(January 2nd, 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.

Quick Tips for our Insider friends!

Hey Insiders,

Happy New Year!

This newsletter is coming to you from Redmond, where we're home for a couple of months, working hard on new content and enjoying the crisp, clear weather. Back in November I mentioned our lovely old dog Coco had passed away, and the house was feeling a bit empty without her, so on December 29th we went to Seattle Area Feline Rescue and adopted three cats! If you're interested in seeing their adventures, follow them on Instagram as RandalsRescues.

And as a regular tradition I published my traditional end-of-year blog posts:

- 2017: the year in books
- 2017 review: the year by the numbers

Start your New Year Off with some training – from the comfort of your own location! We're doing our first live, online Immersion Event NEXT WEEK (January 9-11). It's important content around large tables – in both OLTP and relational-data warehouses! Kimberly is presenting her widely-requested course *IEVLT: Immersion Event on Very Large Tables: Optimizing Performance and Availability through Partitioning*. This will be delivered live via WebEx and is limited to 50 attendees. It's priced at only US$595. **Move fast to claim your seat!** See here for all the details. And, for all of our future, live, online events YOU – our newsletter subscribers – will always get first notice!

Note: you can get all the prior Insider newsletters here.

SQLskills News

Tim's latest Pluralsight course has been published! It's called ***SQL Server: Implementing a Hybrid Cloud Solution in Azure***. Check out the details here.

The first batch of 2018 US classes are open for registration! We'll continue offering a mix of training options as we still feel there's value in each option. We're planning to offer additional online courses but we feel there's still value in networking and in-person courses where you can truly immerse yourself and disconnect from the office, focusing on just learning and training for multiple days… And, of course, conferences like our SQLintersection, where you get breadth and insight into all sorts of different topics and areas that you want to learn deeper. Additionally,

and for now, our online courses are going to be focused on content that's largely different than our in-person courses.

In 2018, we're offering our usual Immersion Events on Performance Tuning (IEPTO1 and IEPTO2) and for the Accidental DBA (IE0), as well as PowerShell, Azure, Clustering and Availability Groups, BI strategies, BI security, and Practical Machine Learning. See here for our 2018 Immersion Event class schedule.

And we're also coming back to Europe in 2018! We're bringing four of our Immersion Events to London in September: IEPTO1 and IEPTO2, plus our new classes: **IEAzure** (on Azure and Azure VMs) and **IECAG** (on clustering and availability groups). See here for details.

Also for our European friends, Erin will be presenting a pre-con workshop at SQLBits in the UK in February – see here for details and check soon as it's almost sold out.

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 40 for 2018 already! If you'd like one of us to present for your user group, check out my blog post here.

Book Review

The most recent book I've read is Garrett M. Graff's *Raven Rock: The Story of the U.S. Government's Secret Plan to Save Itself--While the Rest of Us Die*. This is an extremely interesting book that explains the creation and evolution of the U.S. government's COG (Continuance Of Government) plans, to ensure that government continues in the event of a nuclear war or other major event. It goes into all kinds of details of bunkers, command and control networks, how COG worked when 9/11 happened, and lots more. Highly recommended! (And if you're into this stuff, check out *Command and Control: Nuclear Weapons, the Damascus Accident, and the Illusion of Safety* and *The Dead Hand: The Untold Story of the Cold War Arms Race and Its Dangerous Legacy* as well).

The Curious Case of…

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

Over the holidays I was answering an email question from a prior client who was having performance problems trying to use data stored as *FILESTREAM*. They had a number of symptoms:

- Inserting data into their table with non-NULL *FILESTREAM* data was slowing down
- They were seeing increasing *PREEMPTIVE_OS_CREATEFILE* waits
- Pulling back large *FILESTREAM* data values to their application was slow

The first step is to always check for that wait in my waits library. From there, they'd have had an explanation and been pointed to the *FILESTREAM* whitepaper, which luckily explains their problems!

The first two symptoms were because they hadn't disabled 8.3 name generation for the volume on which they'd placed the *FILESTREAM* data container. This meant that as the number of rows in their table kept increasing and the time to create the *FILESTREAM* file kept increasing. This is due to the inherent lack of scalability in the 8.3 name generation algorithm in NTFS.

