

(February 4th, 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,

This newsletter is coming to you from Redmond, where we're working hard on converting all our demos over to SQL Server 2019 in preparation for classes, conferences, and user groups!

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'll be broadcast on 2/12 (2/13 in EMEA) with 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:

- 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 precon are:

- Paul: 3/31 [Performance Troubleshooting using Waits and Latches](#)
- 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

For the first time in a long time, I didn't finish reading any books in January, so this time I present a kind of meta-book – DK's [*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 recent problems we've helped with on client systems; they might be something you're experiencing too.

Tim came across an interesting issue last week when he noticed a sudden latency and memory usage increase on a server of one of our remote DBA clients. Read about the story [here](#)...

Ponderings...

*(We had some issues with a new client last week who experienced performance issues with a *huge* database, and it turned out that there was a lot of unnecessary data that could be archived out. So, to go along with the #TurnBackTime section below on partitioning, this is an updated version of an editorial I wrote on the subject exactly 5 years ago. With data volumes exploding, this is even more important today!)*

One of the things that DBAs and developers don't think about very often (from experience) is the amount of data that's in a table, tables, or database. In general the data volume is just accepted – of course it's supposed to be there, it's in the database, right?

Well, no. If you step back and think about it, there are a few important questions that you should be asking:

1. Does all this data really **need** to still be in the production database?
2. Is all that data **supposed** to be in the production database?
3. Does each query **have** to process all that data?

If a production database is bloated with unnecessary data, there are a bunch of problems that can result, such as:

- Storage requirements are unnecessarily large (and don't just think of the database itself, also consider AG secondary replicas, a mirror database, replication subscription databases, log shipping secondary databases, test and development database copies)
- Database backups back up older data making them unnecessarily large and longer to execute, even when using backup compression
- Index fragmentation (e.g. index rebuilds) generates more transaction log than necessary (which can have a host of knock-on effects around log performance, log backup size, replication, and so on) and takes longer
- Statistics updates take longer
- Queries may process more data than necessary, affecting query plans, performance, buffer pool space, I/O performance, and more (I remember one client where a weekly proc to calculate attendance metrics for a school system was pulling 15 years of data into a temp table, blowing out tempdb...)

The list goes on and on.

If the answer to the original three questions is 'no' then they can be mitigated by:

1. Careful archiving of unneeded data on a regular basis, either through a partitioning process like the 'sliding window', or through data deletion that's architected to happen in manageable chunks (rather than one enormous delete statement that takes forever to complete, probably blows out the log, and causes bigger problems than those it's trying to solve)
2. Analyzing code to make sure that extra data isn't erroneously being inserted, imported, copied, duplicated, or whatever. This may mean discussing the data in the database with the developers/architects first to figure out where the extra data is.

I've seen plenty of clients have problems from just having too much data that's not needed, and the easiest solution is to migrate it out. Addressing the symptom by doing something like enabling data compression everywhere isn't the answer – you need to address the root cause, just like with any other performance problem.

Call to action: If you're finding that your data volume is just growing and growing and it's starting to cause problems, consider whether all the data is needed or should be there. If you've never thought about these questions, or checked on your production database, you might be surprised to find problem data that you can migrate out to ease your problems.

#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.

Partitioning is the theme for this TBT:

- Kimberly's intro to partitioning post: [SQLskills SQL101: Partitioning](#)
- Original partitioning whitepaper that Kimberly wrote (which is scenario based so it's still very applicable to all versions): [Partitioned Tables and Indexes in SQL Server 2005](#)
- Additional whitepaper written with focus on features rather than scenarios: [Partitioned Table and Index Strategies Using SQL Server 2008](#)
- Kimberly's [recorded partitioning training](#) for the MCM certification

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

- Jonathan: [Securing Data in SQL Server](#)
- Jonathan: [Keyset Does Not Exist – Add-ClusterNode](#)
- Jonathan: [Strongly Type Those Table-Valued Parameters](#)
- Tim: [When Updating Statistics is too Expensive](#)
- Erin: [What Events are in a Trace?](#)

I hope you find these useful and interesting!

Video Demo

In this video, Erin demonstrates DBCC CLONEDATABASE, a command introduced with SQL Server 2016 SP1, and ported back to SQL Server 2014 SP2 CU3. It's a quick and easy method for creating a schema-only copy of a database, which can be used to troubleshoot query and plan issues, test upgrades, and create a verified copy of a database.

The video is 8 minutes long and you can get it [here](#).

The demo code is [here](#).

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 **** NEW and US\$100 off through 1/25/20 ****
- **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