

(February 18th, 2020)

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,

Kimberly put together a series of posts that introduce all of us and link to some photo posts of our lighter side – check it out at [SQLskills – Meet Our Amazing Team!](#)

Back in January the whole SQLskills team recorded a 2-hour webinar with Kevin Kline of SentryOne called ‘*Troubleshooting Horror Stories and the Surprising Lessons They Teach*’. It was broadcast last week and there’s an open Q&A session on 2/19. Check out the registration links at the bottom of <https://www.sentryone.com/events!>

SQLskills News

Live, IN-PERSON classes: our Spring 2020 classes in Chicago are filling up fast:

- IEPTO1: Performance Tuning and Optimization, Part 1 – April 27-May 1
- IE0: Accidental/Junior DBA – April 27-29
- IECAG: Clustering and Availability Groups – April 30-May 1
- IEPTO2: Performance Tuning and Optimization, Part 2 – May 4-8
- IEAzure: Azure SQL Database, Azure VMs, Azure Managed Instance – May 4-7

We’ve also **added a new IEReporting** Immersion Event on Using and Administering SSRS and Power BI, running April 27-28.

You can get all the details and registration information through the class schedule page [here](#).

SQLBits: Kimberly, Erin, and I will all be presenting at [SQLBits](#) in London, March 31-April 4. Our precons are:

- Paul: 3/31 [Performance Troubleshooting using Waits and Latches](#) SOLD OUT!!
- Erin: 4/1 [Performance Tuning with Query Store in SQL Server and Azure](#)
- Kimberly: 4/1 [Statistics4Performance: Internals, Analysis, Problem Solving](#)

This is the only chance to see us in Europe this year as we won’t be bringing classes back to the UK until 2021. We hope to see you at SQLBits!

SQLintersection: Our entire SQLskills team will be in Orlando for performance/troubleshooting/development-focused conference. Remember, better (and earlier!) design will make everyone's lives easier! The Intersection-suite of events can really help your team improve as our focus is primarily on real-world problems that you're seeing! Our Spring show in Orlando features 9 full-day workshops, 46 technical sessions, and 3 SQL-specific general/keynote sessions plus our SQLafterDark evening event and our RunAs Radio closing session with Richard Campbell. And, you always have the option of attending a session from one of our co-located shows: [DEVintersection](#), [Azure/AI Intersection](#), and [SQLintersection](#). **It's our LARGEST show yet!** But, also a smaller show overall so it's a great place to really meet the speakers / attendees and network! We're really looking forward to it.

As for the SQLintersection workshops, check out our line-up:

- 3 pre-conference workshops on Sunday, April 5
 - [Getting Started as a Data Engineer](#) with Tim Chapman
 - [A Lap Around SQL Server for Linux Folks](#) with Jonathan Kehayias
 - [Leveling up with PowerShell for the DBA](#) with Ben Miller

- 3 pre-conference workshops on Monday, April 6
 - [Performance Troubleshooting using Waits and Latches](#) with Paul Randal
 - [Management, Administration, and Best Practices for the Hybrid DBA \(SQL Server 2016-2019, Azure SQL Database, Managed Instance\)](#) with Tim Radney and David Pless
 - [The Developer's Guide to Consistency, Accuracy, Concurrency, and Transactional Control](#) with Kimberly L. Tripp

- 3 post-conference workshops on Friday, April 10
 - [Kubernetes Zero to Hero: Installation, Configuration, and Application Deployment](#) with Anthony Nocentino
 - [Performance Tuning and Optimization for Modern Workloads \(SQL Server 2016-2019, Azure SQL Database, Managed Instance\)](#) with Tim Radney and David Pless
 - [Query Tuning Tools in SQL Server](#) with Grant Fritchey

There are also co-located workshops with each of the conferences so you can mix and match! Our workshops are scheduled so that you and your developers have multiple paths to get the most out of your workshop and conference experience! Check out the entire session line-up online [here](#). And, be sure to register with the discount code 'SQLskills' to get a \$50 discount.

Finally, even if you can't join us in person, I've put out a call for **remote user group sessions for 2020!** If you'd like one of us to present for your user group, check out my blog post [here](#).

Book Review

Here's one from the archives: Peter Frankopan's [The Silk Roads: A New History of the World](#). This is a very interesting book tracing the varied history of the countries along the Silk Road, including empires, explorers, and religions that affected the various routes. A lot of what's in the book I already knew, but having it all presented in one volume in a chronological sequence was excellent. The last 100 pages or so detailed the quite despicable British and American machinations around the countries in the Middle East for their own (mostly oil-related) gains, to the huge detriment of the native populations, which I felt quite ashamed to read about, being of both nationalities. Highly recommended for history fans!

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.

Last week I went through an interesting thought exercise around restoring a table from a backup directly into another database – read through my findings and thoughts [here](#)...

Ponderings...

(From Erin this time – enjoy!)

Last summer, Tim wrote a pondering about the importance of testing changes before implementing them in production. I want to expand on what he wrote and talk about *how* you do testing.

