(December 4th, 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 Las Vegas, where I'm currently on stage at our SQLintersection conference presenting a session on using wait statistics, my favorite subject, and Kimberly's on stage presenting about Indexes – one of her favorites! We've already had two days of workshops and now we have three days of sessions followed by a post-conference workshop day on Friday. DEVintersection / SQLintersection runs twice a year – usually with one in Orlando and the second in Las Vegas; we focus on development, architecture, migrations, and best practices so split your team up across events to divide and conquer; we hope you can join us in 2019!

SQLskills News

Live, ONLINE classes: we've announced six online classes running in the first quarter of 2019! It's the perfect way to keep learning, keep your systems moving forward, and stay motivated! And our format is receiving rave reviews: each class is held over three, *slightly-longer-than-half* days with a combination of lecture/demo and open Q&A. It's where you can really get your questions answered! Our online course line-up is as follows and now includes a **new class on columnstore indexes** taught by Jonathan:

- **IEQS**: Solving Common Performance Problems with Query Store
  - January 15-17
- **IEPUM2017**: Planning and Implementing an Upgrade/Migration to SQL Server 2017
  - January 29-31
- **IEQUERY**: Fixing Slow Queries, Inefficient Code, and Caching/Statistics Problems
  - February 12-14
- **IETLB**: Transactions, Locking, Blocking, Isolation, and Versioning
  - February 26-28
- **IEVLT**: Very Large Tables: Optimizing Performance and Availability through Partitioning
  - March 12-14
- **IECS**: Columnstore Indexes
  - March 26-28 **** NEW ****

These classes will be delivered live from 10am-3pm PST, Tuesday-Thursday for three consecutive days. In total you'll receive roughly 12-13 hours of content including open Q&As, which is similar to two, full, workshop days without leaving the comfort of your home/office! And by dedicating only 3 half-days of your time you still have time to be productive to your company by getting regular work done. Finally, you'll also receive **lifetime** access to the recordings – for reviews and refreshers, you get amazing ROI!

The classes are priced at US$699 each (US$599 for past attendees) and we're offering a combo package of any three for US$1,749 or all six for US$3,249. You can find all the logistical, registration, and curriculum details by drilling down from the class schedule page here.

Last, but not least – we've just announced our LIVE, IN-PERSON classes: we're back in Chicago in April/May 2019 and our classes are open for registration, including:

- IEPTO1: Performance Tuning and Optimization – Part 1
- IEPTO2: Performance Tuning and Optimization – Part 2
- IECAG: Clustering and Availability Groups
- IE0: Junior/Accidental DBA
- IEUpgrade: Upgrading and New Features
- IEPML: Practical Machine Learning
- IEAzure: Azure SQL Database, Azure VMs, and Azure Managed Instance

We're also adding a new IEPowerBI class – details on that coming soon.

All of the in-person classes have discounts (ranging from $100-200) for registering in 2018, and you can find all the logistical, registration, and curriculum details by drilling down from the class schedule page here.

Finally, even if you can't join us in person, I've put out a call for **2019 remote user group sessions** and we've done more than 75 this year already! If you'd like one of us to present for your user group, check out my blog post here.

Book Review

I've read two very good books since the last newsletter.

First is Sebastian Barry's *The Whereabouts of Eneas McNulty*. It follows the life of McNulty, from his childhood in early 1900s Sligo (a northern town in Ireland), through WWI, and back to a stint in the Royal Irish Constabulary. For this work he gets a death sentence pronounced on him by the local rebels, so he flees, ending up eventually in WWII, through Nigeria, and ultimately to the Isle of Dogs (near London). Each time he returns to Sligo, the sentence is renewed, and he flees again - denied the life he wants. Great writing and highly recommended!

The second book is Shaun Bythell's *The Diary of a Bookseller*. It's a year in Bythell's life owning a bookshop in Wigtown (a small town in SW Scotland that re-invented and re-invigorated itself in the late 1990s as a book town) – a place and area I've been to many times. It gives day-to-day accounts of interesting customers, book purchases, and general goings on. It's a charming book, and of course excellent for me as a confirmed bibliophile. 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 an email last week from a person who told me they were very confused about how differential backups work in SQL Server. They had recently moved from ORACLE where they performed differential backups every day using the Oracle Recovery Manager and then was able to use them and a set of logs to recover when needed.

After starting a similar process on a SQL Server instance, they noticed that the differential backups were steadily increasing in size every day, even though the volume of inserted data each day was roughly the same.

