(October 23rd, 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 Redmond, where I'm wrapping up a new Pluralsight course from Glenn while helping proctor the first day of our brand-new, multi-instructor, online IEQUERY class (it'll be offered again next year).

And talking of 2019, **live, in-person classes for Spring 2019 are open for registration** – details below, including a **new IEAzure class expanded to four days!**

SQLskills News

We're going to see a lot of you in the coming weeks, and for the first time in four years, **our own awesome SQLintersection Fall 2018 show doesn't clash with the PASS Summit**. You still have time to join us in sunny Las Vegas as SQLintersection is 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, in-person classes: we have all of our usual classes open for registration for April/May 2019 in Chicago, 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 in the next few weeks.

All the 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

The most recent book I've read is an Icelandic saga (a story of a real historical episode from the 9th-to-11th century), *The Saga of Gunnlaug Serpent-Tongue* (anonymous). This one was composed in the 13th century, and tells a story of two warrior-poets competing for the hand in marriage of a woman, while illustrating customs like gift-giving and honor. Not something with wide appeal, but interesting for anyone into history and who's been to Iceland. A more detailed Wikipedia explanation of this saga is here.

So why am I telling you about this if I don't think it has wide appeal? It's one of the books from *Penguin Little Black Classics* set of 80 small books. From Amazon: "This spectacular box set of the 80 books in the Little Black Classics series showcases the many wonderful and varied writers in Penguin Black Classics. From India to Greece, Denmark to Iran, the United States to Britain, this assortment of books will transport readers back in time to the furthest corners of the globe. With a choice of fiction, poetry, essays and maxims, by the likes of Chekhov, Balzac, Ovid, Austen, Sappho and Dante, it won't be difficult to find a book to suit your mood. Little Black Classics celebrate the huge range and diversity of the Penguin Classics list - from drama to poetry, from fiction to history, with books taken from around the world and across numerous centuries."

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 question last week from someone who was trying to restore after a system crash and the final log backup was introducing corruption. They found this by doing the complete restore and running a *DBCC CHECKDB*, which showed corruption. So they did another restore, skipping the last log backup – no corruption – so they asked me if I could think of anything.

Apart from actual log/data corruption, there's only one thing I can think of that will cause the symptoms they reported. I asked them what that last log backup was, and they said it was a tail-of-the-log backup after they'd had the system crash. I then asked what recovery model they use and they said bulk-logged, all the time.

Aha! Here's what happened…

The tail-of-the-log backup contained a minimally-logged operation. When that happens, the log backup backs up the log generated since the previous log backup AND the data pages changed by the minimally-logged operation. If the data files aren't available, and the *NO_TRUNCATE* option is used to force the log backup, from SQL Server 2008 R2 onwards, even though the log backup *should* back up the data pages as well because there was a minimally-logged operation, it can't but the backup succeeds. When you restore that log backup, it creates a bunch of empty data pages – corruption.

I can argue the pros and cons of that behavior change, but that's what caused their problem.

Bottom line: If you have to perform a tail-of-the-log backup (see here for instructions), pay attention to the messages the backup produces. In the case above, it doesn't throw an error, but will have a line of output that says 'CONTINUE_AFTER_ERROR' was used, even though you didn't specify it. That backup will introduce corruption if you restore it.

Paul's Ponderings

*This year we've been involved in a *lot* of upgrades as people move off of 2005/2008/2008 R2 systems and up to 2016/2017 – some with our help, and some where we've been called in after the upgrade.*

For all the upgrades where we've helped clients prepare and execute the upgrade, they gone according to plan and there haven't been any big surprises in terms of performance/functionality regressions.

Unfortunately, for almost all the upgrades where someone came to us for help after they upgraded, there were big problems, sometimes requiring us to spend a lot of time figuring out the problems, and in several cases, helping the (now) clients go back to their old version.

I was going to write an editorial around this, but I saw that Jonathan had written one back in 2012, so I'm going to reuse that – enjoy!

Every time I watch the movie The Replacements I'm reminded about a topic I've intended to write about repeatedly but I never get around to actually doing it. In the movie, during a locker room meeting, the head coach asks the players what they are afraid of. After a number of players mention off-the-wall topics like spiders and bees, the quarterback says that he is afraid of quicksand. The explanation behind this in the movie has always been something that has resonated with me in my career because I can easily look back and see points in time where quicksand got the best of me. In this case quicksand is a metaphor, it describes a situation where

something has gone wrong that is out of your control and no matter how hard you fight, things just continue to go wrong until you are sinking, just like quicksand.

