(June 17th, 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!

So one thing I completely underestimated when we started this whole moving process was my books. I have a lot of books – more than 2,500 physical books, and many of them hardbacks and large reference books (like the excellent DK series). Which means they're heavy! I've been making runs up to Camp Savage using my dump trailer to haul 1,500lbs of books at a time in 18x18x12" boxes (so they're liftable). Today is the fourth such run once I publish this newsletter!

An onerous process, to say the least. But on the excellent flip-side, it lets me look at all the books and rediscover some I'd bought but forgotten about – like Neal Stephenson's *Termination Shock* from five years ago that I found yesterday and am excited to start this evening!

Summer Sale Time!

As I said in the last newsletter, time flies! Summer officially starts this coming Sunday so on I've be putting our Blackbelt bundles on sale for a week or two:

- **$699 for one-year access**
- **$1,299 for lifetime access**

More than 158 hours of incredible training – see our shop for details!

If you received an email from me back in March about an upgrade on your courses/bundles, that offer still holds – just let me know!

PASS Data Summit Workshop

I'm excited to announce that I'll be presenting at the Pass Summit in November! My pre-con workshop is my absolute favorite to teach *Performance Tuning with Waits and Latches* and I'll also be presenting a session on Advanced Data Recovery Techniques. Bonus is that I'll have Kimberly with me!

All the details are available on the Pass Summit website. See you there!

Book Review

I've read/listened to 18 books since the last newsletter (and just passed 100 books for 2026!) and here are two I'd like to recommend to you.

Firstly, Alfred Lansing's *Endurance: Shackleton's Incredible Voyage*. This is an abridged (removing all the expedition pre-amble) version of the book, detailing the crushing of the Endurance, the groups' time living on the Antarctic ice, and the incredible small-boat journey from Elephant Island to South Georgia, for which Shackleton is best known. Now on to Fiennes' masterly biography of the man himself.

Secondly, Ranulph Fiennes' *Shackleton: The Biography*. I decided to learn more about Shackleton the man, and his other expeditions apart from the famous one where he undertook the rescue journey from Elephant Island to South Georgia. Very interesting to read it from the perspective of another distinguished polar explorer, not just a random armchair biographer. Excellent!

Ponderings…

(I have an editorial I'm working on around reliance on AI and how dangerous it can be, for many reasons, but in the meantime, here's one from five years ago that I can't emphasize enough!)

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?

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. Just last week I saw a case 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 two billion-record table, using the NVARCHAR (3) column (let's say with half "YES" and half "NO" values) would take (1 billion x 8) + (1 billion x 6) + (2 billion x 2) =

13GB, compared with 2 billion * 1 = 1.86GB. That's a huge difference. Some might say that in the grand scheme of things, 11 extra GB is negligible, but if many more columns have inefficient design, that adds up quickly. Think of extra disk space, extra log space, extra buffer pool space, lower data density, larger backup size, longer restore times, and so on. Kimberly has a Pluralsight course that covers this 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 (less so with advances made in SQL Server 2019). Understanding that various allocation bitmaps need to be updated in memory whenever a page or extent is allocated explains why the contention occurs 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 I blogged about) 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.

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.

Very apt considering my PASS announcement: why does your wait statistics script filter out a bunch of wait types? You can read through my explanation here…

#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 transaction log is the theme for #TBT this time. Here are some transaction log resources for you:

- My 8-hour Pluralsight course: *SQL Server: Logging, Recovery, and the Transaction Log*
- My *transaction log blog category*
- My TechNet Magazine article on *Understanding Logging and Recovery in SQL Server* (the formatting is messed up a bit on that web page)
- The *WRITELOG* page from my waits library

And coincidentally, the last blog post I published was also on the transaction log:

- *The SQL Server Transaction Log, Part 4: Log Records*

I hope you find these useful and interesting!

Video Demo

And continuing with the wait statistics theme for this newsletter, here's a demo from Tim's Pluralsight course *SQL Server: Understanding and Using Azure SQL Database*. In the demo, Tim shows how to change my scripts for capturing wait statistics and I/O statistics so they work on Azure SQL Database.

The video is just under 3 minutes long and you can get it in MP4 format here.

The demo code is available here.

Enjoy!

SQLskills Training

We have no plans for live, public classes in 2026, but we've already released two new courses this 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