(September 29th, 2021)

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!

This newsletter is coming to you from Concrete, WA, where I'm still camping and working on our land – it's so peaceful out here – but it's starting to get a little rainy as well as chilly!

Registration is open for our first free SQLskills Insider learning session on Wednesday, October 6 – see here!

Take care and stay safe out there!

SQLskills News

Webcast on Thursday, September 30 – TOMORROW!

Erin is co-presenting a webcast on *Working Around the Workload* on Thursday at noon EDT. It's all about how to run everything apart from the main workload, without affecting the workload itself. You can register for the webcast and read the abstract here.

SQLskills Insider Sessions!

We've been thinking a lot about work-life balance and the continued stresses from the ongoing pandemic and we decided that we would push our upcoming classes to Spring 2022. A huge thank you to all of you who were registered to either move to Spring or change to recorded courses – no one chose a refund. And, it makes sense too – we just recorded a round of our most popular courses (IEPTO1, IEPTO2, and IEAzure++) this past Spring on SQL Server 2019 so there's great, fresh content already available.

So, what can you do this Fall? Dedicate some time to staying up-to-date! Consider taking some classes, look at upcoming FREE conferences like PASS and Ignite, but above all, make time for YOU! Make time for learning and staying healthy, both mentally and physically! In these stressful times, there's no better investment!

For SQL Server (and many other things), it can be great to create a dedicated and regularly scheduled "block" of time for that task. And, we will help fill it with content! Whether you attend one or all of our **FREE SQLskills Insider Sessions**, you should still consider some regular blocked time for YOUR learning and YOUR career.

As for SQL Server, we're going to run different topics or a theme for a few weeks and then we're going to follow them up with an open Q&A. These LIVE, interactive sessions are completely FREE and will always happen at the same time – beginning at 10:00am Pacific Time. We're going to put together a web page with more details and a registration link so you can get reminders for the specific sessions in which you want to attend. The only requirement – we request you be a SQLskills Insider (which you already are if you're reading this!)

These SQLskills Insider Sessions will begin on Wednesday, October 6 and will run almost every Wednesday (unless it conflicts with another event or US holiday) and run for up to 75 minutes of content (sometimes they'll be shorter), up to 15 minutes of Q&A (sometimes a few of us might stay longer) and will be a great way to remind you of things you may already know or just allow you to stay fresh.

To start things off, Kimberly is going to do a series on statement execution, stored procedures, caching (the good and the bad), and recompilation; she'll follow up this series with an Ask Anything Q&A.

- All of these will be 10:00am PT until 11:30 PT (75 mins presentation, 15 mins Q&A)
- Main page with all the details is here
- **SQLskills Insider Session: Statement Execution and Caching**
  - Wednesday, October 6, 2021 – details/registration
- **SQLskills Insider Session: Plan Cache Pollution and Clean Up**
  - Wednesday, October 13, 2021 – details/registration
- **SQLskills Insider Session: Stored Procedures, Part 1 – Execution and Caching**
  - Wednesday, October 20, 2021
- **SQLskills Insider Session: Stored Procedures, Part 2 – Caching Problems and Recompilation**
  - Wednesday, October 27, 2021
- **SQLskills Insider Session: Ask Anything Q&A – Caching and Recompilations**
  - Wednesday, November 3, 2021

We will skip an Insider Session on November 10 as the FREE PASS Conference will be running that day (Wednesday, November 10 through Friday, November 12th!

Our thought: block around this time – EVERY Wednesday! **Dedicate 2 hours a week to Improving Your SQL skills with SQLskills**! Join us when it makes sense (hopefully always! ☺) or, read blog posts, review topics of interest, go spelunking in the documentation on a topic that's always interested you. Stay fresh, stay current – stay ahead of the competition!

Block the time NOW and we'll send you more details in the next newsletter!

PASS Data Community Summit

Both Erin and Kimberly are presenting full-day workshops at the Summit in November:

- Erin: Finding Problems and Stabilizing Performance with Query Store
- Kimberly: Solving Query Performance Problems with Statistics

The main registration page is here.

Book Review

The most recent book I've read (Jess Kidd's *Things in Jars)* was a rare disappointment for me, so here's one from the 2009 archives that does a much better job of bringing old London to Life. I give you David Liss' *A Conspiracy of Paper*. This is excellent historical fiction set at start of the stock market in London, 1719. The South Sea Company is in its heyday and brokers are called 'stock jobbers'. The book follows a Jewish ex-boxer, who is now a thief-taker and investigator unravelling the mystery of his father's murder, intertwined with the workings of the fledgling stock market. Highly recommended! (And also the sequels with the same protagonist…)

The Curious Case of…

This section of the newsletter explains recent problems we've helped with on client systems or been asked about over email or #sqlhelp; they might be something you're experiencing too.

Last week I had a question in email about how to find what update transactions a user had performed from the transaction log. I was going to write a post about it and then realized I already had one from back in 2015. You can read my explanation of how to look in the log for everything a user has done here…

Ponderings…

When performance is poor, a common reaction is to blame "the database;" it's got to be SQL Server itself that's causing the problem... And, of course, there could be a myriad of things that could be wrong – be it incorrect indexing strategies, poor coding practices, statistics and/or cardinality estimation issues, index fragmentation, parameter sniffing, improper configurations, poor database design – the list goes on and on and on.

