

(April 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,

We're just back from an excellent SQLIntersection conference last week. So many attendees told us they prefer the smaller, more intimate show we produce, where they get to spend time with the speakers and really address their questions and concerns! Our next SQLIntersection is scheduled for the first week of December in Las Vegas – more details soon.

We have a new, live, online class coming in May! Erin will be presenting our new ***IEQS: Immersion Event on Query Store***. This will be delivered live via WebEx on May 22-24 (roughly 12-13 hours of content including open Q&As; similar to two full workshop days without leaving the comfort of your home/office!). It's priced at only US\$795; however, as a newsletter subscriber, you can use the discount code 'newsletterqs' when you register **before** 4/16/18 to save US\$100. **Move fast to claim your seat!** See [here](#) for all the details.

And you still have two weeks left to register for Kimberly's delayed, and highly-rated *IEVLT: Immersion Event on Very Large Tables: Optimizing Performance and Availability through Partitioning* course. See [here](#) for all the details, including the incredible feedback from the January class.

Note: you can get all the prior Insider newsletters [here](#).

SQLskills News

In-person US classes: 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, and Practical Machine Learning. See [here](#) for our 2018 Immersion Event class schedule. As noted above, **there will be no Fall in-person classes in the U.S. in 2018.**

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 and Azure VMs) and **IECAG** (on clustering and availability groups). See [here](#) for details.

Also, as for our Immersion Events; our Fall schedule is already packed with events and engagements. As a result, there will be **no October/November in-person classes in Chicago this year** so in the U.S. the only in-person classes will be in Chicago in April/May and Bellevue,

WA in June. This means that each of **our in-person classes will only be offered once in the U.S. in 2018** (with an extra IEPTO1 in June in Bellevue), and these classes are all filling up quickly – please don't wait to the last minute to register otherwise you'll be disappointed.

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 60 for 2018 already! If you'd like one of us to present for your user group, check out my blog post [here](#).

Book Review

Here are a couple of books I've read recently:

The first book is Virginia Woolf's [To the Lighthouse](#). This is the first of Woolf's novels I've read and it won't be the last. It follows two brief times (ten years apart) in the life of an English family in their vacation house in the Hebrides (islands off the NW coast of Scotland). The interesting thing about the book is that it's all about the thoughts that go through the family member's heads as they're interacting with each other - harsh, critical, and insightful.

The second book is Ken Follett's [The Man From St. Petersburg](#). I'm a big fan of Follett and I really enjoyed this book. It's about Great Britain and Russia colluding before the start of WWI and a Russian anarchist trying to assassinate the Russian envoy, with all kinds of historical and familial complications. It's nicely written for the period, and a page-turning drama.

Both are excellent and 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.

Recently I was asked about a restore issue that someone was having after a hardware failure. The failure meant that performing a full database backup failed, but the scheduled, hourly log backups worked, so their plan was to go back to a prior full backup and restore the sequence of log backups. They stated that no other log backups had been performed apart from the scheduled ones.

However, part way through the restore sequence, the first log backup after the most recent failed full backup attempt failed to restore with an LSN error. I was racking my brains trying to figure out what in the failed full backup attempt could have caused the next log backup to be bad in some way, as full backups don't change anything in the database to do with log backups. It made no sense to me, and would be a nasty SQL Server bug if that's what happened.

A little while later, we found out the answer. When the full backup attempt failed, a junior DBA had performed a log backup, appending the log backup to an existing backup file. When the

restore sequence was attempted, this log backup was not restored, and so the restore of the next 'real' log backup failed (and rightly so) with an LSN error, because of the missing log records in the log backup the junior DBA performed.

Once they realized that, the restore sequence worked perfectly. I hadn't thought of that as they'd been very specific about what log backups had been performed.

Bottom line: Be very careful about performing 'out-of-band' log backups as they become part of the restore sequence. If you lose it or don't know about it, your restore sequence stops at the prior log backup, which could be disastrous for your company. Additionally, make sure that your "restore scripts" are automated against the information in *msdb* rather than only an assumption that what was supposed to happen actually did! Finally, check out the new features in SSMS as they can generate the restore sequence from *msdb* OR from backup devices directly (this is a fantastic feature that was added to SSMS for SQL Server 2014 but few people know it exists).

