(February 12th, 2018)

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.

Quick Tips for our Insider friends!

Hey Insiders,

We're presenting another live, online Immersion Event! Due to the popularity of Kimberly's delivery of the new *IEVLT: Immersion Event on Very Large Tables: Optimizing Performance and Availability through Partitioning* in January, and the fact that many of you told us that you didn't have a chance to attend, we're running it again in March!

This will be delivered live via WebEx on March 20-22 (roughly 12-13 hours of content plus open Q&As; about the same as two full workshop days!). It's priced at only US$895. See here for all the details, including the incredible feedback from the January class on how much they loved it. And, for all of our future, live, online events YOU – our newsletter subscribers – will always get an exclusive discount for the first couple of weeks!

Note: you can get all the prior Insider newsletters here.

SQLskills News

Erin recorded a RunAsRadio podcast on Query Store recently – for details see their site at http://runasradio.com/Shows/Show/570.

The first batch of 2018 US classes are open for registration! In 2018, we're offering our usual Immersion Events on Performance Tuning (IEPTO1 and IEPTO2) and for the Accidental DBA (IE0), as well as PowerShell, Azure, Clustering and Availability Groups, BI strategies, BI security, and Practical Machine Learning. See here for our 2018 Immersion Event class schedule.

And we're also coming back to Europe in 2018! We're bringing four of our Immersion Events to London in September: IEPTO1 and IEPTO2, plus our new classes: **IEAzure** (on Azure and Azure VMs) and **IECAG** (on clustering and availability groups). See here for details.

Also for our European friends, Erin will be presenting a pre-con workshop at SQLBits in the UK in February – see here for details and check soon as it's almost sold out.

Our Spring SQLintersection conference is also fast approaching at the end of March, and we have a phenomenal line up of workshops, sessions, and speakers. Check out this blog post for all the details, and use the discount code 'sqlskills' when you register to save $50.

Finally, even if you can't join us in person, I've put out a call for 2018 remote user group sessions. In 2017, we did more than 100 of these around the world and we have set up more than 50 for 2018 already! If you'd like one of us to present for your user group, check out my blog post here.

Book Review

The most recent book I've read is David Enrich's *The Spider Network: The Wild Story of a Math Genius, a Gang of Backstabbing Bankers, and One of the Greatest Scams in Financial History*. Great book! This is the story of the spectacular rise and equally spectacular fall of the interest-rates derivatives trader Tom Hayes. He was involved in wholesale manipulation of LIBOR-like exchange rates for the financial gain of his trades while working in Japan, in a spider's web of other traders and brokers. Some say he was the spider, and others say he was the fly, as you'll see in the book. It's quite a page-turner, and very readable, with very little arcane financial info. Highly recommended!

The Curious Case of…

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

A couple of weeks ago, Tim and I were helping a new client recover from a SAN failure. Through a long story, they didn't have any usable backups, and so Tim was working through their database, scripting out as much as possible – a long and tedious process.

The corruption was centered in the *sys.syrscols* system table, which stores Storage Engine column metadata, and so we were concerned that some of the client's tables would be inaccessible.

Then I had an idea. As the corrupt page was *so* low in the primary data file, a pretty old backup probably had the correct version of that page, as newer tables would have metadata stored on data pages later in the file. My idea was to do a hex editor page swap from the old version of the database to the newer, corrupt database (along the lines of the boot page hex editor page swap I've popularized – see here).

However, it failed, because the physical location in the file of the corrupt page was completely inaccessible by the I/O subsystem. Any attempt to read or write that page failed, using any mechanism: *DBCC WRITEPAGE*, *DBCC PAGE*, *BACKUP*, or even a file copy at the Windows level.

This is an extremely rare occurrence IMHO; I've only ever seen these happen 2 or 3 times in my entire career working with Windows and SQL Server. So that left scripting out, which luckily worked perfectly for all the client's critical tables.

Bottom line: Sometimes there's absolutely no alternative to just scripting out data from a corrupt database. Make sure you know how to do that, as a combination of problems and bad luck could happen in your environment.

Paul's Ponderings

(This is adapted from an article I published in the newsletter back in February 2012 as I wanted to write about the same topic because it's come up a few times recently.)

One of the problems I keep hearing about in the SQL community is application vendors that use SQL Server as the data platform and have erroneous configuration recommendations, requirements, or maintenance routines (like mandating that shrink is run on the application database regularly so space isn't wasted… sigh).