The second symptom was because they were accessing the *FILESTREAM* data through T-SQL, rather than streaming it directly from NTFS. This meant the large *FILESTREAM* data values were being read into SQL Server's memory, and then sent to the application through TDS, which is *way* slower than accessing the *FILESTREAM* files directly.

Bottom line: If you're going to use the *FILESTREAM* data type, make sure that you're following all the best practices to get good performance, otherwise performance can be *really* bad. See the *FILESTREAM* whitepaper I wrote for lots of great tips and tricks!

Paul's Ponderings

(This editorial is adapted from my traditional first-newsletter-of-the-year editorial that I've been doing since 2014.)

At the start of every New Year it's always tempting to come up with a series of resolutions about things you're going to start doing, and it's tempting for me to write an editorial about planning to do new things. It's really easy to come up with ideas for cool things to *start* doing, which is why it's the easy route to take at the start of a new year.

But before you go too crazy with new ideas, consider a post mortem of your goals and plans from 2017. Look back at the year and see what didn't go well, and then figure out how to improve on those things. In my opinion, that's going to have a more positive effect on your life than continuing to do poorly on those things, and also try to add new things as well. (From 2014 to 2016, Kimberly commented *"the SQL development team could learn from this too…",* but with SQL Server 2016, they really did go back and fix a bunch of things that were broken or outdated.)

During 2015, one of the things I had reinforced through the year while I mentored folks in the SQL community is that *everyone* has something with which they need help or advice, and I continue to swap emails with a few mentees when they ask for advice (in 2016 and 2017). I strongly advise you to get a mentor outside your management chain (and preferably outside your company) – someone to whom you can bounce questions and one that has no vested interest – other than just helping you out.

And I'm not just talking about things in your work life; I'm also talking about things in your personal life. Here are some examples of things to think about and consider improving upon for 2018:

- Your work/life balance
- The support your immediate manager gave you
- Your interactions with some of your colleagues
- Your interest in some aspects of your job
- Your career/skills development
- Resistance to change
- Keeping on top of performance fires
- Keeping track of goals and to-do lists
- Saying 'no' when you're already overloaded
- Accepting a poor status quo at work instead of looking for a better job
- Finding time for disaster recovery planning/testing
- Finding time for all the little things that it's easy to procrastinate about
- Finding time for your hobbies
- Finding time and motivation to become/stay healthy
- Keeping in touch with old friends and distant family
- Finding time to reads some books or learn something new that's non-work
- Your work/life balance (again, because it's *so* important)

All of these can cause significant amounts of stress, so fixing them should be the first priority. Only once you have these items fixed, should you embark upon new and additional goals!

(Kimberly added: *And, don't get us wrong, this is just as hard for us as it is for you! And we both failed miserably at a couple of our goals too. But, we just have to look back and see why – and work to make it better from here. No reason to dwell, only to improve!*)

Call to action: Before you set cool, new goals for 2018 (whether at work or at home), do a post-mortem and work to fix the stuff from 2017 that was broken or didn't go so well. Not only will it feel good but it'll also reduce stress. And try to get a mentor!

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.

HWiNFO Supports Upcoming Intel Whiskey Lake and Ice Lake Processors

HWiNFO is a useful, free utility for getting some specific, comprehensive hardware information about your processor(s), memory, storage, and video card. It also includes pretty extensive real-time temperature and health monitoring of a system (but I would not use it as a replacement for the vendor-supplied tools on a server).

It has more detailed and specific information about processor instruction support, along with current clock speed monitoring of all of your processor cores.

Another use for this tool to get pretty reliable advance notice of some upcoming new hardware releases. By looking at the software change log for this tool, you will often get some confirmation that a new generation of processors or chipsets is getting pretty close to release, since the tool adds support for it. This sort of tool support makes it easier to understand what is happening with Intel's Process-Architecture-Optimization (PAO) development model.

Process-Architecture-Optimization is a development model introduced by Intel for their mainstream microprocessors in 2016 following the phase-out of their Tick-Tock model. The change is a result of the increase in cost and complexity of advancing lithography processes in the past decade. Under the new model the amount of time utilized for any given process technology is lengthened as Moore's Law increases in complexity with smaller nodes.

