

(October 12th, 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. The weather just turned too nasty up at camp so I gave up on October 5 after 109 straight nights of camping up there!

Registration is open for our next few SQLskills Insider learning sessions every Wednesday – see [here](#)!

Take care and stay safe out there!

SQLskills News

SQLskills Insider Sessions!

Our first SQLskills Insider Session was last week and 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>.

We have 3 more sessions in Kimberly's caching and recompilation series with a final open Q&A after all 4 are complete.

- Main page with all the details is [here](#)
- **COMPLETED: SQLskills Insider Session: Statement Execution and Caching**
 - Wednesday, October 6, 2021 – [details](#)
- **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 – [details/registration](#)
- **SQLskills Insider Session: Stored Procedures, Part 2 – Caching Problems and Recompilation**
 - Wednesday, October 27, 2021 – [details/registration](#)
- **SQLskills Insider Session: Ask Anything Q&A – Caching and Recompilations**
 - Wednesday, November 3, 2021 – [details AND to send your remaining questions](#)

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, 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!

PASS Data Community Summit

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

- Monday, November 8 Kimberly is presenting: [Solving Query Performance Problems with Statistics](#)
- Tuesday, November 9 Erin is presenting: [Finding Problems and Stabilizing Performance with Query Store](#)

The main registration page is [here](#).

Book Review

I haven't finished any books since the last newsletter so here's one from the 2010 archives: Ross King's [Ex-Libris](#). It's an excellent novel set in 1660 in London and deals with a book seller who gets involved in investigating the existence of a rare hermetic manuscript. It's full of period detail and lots of info on rare books from that time, a book about old books and libraries is always one of my favorites. Some excellent twists and turns keep the book interesting. Recommended (plus his earlier non-fiction work I read in 2009: [Brunelleschi's Dome: How a Renaissance Genius Reinvented Architecture](#).)

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 was explaining to someone how to resolve SQL Server code call stacks without having to download symbols, as part of a debugging exercise, and I remembered Jonathan wrote a post about the SQLCallstackResolver tool on Github. You can read all about the tool in Jonathan's blog post [here](#)...

Ponderings...

(From me this time, but I didn't have a chance to write a new ponderings while packing up camp, so to go with the corruption video below, here's a Ponderings from 2016 that explains a nifty trick you might be able to use to work around transaction log corruption – enjoy!)

This time I want to talk about log file corruption and what you can do to work around it – in this case, corruption that prevents a log backup running.

Sometimes you might see a log backup fail with an error, but a full backup succeeds, and so does a *DBCC CHECKDB*. What's going on?

Let's say there's a corruption in a portion of the transaction log in a VLF, and that portion of log hasn't been backed up since the last log backup.

A log backup has to back up the transaction log records generated since the last log backup – so it will try to back up the corrupt portion and fail.

A full database backup, on the other hand, only has to back up enough of the log to allow the database to be restored and recovered to a transactionally-consistent point. In other words, it only requires the transaction log back to the beginning of the oldest active transaction at the point that the data-reading section of the full backup completes. So it might not have to back up the corrupted portion of the log and it succeeds.

A *DBCC CHECKDB* operation uses a database snapshot to get a transactionally-consistent view of the database on which to run consistency checks. When the database snapshot is created, crash recovery is run on it to make it transactionally consistent. That requires the same amount of log as if a full backup was taken – back to the beginning of the oldest active transaction at the time the database snapshot is created. So if the corrupt portion of the log isn't needed, the *DBCC CHECKDB* will succeed without any problem

The question then becomes: how to recover from this situation?

Assuming that the database files are intact, there is a simple solution. While this solution isn't recommended under normal working conditions, it works here because the log backup chain is *already* broken. This solution is to throw away part of the log chain and essentially begin a new log backup chain. Here's what to do:

1. Stop all user activity in the database
2. Switch to the SIMPLE recovery model (breaking the log backup chain)
3. Perform a manual checkpoint of the database (which should remove the requirement that the damaged portion of log must be backed up)
4. Switch back to the FULL recovery model

5. Take a full database backup, or a differential database backup (thus starting a new log backup chain)
6. Start taking log backups

You might want to manually shrink and grow the log file in between steps 2 and 3 too – in case the log file is on a damaged portion of disk. You also will need to do some root-cause analysis to determine why the corruption occurred in the first place, and to take preventative measures to stop it happening again.

Call to action: If you see weird log corruption messages, try to work around it this way rather than doing something drastic like detaching the database and deleting the log file, or immediately jumping to running *EMERGENCY*-mode repair.

#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.

One of my favorite sessions at SQLBits last year was my Mythbusters talk, so that's the theme for #TBT this time. Here are some myths resources for you:

- The very first Pluralsight course we (I) recorded: [SQL Server: Myths and Misconceptions](#) that has almost four hours covering a plethora of SQL Server myths.
- My blog post series from April 2010 – *A SQL Server DBA Myth A Day* – and [this blog post](#) has links to all 30. Some of the later ones had 20+ myths busted in one post, as I got a bit carried away ☺
- The same blog post series, but organized as a [handy PDF file](#).

Posts since the last newsletter:

- Jon: [SqlClient Configurable Retry Logic – ICYMI](#)
- Jon: [SQL Server Storage Best Practices: Choosing Storage Options](#) (on the SentryOne blog)
- Kimberly: [Calling all data “wranglers” / developers – and those that want to know more about SQL Server!](#)
- Paul: [10 SQL Server Performance Tuning Best Practices](#) (on the SentryOne blog)
- Tim: [Native SQL Server Backup Types and How To Guide](#) (on the SentryOne blog)

I hope you find these useful and interesting!

Video Demo

For this newsletter I present a demo video from my Pluralsight course [SQL Server: Advanced Corruption Recovery Techniques](#). If you don't have any backups and you're faced with data loss, think about ways to reconstruct the data from other sources. As an example, if you have damaged data pages at the leaf level of a clustered index, repair will delete them, but you might be able to get some or all of the to-be-deleted data back from nonclustered indexes on the table. Check out the demo where I show how to do that.

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