(September 10th, 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,

This newsletter is coming to you from London, UK where we're just finishing up day 1 of our IEPTO1 class and Tim's finishing the first day of his IEAzure class. There are a couple of seats still available in our IEPTO2 class next week – see here for details.

If you're waiting to find out our Spring 2019 schedule, we'll be announcing that in the first newsletter in October. Outside of our London events this week and next, we only have online courses running for the rest of 2018; we hope to see you online or at a conference (PASS in November and/or SQLintersection in December)!

SQLskills News

The big news this time is that **SQLskills is now a Microsoft Gold Data Platform Partner**, which enables us to really deepen our relationship with the Microsoft field and be involved in many more Microsoft-driven projects and initiatives.

Erin will be presenting her popular, full-day workshop on Query Store at SQLSaturday Portugal on Friday, September 21st – see here for details.

It's time to start thinking about attending a conference in the Fall, and for the first time in four years, **our own awesome SQLintersection Fall 2018 show doesn't clash with the PASS Summit**. SQLintersection is during the first week of December with 18 top-notch speakers delivering 40 sessions and 9 full-day workshops! We're really looking forward to it; our show is smaller than many others so we really have a chance to chat with folks and get your questions answered. And one of our favorite evening events – SQLafterDark – is always a huge hit! Check out all the details here.

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

- 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!). 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 is priced at US$699 and there's a past-attendee price of US$599. Click on the class links above for all the details.

Finally, even if you can't join us in person, I've put out a call for second-half 2018 remote user group sessions and we've done more than 60 this year already! 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 Donna Tartt's *The Secret History*. What an excellent book! It's about a small group of students at an elite New England college who learn ancient Greek under the tutelage of an eccentric professor. They're very close-knit and descend into moral bankruptcy when they commit a ritual murder. From that point on the group gradually fractures and they all become psychologically disturbed, with shocking results (without giving away too much). It's extremely well-written and very engaging – I literally couldn't put it down during the second half (also thanks to jet-leg on arrival here in London).

Another book I've read since the last newsletter is Bernard Cornwell's *Warriors of the Storm*. This is the ninth book in Cornwell's epic Saxon Tales series (the basis for the excellent T.V. series *The Last Kingdom*) that chronicles the making of the England, from around 900A.D. onwards. The series follows the life of a Dane, Uhtred of Bebbanburg (Bamburgh Castle in NW England), who ends up being the most formidable warrior in the land, fighting against the Danes on the side of Mercia and Wessex. Fantastic stuff and I'm about to start the 10th installment.

Both books are highly recommended (and *The Last Kingdom* T.V. series too)!

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 an interesting conversation recently with someone (who's actually attending IEPTO1 in London this week) about a problem they were having trying to drop a filegroup – where SQL Server refused to do so even though it seemed like nothing was using it and the filegroup had no files!

They eventually called Product Support, who gave them a query to find any tables with allocation units stored on the filegroup:

```
SELECT
    [t].[name]
FROM
    sys.tables AS [t]
INNER JOIN sys.filegroups AS [f]
    ON [t].[lob_data_space_id] = [f].[data_space_id]
WHERE [f].[name] = 'MY_FG_NAME';
```

This showed a number of tables with empty *ROW_OVERFLOW* allocation units on the filegroup.

How can that be, when the filegroup doesn't have any files?

Well, there's been a weird idiosyncrasy in the Storage Engine since SQL Server 2005, where sometimes it will create an empty *ROW_OVERFLOW* allocation unit if a table has multiple variable-length columns, *just in case* one of them needs to be pushed off-row. As the allocation unit isn't being used, it doesn't even have a single IAM page, and so doesn't require any storage at all. Hence the ability for it to be on a filegroup without any data files.

Weird!

Bottom line: Another example of understanding how SQL Server works (or kind of doesn't, in this case!) can save you bashing your head against a brick wall trying to figure something out.

Paul's Ponderings

In this editorial I thought I'd list the top-5 wait types whose meaning people most commonly misconstrue. As a kind of test for yourself, first I'll give you the list so you can make a note of what you think they mean:

- *ASYNC_NETWORK_IO*
- *OLEDB*
- *PAGEIOLATCH_??*
- *PAGELATCH_??*
- *SOS_SCHEDULER_YIELD*

Since I published my wait stats library back in March 2016 I still get random emails from people asking about various wait types and how to alleviate them, and the questions usually show misconceptions about the problematic wait types. Sometimes they've even been led astray by a 3rd-party performance monitoring tool that's giving them incorrect advice.

Now here's the same list with quick explanations and common misconceptions (and links to much more comprehensive info in the waits library):

- *ASYNC_NETWORK_IO*
  - Simple meaning: A thread is waiting for the client application to acknowledge receipt of the data it was just sent. The only time this wait is caused by SQL Server itself is when MARS (Multiple-Active Result Sets) is used.
  - Common misconception: There's a network problem. IMHO it's a very confusingly-named wait type.
- *OLEDB*
  - Simple meaning: The OLE DB transport is being used.
  - Common misconception: It's a linked-server performance issue. Many other things use OLE DB, including most DMVs and *DBCC CHECK** commands.
- *PAGEIOLATCH_??*
  - Simple meaning: A thread is waiting for a page read from disk.
  - Common misconception: There's an I/O subsystem problem that's causing disk reads to slow down. IMHO it's more likely there's something going on in SQL Server that's causing more reads to be issued – like a plan change or memory pressure on the buffer pool that means all the data required by the query isn't in memory.
- *PAGELATCH_??*
  - Simple meaning: A thread is waiting for access to a data file page that is *already* in memory. Common causes are insert-hotspots, random inserts, tempdb allocation bitmap contention.
  - Common misconception: This is a read from disk, and so adding memory or messing with the I/O subsystem will help.
- *SOS_SCHEDULER_YIELD*
  - Simple meaning: A thread is able to execute for 4ms without any contention for a resource. An example might be a scan of a large table that's all in memory with no concurrent changes. It could be an indication of a plan change to a less efficient plan, such as someone reducing MAXDOP too far. It could also be from virtual machine scheduling problems.
  - Common misconception: SQL Server needs more CPUs. This wait isn't the indicator of that – long average signal wait times are an indicator of that, if coupled with workload throughput slowdown.

