hurts*.

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.

Intel Announces Core i7-8086K Processor

On June 5, 2018, Intel announced the special limited edition Core i7-8086K processor to celebrate the fortieth anniversary of the Intel 8086 processor (which was the first x86 processor).

The Core i7-8086K desktop processor has six cores (plus hyper-threading), with a base clock speed of 4.0 GHz and a Turbo clock speed of 5.0 GHz and a 12MB L3 cache. This is basically a souped-up version of the Intel Core i7-8700K processor that is extremely well-suited for a gaming PC, or just a very fast mainstream desktop system.

Intel is having a one-day sweepstakes starting on June 7, 2018 at 5:00PM PST where they will be giving away 8086 of these processors (which will have a SRP of $425.00 each).

Intel Demonstrates 380GB Intel Optane 905P M.2 SSD

Intel demonstrated its upcoming 380GB Intel Optane 905P M.2 SSD during the Computex show on June 5, 2018. The demonstration showed two systems side-by-side while they encoded the same video. One system used the Optane SSD 905P while the other system used an Intel 760P flash-based M.2 SSD.

The Optane SSD supplied data fast enough to utilize 82% of the CPU and 86% of the GPU. The SSD 760p with flash memory and a Silicon Motion SM2262 controller only allowed the system to utilize 13% of the CPU and 17% of the GPU. Since these are in an M.2 2280 form factor, they are suitable for mobile use.

Speaking from personal experience, using an Optane SSD as your boot drive makes a very noticeable difference in everyday performance compared to flash-based M.2 NVMe SSDs. I have two desktop systems with the slightly older Intel Optane 900P PCIe NVMe drives that I have been extremely happy with. The enterprise version of these drives (the Intel Optane DC P4800X) performs extremely well for heavy SQL Server tempdb workloads.

Microsoft Releases SQL Server 2016 SP2 CU1

On May 30, 2018, Microsoft released both SQL Server 2016 SP2 CU1 and SQL Server 2016 SP1 CU9. Ideally, you should try to get on the SP2 branch as soon as you are ready.

These two important tempdb fixes are in both branches:

- PFS page round robin algorithm improvement in SQL Server 2016
- Performance issues occur in the form of PAGELATCH_EX and PAGELATCH_SH waits in TempDB when you use SQL Server 2016

#TBT

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

The transaction log is the theme for #TBT this time. Here are some transaction log resources for you:

- My 8-hour Pluralsight courses: *SQL Server: Logging, Recovery, and the Transaction Log*
- My *transaction log blog category*
- My *sqlperformance.com blog posts* on transaction log configuration and performance troubleshooting
- Kimberly's post on *VLF sizing*
- My TechNet Magazine article on *Understanding Logging and Recovery in SQL Server* (the formatting is messed up a bit on that web page)
- The *WRITELOG* page from my waits library

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

- Paul: *CXCONSUMER wait type – history and what you need to know*
- Glenn: *SQL Server Diagnostic Information Queries for June 2018*

- Glenn: *Speaking at PASS Summit 2018*

I hope you find these useful and interesting!

Video Demo

In this demo, Glenn demonstrates the new cross-platform SQL Operations Studio management and development tool. This tool is still in a preview status, and it is open-source, hosted in GitHub. You can download it here.

The video is just under 3 minutes long and you can get it in WMV format here.

Enjoy!

Upcoming SQLskills Events

We have events coming up in 2018 – from our ***new, live, online*** courses to our own live, in-person Immersion Events to our own conference: SQLintersection; all of our Fall events are open for registration.

Each and every event has a different focus as well as different benefits – from deep-technical training in our Immersion Events to wide-ranging topics at SQLintersection where you can learn more effectively how to keep moving forward in both your environment and your career! And, of course, one benefit you'll always get from in-person events is networking; we hope to meet/see you at an event soon!

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

LIVE, ONLINE Immersion Events:

- **IEPUM2017**: Immersion Event on Planning and Implementing an Upgrade/Migration to SQL Server 2017
  - August 28-30 **(** NEW **)**
- **IETLB**: Immersion Event on Transactions, Locking, Blocking, Isolation, and Versioning
  - October 9-11 **(** NEW **)**
- **IEQUERY**: Immersion Event on Fixing Slow Queries, Inefficient Code, and Caching/Statistics Problems

    - October 23-25 **(** NEW **)**
- **Special savings on the three-course combination when registering before July 1; see the SQLskills News section above for full details!**

LIVE, IN-PERSON Immersion Events:

Bellevue, WA, June 2018

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
    - June 18-22 **(** Buy 2, get 1 free!, *only 8* seats remaining **)**

London, UK, September 2018

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
    - September 10-14
- **IEAzure**: Immersion Event on Azure SQL Database and Azure VMs
    - September 10-11 **(** Buy IEAzure, get IECAG half-price **)**
- **IECAG**: Immersion Event on Clustering and Availability Groups
    - September 12-13 **(** Buy IECAG, get IEAzure half-price **)**
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
    - September 17-21

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