(November 9th, 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 Redmond, WA, where I'm currently moderating Kimberly's PASS Data Community Summit pre-con workshop on statistics (not my favorite topic but definitely one of hers!)

Registration is open for the remainder of our 2021 SQLskills Insider learning sessions every Wednesday – see here!

Take care and stay safe out there!

SQLskills News

SQLskills Insider Sessions!

Our first batch of SQLskills Insider Sessions are done and we had a great / interactive audience. It was really nice to reconnect with so many of you!

We wrote about the motivation behind our Insider Sessions and you can register for just one, or the entire series here: https://www.SQLskills.com/iSessions.

The main page with all the details is here, including the five sessions that Kimberly has already completed around execution, caching, and recompilation.

Coming up (skipping November 10 as it's PASS this week):

- **SQLskills Insider Session: Troubleshooting Deadlocks with Jonathan Kehayias**
  - Wednesday, November 17, 2021 – details
- **SQLskills Insider Session: Query Store Best Practices with Erin Stellato**
  - Wednesday, December 1, 2021 – details
- **SQLskills Insider Session: Understanding Statistics: The Histogram with Kimberly L. Tripp**
  - Wednesday, December 8, 2021 – details
- **SQLskills Insider Session: Azure SQL Overview with Tim Radney**
  - Wednesday, December 15, 2021 – details
- **SQLskills Insider Holiday Party…**
  - Wednesday, December 22, 2021

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, register online, and each week we'll send you a meeting link with joining instructions. And, stay tuned, we'll announce the rest of the year's line-up of topics and speakers soon!

Book Review

The most recent book I've read is Frederick Kempe's *Berlin 1961: Kennedy, Khrushchev, and the Most Dangerous Place on Earth*. From Amazon: "Much has been written about the Cuban Missile Crisis a year later, but the Berlin Crisis of 1961 was more decisive in shaping the Cold War-and more perilous. It was in that hot summer that the Berlin Wall was constructed, which would divide the world for another twenty-eight years. Then two months later, and for the first time in history, American and Soviet fighting men and tanks stood arrayed against each other, only yards apart. One mistake, one nervous soldier, one overzealous commander-and the tripwire would be sprung for a war that could go nuclear in a heartbeat.

On one side was a young, untested U.S. president still reeling from the Bay of Pigs disaster and a humiliating summit meeting that left him grasping for ways to respond. It would add up to be one of the worst first-year foreign policy performances of any modern president. On the other side, a Soviet premier hemmed in by the Chinese, East Germans, and hardliners in his own government. With an all-important Party Congress approaching, he knew Berlin meant the difference not only for the Kremlin's hold on its empire, but for his own hold on the Kremlin. Neither man really understood the other; both tried cynically to manipulate events. And so, week by week, they crept closer to the brink."

Highly recommended!

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 from someone who noticed that while executing a scan of a heap using NOLOCK, there was a BULK_OPERATION lock held on the heap for the duration of the scan. The question was why is the BULK_OPERATION lock needed, as surely there's no way for the NOLOCK scan to read a problematic page? You can read all about the answer here…

Ponderings…

(From me this time – enjoy!)

One of the discussions I've had recently was around why SQL Server doesn't have a way to automatically remove index fragmentation. I've had similar discussions over the years, both inside and outside Microsoft and the SQL team, on SQL Server's lack of automatic facilities to do things that involve a bunch of resource consumption.

While I was on the SQL team, some of these were seriously considered, but I was always very cautious about them. Here's a (non-exhaustive) list of questions to consider, using index fragmentation removal as an example:

- What time of day/week should the process run?
- What thresholds to use for 'do nothing', 'reorganize', 'rebuild'?
- What if there's no disk space to do a rebuild? How to check that? What if a rebuild fails because of lack of space: should a reorg be done instead?
- What if the database has to autogrow to accommodate the new index? Should some kind of shrink be run afterward?
- How much parallelism should an index rebuild use?
- Who manages the transaction log size if the table is really large? Should we automatically take log backups? Where should they be backed up to and how should they be named? What if there's no space for the log backups? What about space on the log shipping secondaries? And the time required to replay the log on the log-shipping secondaries?
- What about the impact of extra transaction log on database mirroring or availability groups? Both in terms of the huge SEND queue preventing log truncation, and the huge REDO queue preventing the mirror/AG replica coming online quickly?
- What if a cluster failover occurs during the operation and the rollback prevents the database coming online within the company's RTO?
- What if there's no Enterprise Edition? Then there's no online index rebuild and so the application/users are blocked while the table is SCH_M locked while the offline index rebuild occurs. Does that imply always do a reorg in Standard Edition?
- What if we know that certain indexes may become fragmented but that doesn't affect the workload, so ignore those indexes?

My point is that it's a lot harder than you think to just put some automatic behavior into SQL Server. Given the capacity for resource usage, if it's wrong once and causes a performance problem, or excessively large backups, or something else negative, it will get a bad reputation and people won't use it.

Call to action: It's very tempting to blame Microsoft for not having an automatic method of doing X or Y. However, when you stop to consider the complexity of the problem and the sheer diversity of environments that the automatic feature would have to work in, it should come as no

surprise that Microsoft tends to *not* invest engineering time in such features. Very often there is a third-party solution to the problem you'd like to see solved – for instance, Ola Hallengren's wonderful database maintenance solution that you can customize to achieve much of what I've considered above. Use a little bit of ingenuity and you can solve the problem yourself, for your environment, and in a way that best suits your application.

#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 FILESTREAM, as that was the most recent Curious Case. Here are some useful links:

- My whitepaper: *FILESTREAM Storage in SQL Server 2008*
- My blog posts:
  - *FILESTREAM directory structure*
  - *FILESTREAM directory structure – where do the GUIDs come from?*
  - *FILESTREAM garbage collection*
- *FILESTREAM* blog category

Posts since the last newsletter:

- Paul: *Tracking Synchronous Statistics Updates* (on SQLPerformance.com)
- Paul: *The Curious Case of… the BULK_OPERATION lock during a heap NOLOCK scan*

I hope you find these useful and interesting!

Video Demo

In this demo video, Jon takes a look at the SID associated with Windows Logins and Groups in SQL Server and how those are mapped to different operations.

The video is about 9 minutes long and you can get it in MP4 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