

(December 18th, 2017)

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,

As we're winding down for the year, we'd like to wish you all a very happy holiday season, whatever and however much you celebrate it!

BIG, BIG NEWS: We're doing our first live, online Immersion Event on January 9-11; start off the New Year learning some important content! Kimberly will be 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. For the next 48 hours, the only people who know about this course are newsletter recipients (we are not tweeting or blogging until Wednesday). Until then, **you can register for the discounted price of US\$495**. On Wednesday morning, we'll blog/tweet about the course and the price will increase to the regular price of US\$595. **Move fast to claim your seat and save US\$100!** See [here](#) for all the details. And, for all of our future, live, online events YOU – our newsletter subscribers – will get first notice and with a discounted price!

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 more. See [here](#) for our 2018 Immersion Event class schedule.

And, we've just added a new in-person Immersion Event: *IEPML: Immersion Event on Practical Machine Learning*. It's available as a full 5-day course, or as a 2-day introductory Part 1 and a 3-day intermediate Part 2. You can read a detailed curriculum [here](#) and all the class registration and logistical details are [here](#).

Remember: Class prices are going up a bit in 2018 so register in the next two weeks for 2017 prices/discounts!

We're coming back to Europe in 2018! We're bringing four of our Immersion Events to London next 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 – 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

As you all know, I'm seriously addicted to books and reading – a true bibliophile. I thought what better book to recommend as the last one of the year than a book about wonderful books! So, one of the most recent books I've read is the DK book [Remarkable Books: The World's Most Beautiful and Historic Works](#). Here's the Amazon blurb about it, that describes is better than I can:

A beautifully illustrated guide to more than 75 of the world's most celebrated, rare, and seminal books and handwritten manuscripts ever produced, with discussions of their purpose, features, and creators.

*From ancient masterpieces such as *The Art of War*, written on the leaves of bamboo, to the stunningly illustrated *Birds of America*, to Chairman Mao's *Little Red Book*, *Remarkable Books* delves into the stories behind the most incredible tomes ever produced, offering an insight into their wider social and cultural context, and is chronologically ordered to demonstrate the synergies between the growth in human knowledge and the bookmaking process. Alongside breathtaking images of the books and manuscripts themselves, close-up views draw out interesting features, which are discussed in greater detail, while biographies tell the lives of the people who produced them.*

This coffee table-worthy book is wrapped in a textured jacket with gold foil, making it a great gift for those with an interest in literature and art and design.

I'm familiar with many of the books covered, especially the illustrated manuscripts from the Middle Ages, as those are of particular interest to me. It really is a fantastic book itself, and I hugely recommend it!

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.

I was working with a client recently who ran into a problem with one database, using the simple recovery model, where the log was growing and would not clear. They checked the `log_reuse_wait_desc` field in `sys.databases` and it said `REPLICATION`.

They had previously set up the database as a replication subscriber but they'd since removed replication, and yet the log still would not clear when they ran a checkpoint (which is what attempts to clear the log when using the simple recovery model).

They had tried the usual things: running `sp_repldone`, running `sp_removedbreplication`, and even re-configuring and removing replication again. Nothing worked.

When they told me this, my response was to ask if they were using Change Data Capture in that database. And they were.

What they didn't know is that Change Data Capture harvests committed transactions on the tables being monitored from the transaction log. And it does this using the replication Log Reader Agent job if replication is configured, or its own capture job if replication is not configured. One of their developers had written custom logic around the call to `sp_cdc_scan` in the capture job and had made a mistake that prevented the scan from starting, so the log just grew and grew.

But why did the `log_reuse_wait_desc` say `REPLICATION` and not CDC? Because there is no separate indication that CDC is preventing the log from clearing – it uses the `REPLICATION` value, which can cause a lot of confusion!

Bottom line: If you ever see a `log_reuse_wait_desc` value of `REPLICATION`, and replication isn't configured, check the `is_cdc_enabled` flag in `sys.databases` as well, and then troubleshoot CDC from there.

Paul's Ponderings

Occasionally one of the team has a topic they'd like to write about for the newsletter, so this time the Ponderings is written by Erin. Enjoy!

Last month at the PASS Summit I presented a session titled [Query Store and Automatic Tuning in SQL Server](#). The session was focused on Query Store and new, related features in SQL Server 2017. Something many people may not be aware of is that [Automatic Tuning](#) covers **two** features: *Automatic Plan Correction* and *Automatic Index Management*. The former feature is available in SQL Server 2017 and Azure SQL Database, the latter is only available in Azure SQL Database.

In fact, there are many features only available in Azure, as the development cadence is much faster than that of the box product. Microsoft published a blog post from Conor Cunningham last week, [How Cloud Speed Helps SQL Server DBAs](#), and I recommend you give it a read, even if you're a DBA or developer for an on-premises system. Why? Because while you may be a data professional for an on-premises system today, you don't know where you're going to be in a year, or two, or three... Nor do you know what your company may decide to implement, what features might be beneficial down the road, and whether your system will always live in a local data center. I'm not saying you need to become an expert in all things SQL Server; you just can't, as the product is too big. But I do think it's wise to have high-level knowledge of what's available.

For example, take the automation which exists in Azure SQL Database for tuning and troubleshooting (see [here](#)). If you drill into this content in more detail, the performance recommendations aren't *just* specific to forcing plans automatically and creating and dropping indexes automatically; there are recommendations related to forced parameterization and schema change... and Microsoft is just getting started.

With over 2 million customers running Azure SQL Database, Microsoft is capturing terabytes (yes, terabytes) of data *daily!* You may have heard the term "*data driven*" and if so, you might associate it with Microsoft marketing and think it's just a buzzword. But the truth is these automated features have been developed *with that telemetry*, and for any of these features, SQL Server relies on the data it captures to make informed decisions – for example should a query should have a forced plan, or should an index be created or dropped?

