(September 25th, 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 somewhere on the M6 motorway in England! We had a really great two weeks of teaching in London and now we're headed up to Scotland to spend a few days with my parents in Helensburgh, the town I grew up in, and do a little fishing.

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

By the way, if you're at the **Microsoft Ignite** conference in Orlando this week, stop by the Azure SQL DB booth and say hi to Tim. He's there presenting the following sessions:

- THR2084 - How to migrate from on-premises to Azure SQL Database Managed Instance
- THR2085 - Azure SQL Database Managed Instance comparison to Azure SQL Database

SQLskills News

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 (*and great fun too*)! 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* (delivery by Kimberly)
- October 23-25: ***IEQUERY**: Immersion Event on Fixing Slow Queries, Inefficient Code, and Caching/Statistics Problems* (delivered 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

Over the summer I read a lot of classics that have been on my to-do list for a long time. One such book is Bram Stoker's *Dracula*. I love the Francis Ford Coppola movie (and I know many don't) so I know the story well, but I've never read the original book until now. The classic vampire story is quite a page-turner, and written from the perspective of several of the protagonists through the medium of their diaries. Having the movie in my head helped me visualize a lot of the scenes from the book and I'm happy to say that although Coppola and the screenwriter(s) used a lot of artistic license, they did a good job of staying faithful to the story. For any fans of Gothic horror (and really anyone who hasn't read it), this is 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.

While teaching the wait statistics module of our IEPTO2 class last week, one of the students asked me about a strange situation they'd seen on one of their instances.

They noticed suddenly that the *VDI_CLIENT_OTHER* wait type had become the most prevalent wait on the server and was steadily increasing in aggregate wait time. They weren't using any third-party backup software that might be using the VDI mechanism for performing backups, so were stumped as to what was causing it.

I asked them if that instance was part of an Availability Group – and yes, it is. Then I asked if they'd created a new secondary replica recently using automatic seeding – yes, they had.

Aha!

The AG seeding mechanism creates threads on the primary and new secondary replicas and then uses the VDI mechanism to perform the seeding. When one of these threads is waiting for something to do, it suspends itself and registers the *VDI_CLIENT_OTHER* wait type.

Furthermore, these threads will stick around until the next instance restart, accruing more and more wait time for *VDI_CLIENT_OTHER*. In other words, it's a benign wait type that can be ignored.

Bottom line: Any time you see a wait type you don't recognize or understand, look it up in my waits library. Take the name of the wait you're interested in and append it to the following URL https://www.sqlskills.com/help/waits/ and you'll get to the specific page for that wait type within the waits library. If that wait is TBD in the library, click the link on the page to send me an email about it.

Paul's Ponderings

I've been busy teaching and updating the waits library this week and based on some discussions with attendees of our London classes, I decided to re-run this editorial from early 2015 as it's still hugely relevant today. Enjoy!

Last week Jonathan pointed me at an interesting story about a psychology experiment involving monkeys and bananas and the reinforcement of constrained, negative thinking. The experiment actually never happened, but the story is quite illustrative. You can see a graphic of the story here: http://i.stack.imgur.com/MyQki.jpg (totally safe for work), and I'll paraphrase quickly below:

- A group of monkeys are trained to know that if any one of them attempts to climb a stepladder with a banana at the top, cold water is sprayed on all of them (i.e. temptation is prevented by the spraying of cold water, so the group prevents any individual from trying to get the banana).
- When a new monkey is introduced, and tries to get the banana, all the other monkeys prevent it as they don't want to be sprayed with water. The new monkey learns that it can't get the banana.
- Eventually all the monkeys are replaced, and they all know not to try to get the banana, but none of them know why, as they'd only been told not to get the banana by the others.

It's a really interesting story of how conditioning can become groupthink (even if the experiment didn't actually happen).

There are obvious parallels in human society, and specifically in the work environments of many of you reading this newsletter. Let me explain:

- A new person (A) joins a team. The new person has a great idea for how to do something differently/better, and everyone else on the team prevents them from expressing their opinion because that won't be allowed, it won't work, and/or they could get into trouble (e.g. from an intransigent and influential boss/architect/senior developer).

- Eventually all the original people leave the team and it's only person A and the people who joined the team after A that are left. They all have the ingrained knowledge that they can't try anything new, or try a specific technology because they won't be allowed to etc.

In other words, being constrained to incumbent technologies and methodologies becomes "the way it's done around here, because it's always been that way". It doesn't matter if the wider world knows that's maybe not the best way to do it, the historical groupthink wins out.

2018 comment: I've encountered that several times just in the last week, with people asking whether they still need to create a separate filegroup for nonclustered indexes, as their colleague says they do.

We see this with new clients over and over, and it can be really difficult to educate the client team as to different (usually better) ways of approaching a problem. Sometimes it's just one person who stymies all thoughts of innovation or change in the team, and sometimes it's the collective groupthink that does it.