Now, I admit that I consider myself a scientist. My undergraduate and graduate degrees are from schools of science, and for my master's degree I wrote a thesis on my research project, which included testing and data analysis, with my conclusion based on what I found. I am a data person at heart.

So, when I see scenarios where test results are shared without accurate testing, I get extremely frustrated. With customers this can lead to incorrect conclusions, additional problems, or simply confusion causing it to take longer than necessary to find a solution. In the SQL Server community I sometimes see these results categorized as “bad” or “dangerous” and the summary suggests things like “I wouldn't use that feature in SQL Server” or “it looks like the SQL Server team really messed up.” If an individual doesn't understand true A/B testing, they can think misinformation is truth.

For example, if we want to understand the impact of changing MAXDOP for the instance, we run a set of queries that represent our workload. We run these queries a specific number of

times, with pre-defined input values, and we capture metrics about the performance (including execution plans) of the queries.

Then, we change MAXDOP for the instance. That is the **ONLY** change we make. And we run the EXACT **same** set of queries, the **same** number of times, with the **same** input values, and we capture the **same** metrics about performance.

Scientifically, we can look at this scenario and identify our hypothesis (what we think will happen), what we will change (independent variable), what we will be the outcome of the change (dependent variable). For our MAXDOP testing, our hypothesis is that we think changing MAXDOP will improve performance, our independent variable is MAXDOP, and our dependent variable is the query metrics we collect.

Now, if you're going to point out that this kind of testing is really difficult to do in a production environment, I wholeheartedly agree. A/B testing in production is ridiculously hard, if not impossible. That's why we don't want to test in production if we don't have to. That's why, as Tim discussed, you need another environment that mimics production. But where possible, we want to isolate the issue and change only one thing. I know this can be complicated – upgrading to a new version of SQL Server on new hardware makes it impossible to only change *one* thing.

But I trust you get my point: you want to make sure you're comparing apples to apples, whether you're testing something for a blog post or an article to prove or disprove an idea that you have, or whether you're trying to determine if a trace flag, or code change, or new feature will solve your production problem. The act of claiming you are testing, but not doing it correctly, is a disservice to the people you want to help, be it your company or the SQL Server community.

#TBT

(Turn Back Time...) Blog posts we've published since the previous newsletter plus some older resources we've referred to recently that you may find useful.

Jonathan's just started a blog post series on replication, so that's the theme for the #TBT section this time. Here are some resources for you:

- Joe's Pluralsight course on [SQL Server: Transactional Replication Fundamentals](#) that has 2 hours explaining what replication is, how to set it up, how to monitor and troubleshoot it, and how to combine it with other HA technologies
- Whitepaper on [SQL Server Replication: Providing High-Availability Using Database Mirroring](#) that I wrote for Microsoft back in 2008, but is still completely relevant today and can be extrapolated to working with Availability Groups too.

- Old whitepaper on [*Proven SQL Server Architectures for High Availability and Disaster Recovery*](#) that I wrote for Microsoft in 2010, and isn't listed on their site any more, but has an interesting set of architectures, including one using peer-to-peer replication.
- And a bunch of blog posts:
 - [*In defense of transactional replication as an HA technology*](#)
 - [*REPLICATION preventing log reuse but no replication configured*](#)
 - [*The Transactional Replication Multiplier Effect*](#)
 - [*When is the Publication Access List required?*](#)

And check out the [Repltalk blog](#) that's been a fount of replication knowledge since 2010.

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

- Kimberly: [*SQLskills – Meet Our Amazing Team*](#)
- Jonathan: [*Understanding SQL Server Replication*](#)
- Jonathan: [*Transactional Replication Use Cases*](#)
- Erin: [*Are Changes to Query Store Logged?*](#)
- Erin: [*#TSQL2sday: My Life Hacks*](#)
- Tim: [*Gaining Access to SQL Server When Locked Out*](#)

I hope you find these useful and interesting!

Video Demo

In this Insider video, Tim walks you through how you can regain control of a SQL Server instance in the event you've locked yourself out. Tim recently had to help a customer with this issue and wrote an updated blog post covering the steps involved for a default and named instance – check it out [here](#). Hopefully this isn't something that you encounter often, but it's a great tip to know exist in case you ever need it.

The video is 4.5 minutes long and you can get it [here](#).

No demo code this time.

Enjoy!

Upcoming SQLskills Events

Our Spring 2020 classes are all open for registration!

You have multiple learning opportunities as every event has a different focus as well as different benefits – from deep-technical training in our Immersion Events to a more broad set of topics at

SQLBits / SQLintersection! 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, Spring 2020

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
 - April 27 – May 1
- **IEReporting:** Immersion Event on Using and Administering SSRS and Power BI
 - April 27-28
- **IE0:** Immersion Event for the Junior/Accidental DBA
 - April 27-29
- **IECAG:** Immersion Event on Clustering and Availability Groups
 - April 30 – May 1
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
 - May 4-8
- **IEAzure:** Immersion Event on Azure SQL Database, Azure VMs, and Azure Managed Instance
 - May 4-7

LIVE, ONLINE Immersion Events:

- More coming in 2020!

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