(July 9th, 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,

We have **three brand-new live, online classes** that are filling up steadily – see below!

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! ***Important Note:*** *If you're even thinking about joining us in London, make sure to get your hotel booked by July 16 as our discounted rate expires and our extended block (we already sold out our room block once),* ***will be released****; you might not be able to stay at the Marriott Kensington (where the event is being held).*

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. Click on the class links above for all the details.

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 second-half 2018 remote user group sessions and we've done 59 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 James R. Clapper's *Facts and Fears: Hard Truths from a Life in Intelligence*. Clapper was the Director of National Intelligence from 2010 to 2017, and is a retired Air Force lieutenant-general who spent his entire career in intelligence. The book is excellent. The first 100 or so pages details Clapper's military career and then it gets into his experiences with and views on a variety of well-known episodes such as Benghazi, Snowden, and the Abbottabad raid that took out Osama bin Laden. The final 100 pages are about the 2016 election, and specifically the Russian interference in the election.

Clapper's views are balanced, professional, and very insightful, as one would expect from someone who was an integral part of (and eventual leader of) the U.S. intelligence community for 55 years, and he writes with a sense of humor too. He holds some of the media, some of Congress, and the current president in disdain for their willful disregard of facts, promulgation of falsehoods, and unwillingness to understand the capabilities and limitations of intelligence gathering and analysis, and considers a variety of moral and ethical questions on facets of intelligence. I found the book hugely interesting and it was obvious that Clapper wanted to write the straight up, unpartisan, and unembellished truth of what the intelligence services did and knew about all these topics, including mistakes that he and they made. Highly recommended!

PS In the past when I've read and reviewed a book that some people disagree with, I've sometimes been accused of pushing a political agenda in this newsletter through my book reviews. I find that accusation highly distasteful. I read what I read because I'm interested in a wide variety of subjects and viewpoints, and I present reviews here only of books I think are worth recommending for the benefit of a wide audience – nothing to do with politics.

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 was going back-and-forth with a prior student last week who was trying to work out how much data had changed in a SQL Server 2014 database since the most recent full backup. He had a really complicated set of code which involved using the *fn_dblog* function to analyze the transaction log. It would look through the log records and build a map of which pages had been

changed and then convert that to a proportion of the database. The problem was that the code was very slow (as *fn_dblog* isn't a very efficient log reader).

I asked what the root problem was, and the student explained that they wanted to programmatically make a decision as to when the next differential backups would be higher than some percentage of the size of a full backup, and so it would be more efficient to just perform a new full backup.

I explained that I'd written and published a script to do just that back in 2008, and it works perfectly on all database versions (see *New script: How much of the database has changed since the last full backup?*) My code uses *DBCC PAGE* to look through the differential bitmaps (one bit per changed extent since the most recent full backup) and spits out a percentage of the database that's changed.

The student was very happy, and started using that code.

Incidentally, in SQL Server 2017, there's a new field in the *sys.dm_db_file_space_usage* DMV called *modified_extent_page_count*, which does exactly the same thing without having to run any code at all.

Bottom line: You might be surprised at what scripts already exist and don't require complex code to do what you want. When you're considering automation, Google/Bing it as someone else might have already laid the foundation for you. More on this below…

Paul's Ponderings

Along the same lines, something I've seen repeatedly since I started working with SQL Server (on the SQL Server team back in February 1999) is people reinventing the wheel. And I don't just mean with T-SQL scripts, I also saw it over and over inside Microsoft, where developers would want to write their own version of code to do function X rather than re-using someone else's code that already does function X.

Writing code, in whatever language, to do something that someone else has already coded is usually a big waste of time, unless the only example you can find is poorly written or you really want to understand how something works, and writing code to use it would provide that understanding.

A prime example of this, that I see constantly, is writing a query involving DMVs. As you probably already know, many of the DMVs don't provide really useful (or actionable) information unless the results are manipulated in some way or joined with another DMV or catalog view.

An example would be one of my favorite DMVs: *sys.dm_io_virtual_file_stats*. On its own, it gives the cumulative read and write latencies for all I/Os done for all files in the instance, and the files are identified by database ID, file ID, and a hexadecimal file handle.

This information isn't really useful in that form. By doing some math on the output, basically converting the cumulative latencies into per-read/write latencies, and joining with *sys.master_files* to get file names and paths, the output becomes readable and easily understood.

Did I write the code to do that? No, I used some code that Jimmy May wrote a long time ago and whenever I show it when teaching, I give credit to Jimmy for it.

There's no shame in using or adapting code that someone else wrote and published. People in the SQL Server community publish code on blogs specifically so you can use it to save yourself time and get useful information from SQL Server.

If I'm ever going to write a new DMV-based query, I always check first to see if someone has already done the work – especially if I know it's going to involve joining with a few other DMVs and system catalogs and I'm going to have to play around to figure out which columns to join on.