I explained that SQL Server does not have incremental data backups like Oracle does. Using SQL Server terminology, Oracle's differential backups contain everything that's changed since the last full or differential backup (i.e. truly incremental), whereas a SQL Server differential backup is everything that's changed since the last full backup (i.e. cumulative, not incremental).

The only incremental backups that SQL Server has are log backups.

Bottom line: There are some pretty fundamental differences between how SQL Server does things compared to other RDBMSs, and backups are one of the most important to understand.

Paul's Ponderings

Last week I was asked why sometimes it's possible to stop a *DBCC CHECKDB* and sometimes it's impossible to do so. This behavior is expected, but not intuitive at all and is to do with how *DBCC CHECKDB* works.

When *DBCC CHECKDB* starts, it usually creates a database snapshot under the covers. A database snapshot is required to provide *DBCC CHECKDB* with a transactionally-consistent, unchanging view of the database, so that *DBCC CHECKDB* knows it is checking the consistency of a static database that should not have corruptions.

A database snapshot is created by issuing a checkpoint on the database, creating the empty database snapshot, and then using the database's transaction log to run crash recovery on the database, but into the database snapshot. In other words, it rolls back any active transactions into

the database snapshot, without actually affecting the real database at all. The database snapshot thus becomes transactionally-consistent.

The time it takes to run crash recovery during the database snapshot creation is proportional to the amount and length of the uncommitted transaction that exist in the database when the database snapshot creation begins. If there is a long-running transaction, then that may take a long time to roll back. This means that the database snapshot creation will take a long time and *DBCC CHECKDB* will also take longer.

In the extreme case when the database snapshot creation is much, much longer than normal, and you decide to kill the *DBCC CHECKDB* process, nothing will happen, and you will have to wait for the database snapshot crash recovery to complete before the process will respond to the kill signal.

The only alternative in this case is to restart the SQL Server instance, which will remove the hidden database snapshot. Note that this does not work in the case of real crash recovery of regular databases – crash recovery in those cases will continue after an instance restart.

So why does the crash recovery process not respond to the 'kill' request?

A thread cannot be killed as such. It can only be signaled that it has been asked to terminate, and it will only notice the termination request when it runs some scheduling code, or makes a call into SQLOS to see if it's exhausted its thread quantum (see here for more details). This is why if you have a non-yielding scheduler (you'll see message 17883 in your error log), you can't kill the non-yielding thread, because it's not making that call into SQLOS and never see that it's been requested to terminate.

Now you'd think that a thread doing long-running crash recovery would have to run some scheduling code at some point, and of course it does, but crash recovery is a special case where thread termination is not permitted. This is because it is not possible to interrupt crash recovery, and there is no distinction in the crash recovery code in SQL Server between *real* crash recovery after an unexpected shutdown and a crash recovery for a database snapshot.

So there you have 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.

Microsoft Releases SSMS 17.9.1

On November 21, 2018, Microsoft released SQL Server Management Studio 17.9.1, which is a small update to SSMS 17.9 with three bug fixes that probably won't affect most people. The following fix is probably the most relevant:

"Fixed a long outstanding issue where restore plan would fail to find a restore plan, or would generate an inefficient restore plan under certain conditions."

I would argue that you are going to be better off just using straight T-SQL for your database restores rather than depending on the SSMS GUI, but if you are going to use the GUI, it is better if it works as designed. You should try to make a continued effort to keep your SSMS version up to date, since newer versions usually work better and they also have new features.

Speaking of new features, SSMS 18.0 Public Preview 5 was released on November 15, 2018. If you want to work with SQL Server 2019 CTP 2.1, you will want SSMS 18.0 in order to get full support for all of the new features.

If you have SSMS 18.0 Public Preview 4 or earlier installed, you will need to uninstall it before you can install SSMS 18.0 Public Preview 5. SSMS 18.0 can be installed side-by-side with SSMS 17.X, and it will also install and run on Windows 7 SP1.

Don't Buy a Gaming PC from Walmart.com

Steve Burke from Gamer's Nexus actually bought a "high-end" ($2099.00) DTW3 gaming PC from Walmart.com, from their new "OVERPOWERED" line of machines. As it turned out, Walmart shipped him a lower-end DTW1 model from the same line, which was probably just a shipping mistake.

More concerning are the terrible component choices, from a very low-end motherboard, power supply, and case, along with terrible build quality for a system that is clearly overpriced for the quality of the components. I hate to think of well-meaning parents or grandparents buying a machine like this from Walmart as a present for a loved one.

