

(April 14th, 2026)

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!

A run-of-the-mill newsletter this time with no explosive news, just more of the same SQLskills goodness you all love – enjoy!

SQLskills News

We've replaced our back-end video streaming system with one that has a lot more end-user functionality, is more robust, and has features like chapters so I can provide indexes into the long Immersion Event daily recordings. The transition has already been made and should be seamless for all viewers. I'll be working through our library to add the chapter indexes over the new few months.

Book Review

Two WWII books for you this time...

The first is William L. Shirer's [*The Rise and Fall of the Third Reich: A History of Nazi Germany*](#). I started reading my print copy of this and got bogged down due to its density, so decided to spend an Audible credit on it – all 50 hours of it! Incredibly interesting, hugely detailed, and gripping! After gleaning a lot of knowledge about Hitler, Nazi Germany, and the German prosecution of WWII, from the earliest days of history classes aged 11 to Beevor's [*The Fall of Berlin 1945*](#) that I listened to and reviewed here last year, I decided to bite the bullet and read what's widely-acclaimed as the definitive book on Nazi Germany. From Amazon: "A modern classic and the definitive study of Adolf Hitler, the rise of Nazi Germany, the Holocaust, and World War II by the foreign correspondent who saw firsthand how the Nazi movement swept Europe. William B. Shirer presents nuanced profiles of major and minor figures as well as exacting details of the political movements, wartime strategy, and atrocities of the Third Reich. This epic work is required reading to understand modern history and how fascist dictators can captivate a population to wicked ends." Highly recommended!

The second is Neill Lochery's [*Lisbon: War in the Shadows of the City of Light, 1939–1945*](#). I've never really read much about Portugal, but I did know it was a neutral power during WWII. As well as being about Lisbon during the war (foreign agents, refugees), the book is heavily focused

on António de Oliveira Salazar and his running of Portugal during that time and how he deftly balanced the opposing sides to the benefit of his country. I had no idea that tungsten (wolfram) was so heavily required for armaments, and that Portugal was the major supplier of Germany, with all the complications that causes, as well as questions about the morality of profiting from payments in stolen gold. Extremely interesting and recommended!

Ponderings...

(I've been run off my feet the last week and haven't written a new editorial, so I thought I'd run this old gem past you – enjoy!)

I've had some discussions other the past few weeks about why it's necessary to run consistency checks regularly (for instance, I recommend a full *DBCC CHECKDB* every week if possible). The problem is the complete unpredictability of corruptions and the fact that as soon as you've examined a page for corruption and moved on, it could immediately become corrupt.

Here's a story. Imagine the Redmond Police Department has a Traffic Control Center and employs Officer Paul. Paul's job is to sit in the control room looking at video images of traffic junctions in the city, waiting for a traffic accident to occur and immediately notifying dispatch that help is required at whatever junction.

It's obviously a really boring job, so it's not uncommon to find Officer Paul staring slack-mouthed at his screen with drool starting to slowly puddle on his desk...

Anyway, imagine that today Officer Paul looks at the image for junction #1 for ten seconds. Nothing happens. He presses the button to watch junction #2 for ten seconds. Nothing happens. He presses the button to watch junction #3. As soon as he starts looking at junction #3, there's an accident at junction #2. He won't know until he cycles through all the junctions and arrives back at #2 again.

There's no way for Officer Paul to get any guarantee that an accident won't happen at a junction he's just looked at, once he stops looking.

This is an analogy of how *DBCC CHECKDB* examines the pages in a database. Each page is read from disk once by *DBCC CHECKDB*, and as soon as the page has been read from disk, it could become corrupt on disk, and you won't know until the next time *DBCC CHECKDB* runs. The only thing *DBCC CHECKDB* can tell you is that at the time it read all the pages in the database, there were no corruptions. It doesn't provide a guarantee that at the time it finishes, there are no corruptions in the database, as it can't possibly know that.

Another variant of this is thinking that if a backup completes, it will restore perfectly.

No.

The I/O subsystem is what I call an 'equal-opportunity corruptor' :-) If it has problems, it will happily corrupt anything stored on it.

You need to have multiple copies of backups and you need to test a full restore sequence regularly. We always recommend having the database option `PAGE_VERIFY CHECKSUM` turned on and then when you perform a backup, you should always use the `WITH CHECKSUM` option. Finally, run a `RESTORE VERIFYONLY ... WITH CHECKSUM` to make sure the backups you are taking are valid. Even better is to take them to another system and restore them.

Call to action: You will all see corruption at some point in your careers, so expect it and prepare for it. You need to run regular consistency checks and you need to regularly test your backups are valid. There are no guarantees around corruption, only increased peace of mind.

(PS 2026: I still stand by this advice, whether you're running on-prem, in the cloud, de-duped, AGs, whatever. Corruption can strike anywhere!)

The Curious Case of...

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

A few years ago Kimberly was working with a client who was having performance problems, with inserts into a particular table taking hundreds of milliseconds for each row being inserted, and index maintenance taking a long time to run. The client could show that all the nonclustered indexes were being used, so couldn't figure out what to do.

After starting to investigate the index definitions, Kimberly noticed that there were duplicate indexes on the table, but all the indexes were showing up as used in the monitoring being performed using the `sys.dm_db_index_usage_stats` DMV.

It turned out that SQL Server can pick any of the duplicate indexes for a query plan, and the client had code that was recompiling and causing different duplicate indexes to be used, hence the output from the index usage monitoring.

Kimberly used her duplicate-index-finding code from her blog post [here](#) and removed 16 duplicate indexes from the table, and then did some manual consolidation to remove a further 4 indexes, leaving a total of 12 nonclustered indexes. That fixed the insert performance problem and removed 20 indexes from having to be defragged during index maintenance too.

Bottom line: Be careful to prevent duplicate indexes being created (remember, some indexes might not look like they're dupes but internally, they are!), as SQL Server will happily let you

create them. Regularly monitoring for unused and duplicate indexes is a good way to help keep your workload and regular maintenance performing well.

#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 IEVLT class](#)

And some of my recent blog posts:

- [SQL101: Phishing attacks](#)

I hope you find these useful and interesting!

Video Demo

In this demo video, Jonathan shows you how to extract a DLL from SQL Server for a SQLCLR assembly for use with a free tool named [ilSpy](#) that can decompile the assembly to allow you to view how the CLR code was written.

The video is about 5 minutes long and you can get it:

- In WMV format [here](#)
- In MOV format [here](#)

And the demo code is [here](#).

Enjoy!

SQLskills Training

We have no plans for live, public classes in 2026, but we've already released the first of many new courses coming over the next year! And of course, all our recorded courses from the last few years are still as relevant as ever.

To help your manager 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 manager explaining why SQLskills training is worthwhile](#)
- [Community blog posts about our classes](#)

You can get all the details in our [shop](#).

Summary

I hope you've enjoyed this issue – I really enjoy putting these together. If there's anything else you're interested in, I'd love to hear from you - [drop me a line](#).

Thanks,
Paul