It is really easy to get trapped in quicksand in IT, and a lot of times it is the result of not having an alternate plan to fall back on should your original plan not work out. I can remember early on in my career as a DBA when I was performing one of my first upgrades from SQL Server 2000 to SQL Server 2005. I had a complete upgrade plan that had been tested repeatedly inside of our test environment using restored copies of the databases on multiple SQL Server 2000 installations that I had built, upgraded in place, rebuilt, restored the databases, upgraded in place again, and then repeated until I had a complete set of upgrade steps that were necessary and I had a rock-solid plan based on testing. This all seemed great until the night that I was actually performing the upgrade on the production server and it failed.

A failed upgrade on its own isn't really a disaster, these things happen, and at this point I had practiced installing and setting up a duplicate copy of this SQL Server repeatedly, except I did that during the day when I had access to the installer media and the server in our data center to be able to drop the CD in the CD-ROM drive. Since I hadn't planned on this upgrade failing, I never considered possibly pre-staging the necessary installation media where I could access it to rebuild the server. At this time I also didn't have direct access to our data center or to the installation media (which was locked in a cabinet in our CTO's office) so I had to coordinate with a member of the server team – first to get access to the installer disks and then to get access to the data center. Now it was time to hit the phone tree for the server admins and see if one of them would meet me at the office to bail me out of this problem I'd created.

Two and a half hours later my frantic voice mails were finally returned by a our Senior Server Admin, who promptly asked why I didn't copy the installer files from the test server I had been working with the last two weeks. In the ensuing panic that followed the upgrade failure I never even thought about alternate locations for the installer media, like the test server I had been using where we copied the installer files to the local disk to keep from having to ferry the CD back and forth to the CTO's office every time I reinstalled SQL Server 2000 on the machine to test the upgrade again. By only planning for the upgrade I had actually left myself wide open to falling into quicksand, and I've repeatedly seen smart people make stupid mistakes in my IT career while engrossed in the ensuing panic that occurs when they don't have an alternate plan for when something goes wrong.

This upgrade taught me a number of lessons about planning and preparing, not for the correct outcome but instead for the worst case scenarios if something goes wrong. Having something break at some point in your IT career is inevitable, a failure becoming a total disaster is not. A good example of this is disaster recovery strategy planning for your environment. A number of companies have thought about their disaster recovery requirements and have a remote data center that has servers, possibly a SAN that is replicated to, but at least copies of the data and

applications ready to run with minimized downtime if necessary in a major disaster at their primary location. However, few companies I've talked to have ever considered what they would do if both data centers were lost to a single event, a disaster that happened on September 11, 2001 when the World Trade Center collapsed since many business had a primary data center in Tower 1 and their DR data center was in Tower 2. If a scenario like this happened today would you be prepared to recover from it? Some questions to think about for this scenario include:

- Can you purchase new hardware and how long will it take to be delivered?
- Is it possible to have a lease agreement pre-established with guaranteed delivery SLAs if needed?
- How long does it take to configure bare-metal installations of your environment from the ground up?
- Where are your backups located?
- How long does it take to get access to them?
- How long to restore?
- If you need network access or network routing how are you going to accomplish this?

From these seven questions alone, you are likely thinking about a number of failure points that you've possibly never considered before, and these haven't even scratched the surface of all the problems that need to be addressed for business continuity planning.

Call to Action: Consider the potential failure points in your day-to-day activities and your operational environment that could easily turn into quicksand should something go wrong, and then formulate a plan ahead of time for how to deal with those potential failures. Then go back and consider what could fail in your alternative plan and how to deal with those as well.

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.

AMD Could Regain Significant CPU Market Share in Q4 2018

AMD continues to ride the momentum of the AMD Ryzen architecture in the server, desktop and mobile processor market segments. In recent months, AMD has scored a number of product wins, with more hardware vendors selling new model systems with AMD processors.

As Intel continues to struggle with their migration from 14nm to 10nm processors, we are also starting to see shortages of many current 14nm Intel processor models, including in the server space. As you might expect, this Intel product shortage has been accompanied by rising prices for many Intel desktop and mobile processor models.