HWiNFO 5.70 adds support for the upcoming 14nm Intel Whiskey Lake platform and the subsequent Intel Ice Lake platform. SemiAccurate has their usual very negative take on Intel's development and release history here.

SQL Server 2014 SP2 CU9 Released on December 18, 2017

On December 18, 2017, Microsoft released SQL Server 2014 SP2 CU9, which is Build 12.05563.0. By my count, this CU has just seven public hotfixes, nearly all of which are for the SQL Engine or SQL performance.

Since SQL Server 2014 SP1 and earlier are no longer "supported service packs", there is no corresponding CU for the SP1 or RTM branches of SQL Server 2014. As always, I think it is a good idea to make an effort to stay current on cumulative updates, as does Microsoft.

Does Older Computer Hardware Actually have Lower Latency?

Dan Luu has done some research and actual experimentation with a high-speed camera (either 240fps or 1000fps) to measure the response latency of some quite ancient computer hardware and phones compared to much newer hardware.

His tests measure the latency between a keypress and the display of a character on a terminal, or in the case of mobile devices, scroll latency in the browser.

His testing shows much lower latency from something like an Apple IIe or a Ti 99/4a than from a modern desktop machine such as a PowerSpec G405. Most modern keyboards only scan their inputs at 100 Hz to 200 Hz, while an Apple IIe scanned its keyboard at 556 Hz. Another factor is the latency of most modern displays compared to old CRT displays. A high quality "gaming monitor" that supports G-Sync or FreeSync has comparable latency to old CRT monitors.

#TBT

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

The theme for the TBT this time is *FILESTREAM*, to match the earlier Curious Case:

- My whitepaper: *FILESTREAM Storage in SQL Server 2008*
- My blog posts:
  - *FILESTREAM directory structure*
  - *FILESTREAM directory structure – where do the GUIDs come from?*
  - *FILESTREAM garbage collection*
  - *FILESTREAM* blog category

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

- Paul: *2017 review: the year by the numbers*
- Paul: *2017: the year in books*
- Paul: *Live online training from Kimberly in early January!*
- Tim: *Microsoft Database Experimentation Assistant with SQL Server 2017*
- Glenn: *SQL101: Avoiding Mistakes on a Production Database Server*
- Glenn: *SQL Server 2014 SP2 CU9*

I hope you find these useful and interesting!

Video Demo

The demo video this time is from Glenn's most recent Pluralsight course: *SQL Server: Understanding, Configuring and Troubleshooting Log Shipping*. In the video, Glenn shows how simple it is to configure log shipping using Management Studio.

The video is five minutes long and you can get it in MOV format here.

No demo code this time.

Enjoy!

Upcoming SQLskills Events

We have lots of events coming up in 2018 – from our online Immersion Event in January to SQLintersection in March to our own in-person Immersion Events in the U.S. and London; all events are open for registration.

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

Online, January 2018

- **IEVLT**: Immersion Event on Very Large Tables: Optimizing Performance and Availability through Partitioning **(** NEW **)**
  - January 9-11

Orlando, FL, March 25-28, 2018

- SQLintersection co-located with the DEVintersection conferences (register with the SQLskills discount to get save $50 on registration)

Chicago, IL, April/May 2018

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - April 23-27
- **IE0**: Immersion Event for the Junior/Accidental DBA
  - April 23-25
- **IEUpgrade**: Immersion Event on Upgrading SQL Server
  - April 23-25
- **IECAG**: Immersion Event on Clustering and Availability Groups
  - April 26-27
- **IEAzure**: Immersion Event on Azure SQL Database and Azure VMs
  - April 26-27
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - April 30-May 4
- **IEBIStrat**: Immersion Event on Developing a BI and Analytics Strategy **(** NEW **)**
  - April 30-May 2

- **IEBISec**: Immersion Event on Securing Your BI Platform **(** NEW **)**
  - May 3-4
- **IEPS**: Immersion Event on PowerShell for SQL Server DBAs
  - April 30-May 2
- **IESSIS1**: Immersion Event on Learning SQL Server Integration Services
  - May 7-11
- **IEPML**: Immersion Event on Practical Machine Learning
  - May 7-11 **(** NEW **)**

Bellevue, WA, June 2018

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - June 18-22

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
- **IECAG**: Immersion Event on Clustering and Availability Groups
  - September 12-13
- **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