How well did you do?

Call to action: Whenever you're doing performance troubleshooting, and whatever metrics you're looking at, you *must* know what the metrics actually mean otherwise you're in danger of jumping to incorrect conclusions, wasting time, and possibly doing the wrong thing to try to alleviate your performance problem.

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.

Microsoft Releases SSMS 17.9

On September 5, 2018, Microsoft released SQL Server Management Studio (SSMS) 17.9. This release includes additional Showplan improvements that include row mode memory grant feedback attributes that are available when you are using the Adaptive Query Processing improvements available in SQL Server 2017 Enterprise Edition. You get these improvements when you use compatibility mode 140.

There is also support for the new vCore SKUs for Azure SQL Database, which allows you to view and change your service objective for these new vCore SKUs, (at least some of them). The dropdown for "Current Service Level Objective" only shows a subset of the Gen 4 vCPU SKU levels, and does not show the Gen 5 vCPU SKU levels at all.

There are also a pretty large number of bug fixes in this release. You can use SSMS 17.x to manage SQL Server 2008 through SQL Server 2017, along with Azure SQL Database. SSMS 17.x will run on Windows 7 SP1 and later, so it should be pretty easy for most people to be able to use. You can download it here.

Networking Improvements in Windows Server 2019

Microsoft has a series of blog posts that go into quite a bit of detail about networking improvements that are included in Windows Server 2019.

Two of the more interesting features are related to improving network performance for virtual workloads. These include:

- Receive Segment Coalescing (RSC) in the vSwitch
- Dynamic Virtual Machine Multi-Queue (d.VMMQ)

Receive Segment Coalescing (RSC) is a hardware off-load to the NIC that was disabled in Windows Server 2016 when you used a virtual switch. In Windows Server 2019, it works with virtual workloads and is enabled by default. This increases throughput and reduces CPU utilization on the host.

Dynamic Virtual Machine Multi-Queue allows Windows Server 2019 to automatically and ***dynamically*** use multiple processor cores to efficiently handle very high network bandwidth requirements, scaling the number of processor cores used as needed. This feature requires a driver update to your NICs to a Dynamic VMMQ capable driver, which you will have to get from your NIV vendor.

Microsoft has been quite open with multiple public preview releases of Windows Server 2019, along with many blog posts outlining improvements in Windows Server 2019. Windows Server 2019 is due for GA release in the 2H of 2018.

AMD Ryzen Threadripper 2950X Processor

The second generation, 12nm AMD Ryzen Threadripper 2950X processor went on sale on August 31, 2018. This is the 16 core/32 thread replacement for the 14nm AMD Threadripper Ryzen Threadripper 1950X 16 core/32 thread processor from 2017.

The older 1950X has a base clock speed of 3.4GHz with a Turbo Core clock speed of 4.0GHz, while the new 2950X has a base clock speed of 3.5GHz with a Turbo Core clock speed of 4.4GHz. That is not much of a jump in base clock speed, but there are several architectural improvements that will allow more cores to spend more time running at higher clock speeds compared to the 1950X, which means you will actually get a noticeable performance improvement.

This new processor will work in existing X399 motherboards after a BIOS update. The general consensus is that the 2950X is a better choice for most workloads than the flagship AMD Ryzen 2990WX 32 core/64 thread processor.

I'd like to hear if people are interested in a component selection/build guide for a HEDT workstation using a Threadripper processor. Please let me know!

#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.

As I was processing Erin's Insider Video I was reminded of four Pluralsight courses that Joe Sack recorded that are all about identifying and fixing esoteric query plan performance issues, so these are what I'm highlighting this time:

- Joe's course: *SQL Server: Common Performance Issue Patterns*
- Joe's course: *SQL Server: Common Query Tuning Problems and Solutions – Part 1*
- Joe's course: *SQL Server: Common Query Tuning Problems and Solutions – Part 2*
- Joe's course: *SQL Server: Troubleshooting Query Plan Quality Issues*

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

- Paul: *SQLskills is now a Microsoft Gold Data Platform Partner*
- Erin: *ALTER DATABASE SET QUERY_STORE command is blocked*
- Glenn: *Perfect Storm for Upgrading to a Modern Version of SQL Server*

I hope you find these useful and interesting!

Video Demo

Tracking down what's really going on with UDFs can be a frustrating task, but with Extended Events it gets a bit easier. In this video, Erin shows you how to use Extended Events to understand what is really executing with a UDF, and also shows some additional performance data that exists in 2016 and higher. (Jonathan's blog post that Erin references is here.)

The video is about 10 minutes long and you can get it in WMV format here.

The demo code is 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. Our first set of 2019 events will be announced in early October.

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:

- **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 **)**

LIVE, IN-PERSON Immersion Events remaining for 2018:

London, UK:

- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - September 17-21 **(** 2 SEATS REMAINING **)**

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