I completely understand that many people either can't or don't want to build their own machine from parts, but if you want to buy a pre-built machine, there are many much better choices of system integrators. One good example is iBuyPower. You can get a much better machine for less money from a good system integrator rather than Walmart.com. You can watch the complete video here.

BTW, here is a review of a terrible video of how NOT to build a PC that was published by The Verge. The original video was so full of inaccurate information, that it was removed by the publisher.

Samsung Introduces First Consumer SATA QLC SSD

Samsung is releasing the Samsung 860 QVO SSD, which is a quad-level cell (QLC) SATA SSD that will be available in 1TB, 2TB, and 4TB capacities. QLC NAND flash is a less expensive type of NAND flash compared to triple-level cell (TLC) or multi-level cell (MLC) that is found in most consumer-level flash storage.

Unfortunately, as you add more bits per cell, both write endurance and write performance decrease quite rapidly. Both QLC and TLC storage devices usually use some single-level cell (SLC) as a write cache to try to compensate for this issue. This works pretty well for common desktop usage.

Theoretically, QLC NAND flash should be about 25% less expensive than TLC NAND flash, which will make it more feasible cost-wise as a replacement for magnetic hard drives. Early yields for QLD NAND flash are reported to be low (which increases their cost), so we are not seeing that 25% cost reduction just yet.

My advice right now is to stick with TLC NAND devices (such as the Samsung 860 EVO) that give much better performance and endurance for nearly the same cost. There is actually very little real-world difference between the higher-end MLC-NAND based Samsung 860 PRO and the Samsung 860 EVO line, which makes the 860 EVO a smart choice if you want to save some money.

Anandtech has a good review here.

#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.

Following on from the *Curious Case* above, the theme for the TBT this time is backups:

- My Pluralsight course: *SQL Server: Understanding and Performing Backups*
- My TechNet Magazine article from 2009: *Understanding SQL Server Backups*
- My SQL Server Magazine article from 2011: *Advanced Backup and Restore Options*
- My TechNet Magazine article from 2009: *Recovering from Disasters Using Backups*
- Backup blog posts from our accidental DBA series
- Blog post: *New script: How much of the database has changed since the last full backup?*
- Blog post: *Importance of having the right backups*
- Blog post: *Importance of validating backups*
- My blog post category on *Backup/Restore*

I hope you find these useful and interesting!

Video Demo

In this Insider video, Tim shows you a new preview tool for your Azure SQL Databases called Data Discovery and Classification. This tool allows you to discover and classify sensitive data which can be very helpful with meeting data privacy standards and compliance requirements. It also helps with various security scenarios such as monitoring, auditing, and alerting on suspicious access to sensitive data. The practice of classifying your data also helps to control access and harden the security of your databases that hold sensitive data.

This tool is part of the Advanced Threat Protection offering in the Azure Portal. Security is a big focus for the Azure Cloud and the tools being built can help enormously with keeping your data assets secure.

The video is about 6 minutes long and you can get it in WMV format here.

Enjoy!

Upcoming SQLskills Events

Our first set of 2019 live, in-person events has been announced for Chicago in April/May and we've also just added our full 2019 Q1 lineup of live, online classes.

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:

- **IEQS**: Immersion Event on Solving Common Performance Problems with Query Store
  - January 15-17
- **IEPUM2017**: Immersion Event on Planning and Implementing an Upgrade/Migration to SQL Server 2017

    - January 29-31
- **IEQUERY**: Immersion Event on Fixing Slow Queries, Inefficient Code, and Caching/Statistics Problems
    - February 12-14
- **IETLB**: Immersion Event on Transactions, Locking, Blocking, Isolation, and Versioning
    - February 26-28
- **IEVLT**: Immersion Event on Very Large Tables: Optimizing Performance and Availability through Partitioning
    - March 12-14
- **IECS**: Immersion Event on Columnstore Indexes
    - March 26-28 **** NEW ****

LIVE, IN-PERSON Immersion Events:

Chicago, IL, April/May 2019

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
    - April 29-May 3
- **IECAG**: Immersion Event on Clustering and Availability Groups
    - April 29-30
- **IEPowerBI** – details coming soon!
    - April 29-30
- **IE0**: Immersion Event for the Junior/Accidental DBA
    - May 1-3
- **IEUpgrade**: Immersion Event on Upgrading SQL Server
    - May 1-3
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
    - May 6-10
- **IEPML**: Immersion Event on Practical Machine Learning
    - May 6-10
- **IEAzure**: Immersion Event on Azure SQL Database, Azure VMs, and Azure Managed Instance
    - May 6-9

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