Paul's Ponderings

Continuing the theme of lost backups, the editorial this time is an amalgamation of some of the narration from my Pluralsight course [SQL Server: Understanding and Performing Backups](#).

If at any time you want to take a one-off, out-of-band backup, you need to be very careful how you do it.

A differential backup is all the changed data since the last 'real' full backup. By 'real' I mean not specifying the copy only option for the full backup. If you perform a real full backup, that full backup becomes the basis for subsequent differential backups and might be required during a disaster recovery operation.

A log backup is the entire log generated since the last real log backup, where real means a log backup where you didn't specify the copy only option. If you perform a real log backup, that log backup becomes part of the log backup chain, and again, it may be required during a disaster recovery operation.

If you or someone performs an ad hoc, or one-off, real full backup or an ad hoc real log backup for some reason, and then that backup is lost or damaged, your ability to restore can be compromised (as in the *Curious Case* I explained above). Remember that full backup will reset the differential bitmaps, meaning all subsequent differential backups require that particular full backup as the start of a restore sequence. And if it's a log backup, it's backed-up a portion of the log than won't ever be backed up again so is needed as part of the log backup chain.

As an aside, a few other things can lead to the log backup chain being broken apart from a missing or damaged log backup:

- If you switch to the simple recovery model, that immediately breaks the log backup chain. When you switch back to the full or bulk-logged recovery model, you need to perform a full or differential backup to re-establish the basis for being able to perform a log backup. If you don't do that, SQL Server won't let you perform a log backup.
- If you create a database snapshot, then perform a revert from snapshot operation, basically pushing all changed pages from the snapshot back into the source database, that breaks the log backup chain. It essentially removes the transaction log and creates a new one with only two VLFs and a total size of 0.5MB. You'll need to perform a data backup before you can perform log backups. You might choose to do this as an alternative to a lengthy restore sequence, where the revert from snapshot will be a lot faster.
- If you need to rebuild the transaction log, for instance because it's been damaged or destroyed, or from using the emergency-mode repair option of *DBCC CHECKDB*, that breaks the log backup chain and has the same requirements for reestablishing it as above. You can learn more about recovering from log damage in my Pluralsight course on [Advanced Corruption Recovery Techniques](#).
- Finally, if you're running a version of SQL Server prior to SQL Server 2008, you could manually truncate the log, which would break the log backup chain. You can prevent people from doing this on these older versions by enabling trace flag 3231 as a startup trace flag.

To avoid having one-off backups affecting your differential backups or your log backup chain, you can use the *COPY_ONLY* option. For a full backup, this option makes the backup NOT clear the differential bitmaps and become the basis for subsequent differential backups. It's as if that full backup never happened.

And for a log backup, this option makes the backup NOT become part of the log backup chain. It's as if that log backup never happened.

Call to action: On any database where automated, regular backups are being performed, always use the *COPY_ONLY* option for ad hoc full and log backups unless you have a very good reason not to. Better to be safe than sorry.

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 Launches Speculative Execution Side Channel Bounty Program

Microsoft has [announced](#) a new bug bounty program that pays up to \$250,000.00 for newly discovered speculative execution side channel vulnerabilities that are reported to Microsoft between March 14, 2018 and December 31, 2018. This is how Microsoft describes the program:

“Through this program, individuals have the opportunity to submit novel speculative execution side channel vulnerabilities and mitigation bypasses that affect our latest Windows and cloud platforms. Under this program, qualified submissions are eligible for payment of up to \$250,000.00 USD. All bounties will be awarded at Microsoft's discretion.”

Intel [announced](#) a similar program back in February. I think these are both great developments, since it gives large financial incentives for security researchers to try to find and responsibly report these types of vulnerabilities.

If you want to learn more about Microsoft has recently written about these type of attacks, [this is a great post to get started with](#). They also have [updated guidance](#) on how to mitigate speculative execution side channel vulnerabilities for pretty much all Microsoft products and services.