I don't like reinventing the wheel – I'm too busy to spend my time doing that, and I think you are too. Here are a few links to scripts you may find helpful:

- Glenn has a whole series of useful queries that he publishes as his Diagnostic Information Queries (a toolkit) – see here for the latest version
- Kimberly's super-duper *sp_helpindex* rewrite
- Kimberly's script to identify duplicate indexes (with background reading)
- Jonathan's scripts to identify problems using analysis of the plan cache
- Erin's collection of scripts for producing a performance baseline
- Some of my more useful scripts:
    - Wait statistics main script and lots of links to other scripts
    - I/O latencies
    - How much has changed since the last full backup
    - Is a database really in the full recovery mode
    - All open transactions for a database
    - Wasted space in the buffer pool
    - Index counts for all tables
    - Potential space savings from changing a cluster key
    - Identifying queries with SOS_SCHEDULER_YIELD waits

And I remember the very first public SQL Server script I published, which was Example E in the Books Online entry for *DBCC SHOWCONTIG*, which I'd just rewritten for SQL Server 2000,

and the script made use of the new *DBCC INDEXDEFRAG* I'd written. It's still there in the official docs. Nowadays I'd always recommend someone uses (or uses code from) Ola Hallengren's excellent Maintenance Solution (yes, other people have published scripts to do similar, but Ola's are the gold standard IMHO so I won't recommend anything else).

Call to action: Any time you're trying to figure out a script to do X or Y or use DMV Z, if you just want the job done and you're not trying to learn the ins-and-outs of the DMV or how SQL Server works, do a quick Google or Bing search to see if someone already written a script, and then if you can't find anything ask on Twitter using the #sqlhelp tag and if a script exists, someone will point you to it. Don't reinvent the wheel if you really don't need to.

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 CEO Brian Krzanich Resigns

On June 21, 2018, Intel CEO Brian Krzanich resigned after a "past consensual relationship" with a subordinate employee was disclosed. This is the public reason for Krzanich's exit, but there has been a lot of public speculation that the actual reason was the continuing severe problems with Intel's 10nm manufacturing process.

Intel is several years behind their original schedule for moving from a 14nm to a 10nm manufacturing process, which has also delayed the rollout of new processor microarchitectures that depend on 10nm manufacturing.

This multi-year stumble by Intel has opened a big window for AMD to get seriously competitive in the server market. I've seen predictions that AMD will capture 20-25% of the server market in the next two-three years, especially after the second generation, 7nm AMD EPYC "Rome" server processors are introduced in late 2018 to early 2019.

Trace Flag 3427 for SQL Server 2016

Microsoft's Pam Lahoud has written a good blog post about some recent changes and improvements to tempdb behavior for SQL Server 2016 and SQL Server 2017. This is a very useful and detailed post that you should definitely read.

The short takeaway from the post is that it is important that you are on the latest CU, have multiple tempdb data files that are the same size, having larger tempdb data files has new benefits because of PFS page round-robin activity, and, you should enable global TF 3427 if you are on SQL Server 2016.

Building a Workstation from Parts

I typically prefer to use a desktop workstation for my main machine since I can get more compute, memory, storage capacity and storage performance compared to a laptop. You can also get more performance and capacity in a desktop workstation for less money than a high-end workstation laptop.

I also like to build my desktop workstations from parts rather than buying a pre-built system. Historically, you could save some money by building your own system (plus you could pick the actual components you wanted). Saving money has been more difficult over the past 12-18 months, since we have seen increasing prices on RAM, discrete video cards, and flash storage. The system vendors get better pricing on these components, so their system prices have been quite competitive.

Thankfully, this trend is starting to reverse for discrete video cards because GPU-based crypto mining is less profitable than it used to. Flash NAND shortages are not as much of an issue now, so flash storage prices are declining again. Unfortunately, DRAM prices for DDR4 RAM are still quite a bit higher than they were two years ago. Still, the overall cost of building a desktop workstation is less than it was a few months ago, and it should continue to decline in the coming months.

#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 #TBT this time is around stored procedures, so here are some resources for you:

- Kimberly's post: *SQLskills SQL101: Stored Procedures*
- Kimberly's course: *SQL Server: Optimizing Stored Procedure Performance*
- Kimberly's course: *SQL Server: Optimizing Stored Procedure Performance – Part 2*
- Kimberly's *Optimizing Procedural Code* blog post category

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

- Paul: *What is the FCB_REPLICA_SYNC spinlock?*
- Paul: *Calling all user group leaders! We want to present for you again in 2018!*
- Glenn: *SQL Server Diagnostic Information Queries for July 2018*

I hope you find these useful and interesting!

Video Demo

One of the benefits of teaching is the different scenarios we hear about from attendees. In the video this time, Erin looks at a case of an *UPDATE* statement having multiple plans in Query Store. You've seen *SELECT* statements with multiple plans, but have you seen that for a data modification? Check out the video to see how she tracked down the cause, and what you may need to watch out for in your environment.

The video is about 8.5 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.

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 **)**

LIVE, IN-PERSON Immersion Events:

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