

(June 5<sup>th</sup>, 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 this hits your inboxes we're at home in Redmond for a couple of weeks enjoying the sunshine before heading off to Dublin for our first ever SQLSaturday! Our Spring SQLIntersection finished last week and was a great success. We had tremendous feedback that people enjoyed the event and especially the ability to ask us questions and interact with the speakers at evening events such as SQLafterDark. We gave away over US\$28k in prizes at the end-of-the-evening raffle and now it's time to start planning for our Fall show. We hope you can join us for that in Las Vegas in October – more details coming soon!

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

## SQLskills News

We've released our **October line-up of classes in Chicago**, including IE0, IEPTO1, IEPTO2, IEPDS, IESSIS1, plus THREE new courses. We're debuting a [new two-day class on Azure SQL Database and Azure VMs](#), a [new three-day class on Upgrading SQL Server](#), and a [new two-day class on Clustering and Availability Groups](#). See [here](#) for the complete 2017 SQLskills Immersion Event class schedule.

We've started a new initiative where we're all blogging about introductory topics, to help the burgeoning numbers of non-DBAs or junior DBAs who are responsible for SQL Server instances. We're calling it SQL101, and hopefully it'll even help refresh or reinforce topics for the more seasoned DBAs in the community. The blog posts will be automatically collected [here](#).

Even though we're not teaching any Immersion Events in Europe this year, Kimberly and I will both be [presenting at SQLSaturday #620 in Dublin next week](#). This is coming up quickly so we hope you can join us for our two-day SQLskills Performance Track prior to SQLSaturday in Dublin – check out the details [here](#).

**Finally, even if you can't join us in person**, we're still taking requests for remote user group sessions for this year. We have over 70 scheduled so far; 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 Peter F. Hamilton's [A Night Without Stars](#). Hamilton is one of my favorite authors, and this book is a continuation of his epic Commonwealth series, and the second in a two-part story arc. Start by reading the first book, [The Abyss Beyond Dreams](#), and then read this one, as this book doesn't really stand on its own IMHO without reading the first book. The story involves wormholes, aliens, sentient A.I.s, lots of drama, and long-standing characters from within the larger series. Highly recommended for sci-fi fans!

### **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.*

Jonathan was working with a client recently who was seeing a lot of RESOURCE\_SEMAPHORE waits and workload degradation. These occur most commonly when a query cannot get the required query execution memory grant.

After some investigation, he found a query that was being executed by many concurrent connections, each requiring a 360MB memory grant. By changing the query to use a temp table for an intermediate result set, the execution memory grant decreased to 90MB.

However, the introduction of the temp table, with multiple concurrent connections, caused tempdb allocation bitmap contention (PAGELATCH\_UP waits on the PFS pages, in this case). So he increased the number of tempdb data files from 4 to 16, which eliminated the contention and had no detrimental effect to any other part of the workload.

**Bottom line:** sometimes the best way to solve a performance problem can introduce another performance problem instead, but that doesn't mean it's an incorrect solution, if there's an easy way to fix the introduced problem as well, for the overall benefit of the workload.

### **Paul's Ponderings**

I was talking to a SQLintersections attendee last week and he told me how he'd had a really hard time changing jobs recently because he'd been working in a SQL Server 2008 R2 shop for many years and had no experience with SQL Server 2012 or newer and prospective employers weren't impressed with that. He'd eventually acquired experience on his own time and then was able to successfully interview for a new job.

My guess is that quite a few of you out there are in a similar position.

Unless you're content to stay in that same job using nothing but old versions of SQL Server (and 2008 R2 is really old – it was released almost seven years ago and mainstream support for it ended more than two years ago), you need to get some experience with newer versions.

I've helped a few friends and former mentees get new jobs over the last year, and every single one of them required experience with availability groups, which was introduced in SQL Server 2012. Given that it's the replacement for database mirroring, and is the flagship high-availability technology in SQL Server, this is an obvious job requirement for someone looking for a new DBA.

Now with SQL Server 2016's release last year, there's a new must-have-experience-with feature: query store. Anyone running on SQL Server 2016 is going to want to use this for workload tuning, so that's going to become a DBA-job requirement too.

There are many completely free ways you can learn about SQL Server:

- Reading blog posts
- Picking a few features and studying them in [Books Online](#)
- Following forums on sites like [SQL Server Central](#)
- Attending free webcasts from the various [PASS Virtual Groups](#)
- Attending [SQLSaturdays](#)
- Attending your [local user group](#)

It's also really easy to get your hands on the most recent version of SQL Server as Developer Edition is completely free when you sign up for a free [Visual Studio Dev Essentials account](#). And with that you even get Pluralsight for free for three months too! (including Erin's [new course on Query Store](#)...)

**Call to action:** The excuse of “*we only use SQL Server <insert old version> so I don't have any experience with later versions*” isn't really valid as far as I'm concerned. With a plethora of free software and learning paths available, the only thing you need to provide to get up-to-date experience is your own time. That's definitely worthwhile if you're looking to change to a better job.

### **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.*

There were lots of new product announcements from the Computex 17 show in Taipei last week.