The latest SQL Server specific guidance is at [Protect SQL Server from attacks on Spectre and Meltdown side-channel vulnerabilities](#).

Brief Power Outage in South Korea Worsens NAND Shortage

Samsung's NAND flash plant in Pyeongtaek suffered a 30-minute power outage on March 9th, 2018, which damaged 50,000 – 60,000 NAND wafers. This represents approximately 11% of Samsung's NAND production for the month of March, which also maps to approximately 3.5% of global production for the month.

This may seem relatively insignificant, but there have been on-going NAND shortages due to increased demand from smartphones, tablets, laptops, and even servers. If you have been watching pricing for NAND flash devices over the past year, the general trend for SSD pricing has been upward.

Upcoming Intel Cascade Lake Xeon Scalable Processors will have on-silicon Spectre Fixes

Intel has announced that their upcoming (2H 2018) 14nm Cascade Lake Xeon Scalable Processors will have on-silicon (rather than just firmware) fixes to protect against the Spectre variant 2 and variant 3 vulnerabilities. These processors will be the successor to the current Skylake-SP generation of Xeon Scalable processors, and they should work in existing generation servers, such as the Dell PowerEdge R740, with a BIOS update.

Unfortunately, Intel has not yet released very much other information about the Cascade Lake Xeon processor family, beyond the fact that it will support Intel Optane DIMMs. Hopefully, there will be more than the usual amount of minor architectural improvements, clock speed increases and faster memory support, due to the increased competition from the AMD EPYC server processors.

#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 theme for the TBT this time is log shipping. It's been around forever, is simple to configure, and has a bunch of uses. Here are some useful links:

- Glenn's Pluralsight course: [SQL Server: Understanding, Configuring and Troubleshooting Log Shipping](#)
- Whitepapers:
 - [Database Mirroring and Log Shipping Part I – Prescriptive Guidance](#)
 - [Proven SQL Server Architectures for High Availability and Disaster Recovery](#) (written by me)
 - [Database Mirroring and Log Shipping Working Together](#)
- Blog posts:
 - Paul: [SQLskills SQL101: Log shipping performance problems](#)
 - Glenn: [Upgrading SQL Server–Migrating with Log Shipping](#)

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

- Paul: [SQLskills SQL101: Is the recovery interval guaranteed?](#)
- Paul: [No Fall U.S. classes in 2018](#)
- Glenn: [SQL Server 2017 Cumulative Update 5](#)
- Glenn: [Recent SQL Server Software Updates](#)
- Tim: [Calling all user group leaders! I want to present on Managed Instances for you!](#)
- Jonathan: [Simplifying Availability Group Troubleshooting](#)

I hope you find these useful and interesting!

Video Demo

Microsoft just released Azure SQL Database Managed Instances in public preview on March 6th. In this insider video, Tim shows you how easy it is to back up on-premises SQL Server databases and migrate them to Azure SQL Database Managed Instances using Azure Blob storage. According to the documentation, you can restore databases going back as far as SQL Server 2005 (although those compatibility levels will be flipped to 2008).

The video is about 4.5 minutes long and you can get it:

- In WMV format [here](#).
- In MOV format [here](#).

And the demo code is [here](#).

Enjoy!

Upcoming SQLskills Events

We have lots of events coming up in 2018 – from our online IEVLT and IEQS courses to our own LIVE, in-person Immersion Events in both the U.S. and London; all events are open for registration. Every event has a different focus and different benefits – from deep-technical training in our online courses and in-person IEs to wide-ranging topics at SQLintersection where you can learn more effectively how to keep moving forward in both your database and your career! And, of course, one benefit all our in-person events provide is networking.

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, April/May 2018

- **IEVLT**: Immersion Event on Very Large Tables: Optimizing Performance and Availability through Partitioning (**Final planned delivery in 2018**)
 - April 17-19
- **IEQS**: Immersion Event on Query Store **** NEW ****
 - May 22-24

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
- **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