In fact, there was a client (years ago), that Kimberly actually "fired" because of this. One meeting was enough but she told her contact ("the management") that she'd try again and meet again for a second session. In her post-mortem with management, her main quote was – your problems are more than technical (specifically person "X") and until that's resolved, I can't help you.

Here are some simple examples of ingrained behavior we've seen and corrected:

- Always using *NOLOCK* to work around blocking and deadlocking problems (and then it's used everywhere, and no one is even warned (in reports, etc.) about the negatives of dirty reads)
- Only ever rebuilding an index to remove fragmentation, without considering fill factors or trying to reorganize rather than rebuild
- Always using a heap and never using a clustered index because "heaps are faster"
- When creating a primary key, always using a clustered index as that's the default so "Microsoft must know that's always the right way to do it"
- For each column that's used in a *WHERE* clause, creating a single-column nonclustered index
- Always using *OPTION (RECOMPILE)* to prevent plans being put into the plan cache and taking precious memory away from the buffer pool
- Always using *sp_executesql* to put a plan into cache for ad hoc statements so that their plans get reused
- Always creating one file per processor core for *tempdb*, because that was the Microsoft guidance for SQL Server 2000

As I hope you can see, these are all very blinkered (and in some cases downright wrong) views on a variety of SQL Server topics. There are cases when some of these are the right thing to do, but not ALL the time just because 'that's the way we do it here'.

Does this remind you of anywhere? Or of anyone?

It's a dangerous way for a team to work because it can lead to all kinds of problems (such as poor performance, failed disaster recovery, good people leaving the team, business failure) that can be extremely hard to fix unless someone takes a stand or someone from outside the team helps to break the impasse. And even then, it's a delicate process to provide education and slowly change the thinking of the team, or of that one person who dominates the team's thinking.

Call to action: Consider the environment you work in and whether the situation described above is how your team operates. Are you the person to take the stand to try to change the groupthink? Do you do it or do you move on? Or are you the person who's unreasonably dominating the team's thinking? Do you allow yourself to be changed or continue in your intransigence? It's a difficult situation to recognize you're in, whatever your role is, and a hard choice to decide what to do about it.

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.

Performance Best Practices Guide for vSphere 6.7

The most current version of VMware vSphere is version 6.7, which was released on April 18, 2018. If you are a VMware shop, you should be aware of this release, and are hopefully planning on moving to it in the near future. This is especially true if you are on a much older release, such as 5.5 or older.

To help make the case for an upgrade and to help make sure you have everything configured correctly, VMware has a released an updated version of their performance best practices guide, Performance Best Practices Guide for vSphere 6.7.

Another handy document is "What's New in Performance – VMware vSphere 6.7", which is a much shorter document that outlines the key performance-related improvements in this release.

These include CPU scheduler enhancements, persistent memory support, and performance improvements to support Microsoft virtualization-based security (VBS).

Having persistent memory support is going to be extremely valuable for SQL Server workloads. SQL Server 2016 and newer support the persisted log buffers feature that lets you create a

second, 20MB transaction log, to cache the tail of the transaction log on a DAX storage volume that you create on persistent memory.

Intel Offering Xeon Discounts to Battle AMD EPYC

Typically, Intel only offers discounted server CPU prices to large hyper-scalers like Microsoft, Amazon, and Google. Word has been leaking out that Intel has been offering discounts to smaller companies, at quantity levels well below 1,000 CPUs.

Systems based on AMD EPYC 7000 series processors are less expensive than similar Intel-based systems, and they actually offer some technical advantages for some SQL Server and virtualization workloads. They have up to 64 logical processors per socket, up to 2TB of RAM per socket, and 128 PCIe 3.0 lanes per socket. Their single-threaded performance is within 10-15% of Intel performance in many cases.

As I have been saying for months, you should consider an AMD-based server, especially if you have a DW-type workload. The next generation AMD EYPC "Rome" processors (due in mid-2019) will be even more competitive.

#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.

We're ramping up work with Microsoft clients on upgrades off of 2008/2008 R2, 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 one-a-day upgrade blog post series

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

- Paul: *New VLF status value*
- Paul: *New live online training class in October: Fixing Slow Queries, Inefficient Code, and Caching/Statistics Problems*
- Paul: *New live online training class in October: Transactions, Locking, Blocking, Isolation, and Versioning*

I hope you find these useful and interesting!

Video Demo

The video demo this time is from Erin's Pluralsight course *SQL Server: Analyzing Query Performance for Developers* and shows you how to find trace flag information from query plans.

The video is about 7 minutes long and you can get it in MP4 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. Just FYI, our first quarter events will all be online and in the 2nd quarter we'll have a variety of in-person events back in Chicago. We hope to see you there!

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 **)**

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