#### **New Intel High-End Desktop Processors**

Intel [has announced the upcoming availability](#) of several new families of high-end desktop (HEDT) processors, including the Skylake-X and Kaby Lake-X lines. Both of these families are part of the new “Basin Falls” high-end desktop platform, which uses the new X299 chipset, and features a new processor socket (LGA2066).

The Kaby Lake-X processors will only have four cores, plus hyper-threading, and they will have slightly higher base and turbo clock speeds than the existing Kaby Lake-S mainstream desktop processors. They will not have integrated graphics, so more power will be available for the actual processor cores.

The Skylake-X processors will have “low core count” (LCC) models with six, eight, ten, and twelve physical cores (plus hyper-threading), and high core count (HCC) models with fourteen, sixteen, and eighteen physical cores (plus hyper-threading). The X299 motherboards will have eight DIMM slots.

The Kaby Lake-X quad-core CPUs will feature 16 PCIe 3.0 lanes, while the Skylake-X family will split between 28 PCIe 3.0 lanes for the Core i7 processors and 44 PCIe 3.0 lanes for the high-end Core i9 processors.

This sounds pretty impressive, but keep in mind that the high-end Core i9 models will have prices of between \$1000 and \$1999. Still, if you are willing to spend some money, you will be able to have a desktop machine capable of running multiple large VMs quite well.

### **Intel Optane SSD 900p**

Intel [teased a keynote appearance](#) of their upcoming large-capacity consumer Optane SSD which is supposed to be called the Intel Optane SSD 900p. It will [have models with capacities](#) of 280GB, 480GB, 960GB, and 1.5TB. Optane, which is Intel’s branding for 3D XPoint, offers extremely high performance at low queue depths, along with better latency than any existing NAND flash storage. These drives are likely to be pretty expensive, and Intel says they will be available later in 2017.

### **AMD Ryzen Threadripper Desktop Processors**

AMD is [due to release](#) a new line of high-end desktop processors with the code-name “Threadripper” in July, 2017. These will have up to sixteen physical cores and AMD’s SMT technology, for 32 threads per processor. These processors will use the new X399 chipset and will have eight memory slots. They are likely to be less expensive than similar core count Intel Skylake-X processors. There is also pretty wide-spread speculation that the Intel Skylake-X HCC processors may be a competitive reaction to the AMD Threadripper line.

### **#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.*

One of my favorite topics on which to present is the transaction log. It’s so important that I try to present it at every SQLIntersection conference. Just because it’s not a new feature and hasn’t

changed much in recent releases, that doesn't mean DBAs and developers can ignore it. It's one of the most misunderstood parts of SQL Server and yet it's one of the most important – both for performance AND for availability. But, even more than that, it's the only way to guarantee up-to-the-minute recovery with no data loss.

So the transaction log is the theme for #TBT this time. Here are some transaction log resources for you:

- My 8-hour Pluralsight courses: [SQL Server: Logging, Recovery, and the Transaction Log](#)
- My [transaction log blog category](#)
- My [sqlperformance.com blog posts](#) on transaction log configuration and performance troubleshooting
- Kimberly's post on [VLF sizing](#)
- My TechNet Magazine article on [Understanding Logging and Recovery in SQL Server](#) (the formatting is messed up a bit on that web page)
- The [WRITELOG](#) page from my waits library

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

- Paul: [SQLskills SQL101: How can a SELECT cause a database to change?](#)
- Glenn: [SQLskills SQL101: SQL Server Core Factor Table](#)
- Erin: [Query Store Trace Flags](#)
- Jonathan: [Tracking Compiles with Extended Events](#)
- Jonathan: [SQL Server 2016 Distributed Replay Errors](#)
- Jonathan: [Using Group Managed Service Accounts for SQL Server](#)
- Jonathan: [SQLskills SQL101: NUMA Support in SQL](#)

I hope you find these useful and interesting!

## **Video Demo**

In this video Tim continues his series on how to perform maintenance task on Azure SQL Databases. In the [first video](#) he demonstrated the use of a linked server, and in the [second video](#) he showed how to use database maintenance plans and PowerShell. This time Tim shows how to automate maintenance using Azure Automation and Elastic Jobs.

The video is around 6.5 minutes long and you can get it:

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

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

Enjoy!

## **Upcoming Immersion Events**

All 2017 classes are available 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](#)
- [Community blog posts about our classes](#)
- [Immersion Event FAQ](#)

Bellevue, WA, July/August 2017

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - July 31-August 4
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
  - August 7-11

Chicago, IL, October 2017

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - October 2-6
- **IE0:** Immersion Event for the Junior/Accidental DBA
  - October 2-4
- **IECAG:** Immersion Event on Clustering and Availability Groups
  - October 5-6 **\*\* New class! \*\***
- **IESSIS1:** Immersion Event on Learning SQL Server Integration Services
  - October 2-6
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
  - October 9-13
- **IEPDS:** Immersion Event on Practical Data Science
  - October 9-13
- **IEAzure:** Immersion Event on Azure SQL Database and Azure VMs
  - October 9-10 **\*\* New class! \*\***
- **IEUpgrade:** Immersion Event on Upgrading SQL Server
  - October 11-13 **\*\* New class! \*\***

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](mailto:Paul@SQLskills.com) and [Kimberly@SQLskills.com](mailto:Kimberly@SQLskills.com)