While this Intel shortage/price increase has been happening, AMD has seen price decreases for their latest 12nm desktop Zen+ processors (such as the Ryzen 7 2700X).

Recent rumors have indicated the upcoming 7nm AMD Zen 2 CPU architecture is showing a 13% instructions-per-clock-cycle (IPC) lift over the current Zen+ processors for scientific tasks, at the same clock speed. If that information holds up, it would be pretty significant, especially with the expected clock speed increases we will see as AMD moves to 7nm with Zen 2. Both AMD desktop and server processors that use the 7nm Zen 2 architecture are expected to be released in 2019.

New Intel Desktop Processor Families

Intel formally announced its 9th generation Core mainstream desktop processors on October 8, 2018 at its Fall Launch Event in New York. So far, they have announced three members of this family.

- Core i9-9900K (8-cores/16-threads) for $488.00
- Core i7-9700K (8-cores/8-threads) for $374.00
- Core i5-9600K (6-cores/6-threads) for $262.00

The Core i9-9900K has a base clock speed of 3.60 GHz, a Turbo clock speed of 5.0 GHz, and 16MB of L3 cache. The Core i7-9700K has a base clock speed of 3.60 GHz, a Turbo clock speed of 4.9 GHz, and 12MB of L3 cache. Finally, the Core i7-9600K has a base clock speed of 3.70 GHz, a Turbo clock speed of 4.6 GHz, and 9MB of L3 cache. These processors are aimed at gaming and general desktop usage. They are competitors to AMD's Ryzen 2xxx Zen+ mainstream processors.

You may have noticed that Intel has dropped hyper-threading (HT) from the Core i7 line (which is something it has always had, and was one of the main differentiators over the Core i5 line in the past). The lack of HT means a loss of about 25-30% of your overall CPU capacity when you have the same number of physical cores.

These new processors have soldered thermal interface material (rather than thermal paste) which will help with heat dissipation, and let them run more cores at slightly higher (100-200MHz) clock speeds more often. They also have hardware-level mitigation for some of the Meltdown CPU exploits, which will result in better performance than software mitigations.

These new processors will all work in existing Intel 300 series chipsets (with an updated BIOS), but Intel is also introducing a new Z390 chipset, which is a very slight improvement over the previous Z370 chipset, primarily with native USB 3.1 Gen 2 support and built-in 802.11ac Wi-Fi support.

#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.

Parallelism and CXPACKET waits is the theme for this TBT:

- My post: *Knee-Jerk Wait Statistics: CXPACKET*
- My post: *More on CXPACKET Waits: Skewed Parallelism*
- My post: *CXCONSUMER wait type – history and what you need to know*
- The waits library entry for *CXPACKET*
- Jonathan's post: *Tuning 'cost threshold for parallelism' from the Plan Cache*
- Jonathan's post: *SQL 101: Parallelism Inhibitors – Scalar User Defined Functions*
- My post: *Updated sys.dm_os_waiting_tasks script to add query DOP*
- My post: *Who is overriding MAXDOP 1 on the instance?*

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

- Paul: *New class: Immersion Event on Azure SQL Database, Azure VMs, and Azure Managed Instance*
- Paul: *Calling all user group leaders! We want to present for you in 2019!*
- Paul: *Lazy log truncation or why VLFs might stay at status 2 after log clearing*
- Erin: *Baselines for SQL Server and Azure SQL Database*
- Erin: *Removing a database from a replica in an Availability Group*

I hope you find these useful and interesting!

Video Demo

In this insider video, Tim shows you how to duplicate/copy your Azure SQL Databases if you need to make a copy. He shows you two ways within the Azure Portal using restore services and the built-in copy option and then shows you how to use a new command for Azure SQL Database in SQL Server Management Studio using *CREATE DATABASE database_name AS COPY OF source_database*. Knowing how to make a copy of a database by means other than restoring from backup can help set you apart from your peers.

The video is about 3 minutes long and you can get it in WMV format here.

Enjoy!

Upcoming SQLskills Events

We still have events coming up in 2018 – from our ***new, live, online*** IEQUERY course to our own conference: SQLintersection; all of our Fall events are open for registration.

Our first set of 2019 live, in-person events has been announced for Chicago in April/May and we'll be adding a couple more in-person classes in a week or so. Just FYI, our first quarter events will all be live, online classes and we'll be announcing those around late 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, 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