(April 28th, 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 still hunkered down for an extended period of social distancing. Tim and I are still helping people out using our six 3D printers to make mask straps that take pressure off of ears (see here) we've printed and shipped/donated almost five thousand between us so far.

Many thanks to Redgate for donating $500 to each of us to help cover costs! If you know anyone who could use some free straps, hit me up in email – don't be shy.

And we also created a SQL Server 3D Printing Facebook group – check it out here.

I hope you all stay safe and healthy!

SQLskills News

Kandio job candidate assessments: we've teamed up with Kandio to produce technical assessments to help companies screen candidates for job recruitment. Fancy trying you hand to see how you score? See here to try the DBA and Developer assessments for free.

Online class recordings: you can buy recordings of all our online classes, for as little as US$299 each. See here for all the details.

New client discount: we've dropped the price of our health check to US$2,495 for all new clients between now and the end of April. Send us an email here for more details

Live, IN-PERSON classes: our Chicago classes have moved to October:

- IEPTO1: Performance Tuning and Optimization, Part 1 – October 19-23
- IE0: Accidental/Junior DBA – October 21-23
- IECAG: Clustering and Availability Groups – October 19-20
- IEReporting: Using and Administering SSRS/PowerBI – October 21-22
- IEPTO2: Performance Tuning and Optimization, Part 2 – October 26-30
- IEAzure: Azure SQL DB, Azure VMs, Azure Managed Instance – October 26-29

You can get all the details and registration information through the class schedule page here.

SQLBits: Kimberly, Erin, and I will all hopefully be presenting at SQLBits in London in September. Our precons are:

- Paul: Performance Troubleshooting using Waits and Latches SOLD OUT!!
- Erin: Performance Tuning with Query Store in SQL Server and Azure
- Kimberly: Statistics4Performance: Internals, Analysis, Problem Solving

SQLintersection: The Spring show has been canceled and everything has moved to the Fall show in Las Vegas in December. More details here nearer the time.

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

From the archives: Back in 2016 I read Charles Petzold's *Code: The Hidden Language of Computer Hardware and Software*. This book is really excellent! It's a very cleverly written introduction and exploration of how computers work at the logic level, and takes the reader from the basics of electricity and binary to the intricacies of memory busses, CPUs, and assembly language. I quickly skimmed the first hundred or so pages until I got to the part about building counters from relays and it started to refresh my memory with things I'd learned back in 1990 when I did my B. Eng (Hons) degree in computer science and electronics in Edinburgh.I read this book as a way to kick start getting back into computer design as I want to build a CPU and computer system out of TTL logic (one of my many, many 'spare time' goals). First though, when I get home I'm going to build some logic circuits out of relays – just for the fun of hearing all the little clicks as the relays change state :-)! I highly recommend this book for anyone who wants to know a bit more about how computers work.

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 someone told me they'd read an article that stated that a database's recovery interval setting is always guaranteed, no matter what, and wondered how that can be true. I already had a post on that, so you can read through my explanation here…

Ponderings…

(My turn for the Ponderings this time, adapted from an editorial I posted back in 2013.)

One of the questions we get asked every so often is why a developer or DBA should know some of how SQL Server works – the internals – and isn't that just knowledge to satisfy intellectual curiosity? As we're many of us are stuck at home, maybe with some extra time on our hands, now is a good time to think about learning some new info that can help with your career.

Here are a couple of examples of why it's useful to know how SQL Server does things under the covers, and how it can impact your performance.

Firstly, understanding how SQL Server stores column values in records. This is really relevant when designing an efficient table schema. Over the years we've seen several cases where a client was indicating "whether this value has been deleted or not" using an *nvarchar (3)* column, using the values "YES" and "NO". This will take 6 bytes for "YES", four bytes for "NO", and at least two extra bytes in either case for the pointer in the variable-length column offset array in the record. Instead, the indicator could be stored as a *bit* column, which takes only 1 bit and can be combined within one byte for up to eight separate *bit* columns in a record (NOTE: when there's only one bit column then the row will still reserve one byte; rows are always on byte-boundaries, not bit-boundaries).

Having said that, a one billion-record table, using the *nvarchar (3)* column (let's say with half "YES" and half "NO" values) would take 500 million * 8 + 500 million * 6 = 6.5GB to store that column, compared with 1 billion * 1 = 0.93GB for a *bit* column. That's a huge difference. Some might say that in the grand scheme of things, 5.5GB extra is negligible, but if many more columns have inefficient design, the extra space adds up quickly. Think of extra disk space, extra log space, extra buffer pool space, lower data density, larger backup sizes, longer restore times, and so on. Kimberly recorded a Pluralsight course that covers this topic in detail: SQL Server: Why Physical Database Design Matters.

Secondly, understanding how page/extent allocation works. This is most relevant when dealing with tempdb *PAGELATCH_XX* contention. Understanding that various allocation bitmaps need to be updated in memory whenever a page or extent is allocated explains why the contention occurs (on most versions of SQL Server) when many concurrent connections are creating and dropping small temp tables. Understanding how adding more tempdb data files allows the round-robin allocation algorithm to be used explains why adding more tempdb data files can reduce the overall *PAGELATCH_XX* contention on in-memory copies of allocation bitmaps (by having more bitmaps to share amongst the threads instead of just one) and lead to a workload throughput increase. See this blog post for more details.

Even something as seemingly obscure and esoteric as how SQL Server computes allocation unit IDs (which was the Curious Case from the previous newsletter) is useful to know when dealing with database corruption and you need to manually find data file pages belonging to a corrupt table because *sys.sysallocunits* is damaged so queries and *DBCC CHECKDB* fail.

Call to action: Don't dismiss knowing how SQL Server works as useless information. Knowing how SQL Server stores data, uses indexes, uses statistics, manages the transaction log, manages transactions, uses locks, uses latches, and so on can really help you design and tune your databases and workloads for optimal performance. And there's a wealth of information out the, ranging from blog posts to books to online training to in-person training. With the SQL Server community really booming over the last few years, there's never been a time where more information is available on how SQL Server works and how you can capitalize on that knowledge.

#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.

The #TBT this time is deadlocks so here are some resources for you:

- Jonathan's Pluralsight courses: *SQL Server: Deadlock Analysis and Prevention*
- The deadlock entry from our Accidental DBA blog-post series
- Jonathan's very long SimpleTalk article on *Handling Deadlocks in SQL Server*, which is an excerpt from his popular (and free) ebook *Troubleshooting SQL Server: A Guide for the Accidental DBA*

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

- Paul: *Try one of our new candidate screening tests at Kandio!*
- Tim: *Updated Azure SQL Database Tier Options*

I hope you find these useful and interesting!

Video Demo

Need some help troubleshooting in Azure SQL Database? In this video Erin steps through using three scripts that you might be familiar with that you can use to look at performance for any Azure SQL Database.

The video is around 7 minutes long and you can get it in WMV format here.

Enjoy!

Upcoming SQLskills Events

Our 2020 classes have moved to October and 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, Fall 2020

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - October 19-23
- **IEReporting**: Immersion Event on Using and Administering SSRS and Power BI
  - October 21-22
- **IE0**: Immersion Event for the Junior/Accidental DBA
  - October 21-23
- **IECAG**: Immersion Event on Clustering and Availability Groups
  - October 19-20
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - October 26-30
- **IEAzure**: Immersion Event on Azure SQL Database, Azure VMs, and Azure Managed Instance
  - October 26-29

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