Even when starting the troubleshooting process, it's common to use a step-wise methodology that begins with wait statistics analysis; the focus is to find out where (and ultimately, why) SQL Server is waiting. And, while this is usually a good strategy, it's not always SQL Server that's the problem.

Moreover, it's not typical to focus on the application and ask whether the application is at fault. The problem could be that the application is asking for too much data and driving a very expensive workload on SQL Server. SQL Server is only reacting, and it can't keep up with all the poorly-designed application requests.

To help you think outside the box, here are some things to consider about the way your application queries the data in SQL Server. These could be adding unnecessary load and negatively stressing CPU, memory, and I/O:

- **Processing:** For the data that's being pulled from SQL Server, is the application processing the data one row at a time (commonly called RBAR, or row-by-agonizing-row, processing)? If so, this means that SQL Server has a thread waiting for the application to acknowledge the data sent to it, and can lead to high *ASYNC_NETWORK_IO* waits. The application would be better to cache the incoming data locally, reply back to SQL Server as soon as possible that it has the data, and then they both continue processing.
- **Selections:** Is the application going to filter the data locally before using it or displaying it? In that case, it should be waaaay more efficient to push the predicate up to SQL Server (called a ***selection***) and have the minimum possible data returned to the application. SQL Server is very, very good at filtering data, given the right indexes to support the filter predicates. It's generally a very good idea to reduce the amount of data sent to the client!
- **Projections:** Are all the table columns being returned absolutely necessary? Are you trying to build a "one size fits all" dialog? By using a targeted *SELECT* list (called: a ***projection***) rather than just *SELECT **, this cuts down on the data being processed and returned. And, with fewer columns requested, SQL Server often has more optimal ways of getting to this data – also improving performance.
- **OSFA:** Is that dialog one where every possibly option is given and then only one "generalized" stored procedure sits behind it? This is very common and often results in horrible performance and terrible parameter sensitivity problems. Kimberly wrote a really good blog post about this here.
- **Ordering:** Does the data being returned really need to be sorted with an *ORDER BY*? If not, this might cut out a sort operation. Often sort operations can be expensive as they may end up requiring a costly sort-spill to tempdb.
- **Just in case:** Can the *SELECT* be postponed until it's really required? If an application is issuing a *SELECT* "just in case" the user clicks an application button, then it might be wasted processing. It may be better to wait until the button is actually pushed before issuing the *SELECT*, removing all of the processing when the button is not pushed.
- **Consider caching:** If the same data is being queried again and again, cache it locally (or mid-tier) and only issue a new *SELECT* when the data changes. This is ideal when data does not change frequently or if up-to-the-minute data is not required.
- **ORMs and statement builders:** Be careful here… potential for plan cache bloat and lots of compile time (if the statements are truly dynamic) or potential for parameter sniffing problems if the statements use *sp_executesql* behind the scenes. Sadly, there's isn't a simple solution to the ORM problem. Kimberly talks a lot about this in her Pluralsight course: *SQL Server: Optimizing Ad Hoc Statement Performance*.

These are just a few things to think about when analyzing how an application uses SQL Server. Making some of these changes can have a profound effect on the amount of work SQL Server has to do, especially if a single change in the application query logic is multiplied by hundreds or thousands of instances of the application running simultaneously.

Call to action: If your SQL Server is stressed from excessive application queries pulling large amounts of data, instead of first trying to tune those queries, go to your developers and ask them whether each query is really necessary, and whether it has to pull back as much data as it is. And, if you're the developer; ask yourself what you can do to streamline your application's requests to SQL Server. You'll be amazed at the difference you can make!

#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 theme for the TBT this time is backups:

- My Pluralsight course: *SQL Server: Understanding and Performing Backups*
- My TechNet Magazine article from 2009: *Understanding SQL Server Backups*
- My SQL Server Magazine article from 2011: *Advanced Backup and Restore Options*
- My TechNet Magazine article from 2009: *Recovering from Disasters Using Backups*
- Backup blog posts from our accidental DBA series
- Blog post: *New script: How much of the database has changed since the last full backup?*
- Blog post: *Importance of having the right backups*
- Blog post: *Importance of validating backups*
- My blog post category on *Backup/Restore*

Posts since the last newsletter:

- Paul: *SQL Server Latches – Other Latches To Know About* (SQLperformance.com)
- Kimberly: *Calling all data "wranglers" / developers – and those that want to know more about SQL Server!*

I hope you find these useful and interesting!

Video Demo

This week Erin walks us through a database setting you don't usually want enabled, auto-close, and reviews what to look for and potential issues it can cause.

The video is about 7.5 minutes long and you can get it in WMV format here.

The demo code is here.

Enjoy!

Upcoming SQLskills Events

We’ve moved all our classes to early 2022 – schedule details will be published shortly.

With our new streaming system, you can now choose to attend a live, online event or purchase a recording to watch at your leisure, either individually or as part of a bundle. And all attendees of live events get lifetime access to the class recordings too!

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

LIVE, Online Immersion Events:

Spring 2022

- Details coming soon.

You can get all the details on our training options page or just go directly to our new shop.

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