This data driven technology reduces the amount of work that you as a DBA or developer or application administrator has to do. You don't have to try to roll your own process to capture query data or index usage information; SQL Server will do it for you. And you don't have to spend time analyzing all that data and figuring out what's useful; SQL Server will do it for you.

Now, there are many people who become concerned when they hear about these new features in Azure. Some don't trust the automated processes. Others hear the term "automation" and believe that the product team is trying to code them out of the job. With regard to automation, it's important to understand that the same old tools aren't being dropped in here. Automatic Index Management is not just the Database Tuning Advisor running in the background. It's an improved and more complex model that's being used, and it is constantly evolving. The cloud-

first approach allows Microsoft to continually capture telemetry about how customers are using features, what customer workloads look like, and whether the changes that get implemented truly provide benefit. This is an advanced, more robust technology than we've seen before, and it's absolutely worth checking out.

Secondly, Microsoft is not trying to take away your job (in fact, it's trying to help customers who have hundreds or thousands of databases but do *not* have a database administrator... and yes, those customers do exist). If you are tied to the on-premises version of SQL Server and have no interest in what's going on in the cloud, that's fine. I get it. In fact, that's how I felt about six months ago. I love the box product; it's what I've worked with since I started my career in SQL Server. It's "[the fork I know](#)".

But as I've started working with these newer features in SQL Server, it's truly hit home that code is coming first to the cloud, and it gets refined and then eventually seems to make it to a major box release. There are many features only in Azure SQL Database, and there's truly *no guarantee* they will ever be pushed to the on-premises version. As someone who likes to work with new technology, and loves the idea that decisions are being made based on data, I've found that Azure is a cool place to be.

Call to action:

Am I telling you that you need to drop everything and change direction? No.

Am I telling you that DBAs will no longer be needed? No.

What I **am** telling you is that there are many features in Azure worth checking out, and even though it may not be an option for you today because of limitations, don't be naïve and think those limitations are permanent. Azure has grown in leaps in bounds over the last couple years, and there's nothing to suggest that will end anytime soon.

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 17.4

Microsoft has released [SQL Server Management Studio 17.4](#), which is the latest standalone version of SSMS. This version of SSMS is completely free, and is officially supported for SQL Server 2008 and newer, and for Azure SQL Database. It runs on x64 versions of Windows 7 SP1 and Windows Server 2008 R2 and newer.

SSMS 17.4 has a very significant new feature called "[Vulnerability Assessment](#)", which lets you scan your databases for vulnerabilities and deviations from security best practices. You can get

to this feature by right-clicking on a database in Object Explorer, selecting Tasks, selecting [Vulnerability Assessment](#), and then selecting Scan for Vulnerabilities.

The resulting scan is pretty quick, and it exposes a lot of information about what it is scanning for and what it finds, including the actual T-SQL to correct many of the issues that it finds. Vulnerability Assessment only works on SQL Server 2012 and newer instances.

This new feature makes SSMS 17.x even more valuable. I think it is far superior to legacy versions of SSMS, and you should probably start using it to manage even your legacy versions of SQL Server, rather than using the version of SSMS that is part of your legacy version of SQL Server.

Microsoft is Retiring Microsoft Connect on December 31, 2017

On December 11, 2017, Microsoft [announced](#) that “Connect will be decommissioned and shut down after 12/31/2017 due to lack of compliance with GDPR privacy regulations”. They also state that “We have captured all feedback in our internal database, and will replicate it as much as possible into a new forum on UserVoice”. Finally, they state “We will post a link to the new UserVoice site here once it is up and running.”

In case you are wondering, [UserVoice.com](#) is a commercial product meant to solicit and manage customer feedback about products. Hopefully, the SQL Server UserVoice site will be a substantial improvement over the miserable experience that was Microsoft Connect. Microsoft Connect was painfully slow, and not very much fun to use.

#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 query plan analysis:

- Erin's Pluralsight course: [SQL Server: Analyzing Query Performance for Developers](#)
- Joe's Pluralsight course: [SQL Server: Query Plan Analysis](#)
- Erin's Pluralsight course: [SQL Server: Introduction to Query Store](#)
- Joe's Pluralsight course: [SQL Server: Troubleshooting Query Plan Quality Issues](#)
- My blog post: [Query plan analysis first steps](#)
- Joe's whitepaper: [Optimizing Your Query Plans with the SQL Server 2014 Cardinality Estimator](#)

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

- Paul: [2012/2014 bug that can cause database or server to go offline](#)
- Paul: [New Pluralsight course: Implementing a Hybrid Cloud Solution in Azure](#)
- Paul: [New class: Practical Machine Learning](#)
- Glenn: [SQL Server Missing Indexes Feature and Trace Flag 2392](#)
- Tim: [SQLskills SQL101: Using Backup to URL for Azure VM Migrations](#)

I hope you find these useful and interesting!

Video Demo

Wait statistics data has been available for capture by Query Store since mid-summer for Azure SQL Database, but the feature just hit on-premises installations with the release of SQL Server 2017 this past October. If you haven't explored this new data in Query Store yet, watch this video where Erin shows how to see what information is captured and how to find it once you've upgraded.

The video is just over ten minutes long and you can get it:

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

The demo code is [here](#).

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 (**discounted until 12/20/17, 9am PST**)

- **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 **(all classes have discounts for 2017 registrations!)**

- **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 **(all classes have discounts for 2017 registrations!)**

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

London, UK, September 2018 **(all classes have discounts for 2017 registrations!)**

- **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