Sometimes this is because the application has been ported from another platform, like Oracle, and the application developers assume that SQL Server behaves the same way. And sometimes this is because the application developers wrote their code for SQL Server but have no idea how to properly design a database, T-SQL code, or app-layer code.

This is bad for many reasons, including:

- Myths and misconceptions about SQL Server are perpetuated by "authoritative" sources (i.e. these vendors)
- Vendor clients without a DBA may not know any better and so may incur performance problems from implementing the erroneous configuration and maintenance routines
- The DBAs at some vendor clients may have their knowledge called into dispute by their managers because the vendor disagrees with them

The big problem comes when the vendor doesn't really understand the basis for the recommendation that they're making. It can be extremely hard to convince such a vendor that their advice is incorrect – I've even had to resort to the dreaded *"I wrote/used-to-own that code – your understanding of it *is* incorrect"* argument to make people see sense.

Sometimes it can be SQL Server tools vendors who get it wrong – and I've been bitten by this myself in the past. One of our clients contacted me saying that a tool they'd bought was complaining about some indexes I'd created on the client's system, saying they were duplicate indexes that should be removed. Understandably I was really annoyed by this as the tool was giving my client the impression that I'd done something wrong. It turns out that the tool was not considering *INCLUDE*d columns when flagging an index as a duplicate. I blogged about my experience here and the problem was subsequently fixed in that tool.

Even Microsoft, as a vendor of applications that use SQL Server, has made mistakes sometimes too. I taught the SQL Server part of the SharePoint MCM certification since it started (Jonathan

and Kimberly taught some of it too) and in 2011 I was very dismayed to read the new SharePoint 2010 SQL Server maintenance whitepaper. It was full of poor advice so I rewrote it for Microsoft and they republished it. If only all vendors would accept advice and help to rationalize the information they give their clients!

Call to action: Next time you come across some guidance from a vendor that is clearly wrong or misleading, call them on it! Education is the key to getting things like this fixed. And feel free to pull me into any email threads with vendors (or colleagues) who won't listen to reason – I'll help where I can.

Glenn's Tech Insights

This section of the newsletter highlights recent news and views from the hardware and Windows worlds that we think will be interesting to SQL Server community members.

Detailed Background Information About Meltdown and Spectre

Jon Masters (who is a Computer Architect at Red Hat) has put together a 90-slide PowerPoint deck called "Exploiting modern microarchitectures: Meltdown, Spectre, and other attacks". The PDF version of this deck is available here.

There is a lot of detailed background information about CPU microarchitectures, how CPU caches work, and how branch prediction and speculative execution work. There are also a lot of clear, useful diagrams and code examples. This is a very deep, but effective deck, which makes it very useful for later reference and study.

Dell EMC Launches Three New AMD EPYC-based PowerEdge Servers

In a big win for AMD, Dell EMC has launched three new PowerEdge server models that all use AMD's EPYC processor. The first one is the Dell EMC PowerEdge R6415, which is a 1U, one-socket server that supports up to 2TB of RAM and ten NVMe drives. The second model is the Dell EMC PowerEdge R7415, which is a 2U, one-socket server that supports up to 2TB of RAM and 24 NVMe drives. The third model is the Dell EMC PowerEdge R7425, which is a 2U, two-socket server that supports up to 4TB of RAM and 24 NVMe drives.

Because of the underlying architecture of the AMD EPYC processor family, it is possible to have extremely high memory density (up to 2TB) and PCIe 3.0 lane capacity (up to 128) in a one-socket server. This makes it possible to run many storage-intensive workloads on a less expensive one socket server.

For the PowerEdge R6415, it is also possible to choose a relatively low core count processor, such as the eight-core AMD EPYC 7251 to minimize your SQL Server license costs. You can go all the way up to the 32-core AMD EPYC 7551P processor in that same model server.

Having additional choices for server models, and having a viable alternative to Intel-based servers is great news for SQL Server professionals.

#TBT

(Turn Back Time…) This section highlights some older resources we've referred to recently that you may find useful, plus blog posts we've published since the previous newsletter.

The theme for the TBT this time is on-disk structures:

- Kimberly's Pluralsight course: *SQL Server: Why Physical Database Design Matters*
- Kimberly's Pluralsight course: *SQL Server: Indexing for Performance* (which goes into lots of detail on index structures)
- Paul's blog posts on the basics:
  - *Inside the Storage Engine: Anatomy of a record*
  - *Inside the Storage Engine: Anatomy of a page*
  - *Inside the Storage Engine: Anatomy of an extent*
  - *Inside the Storage Engine: IAM pages, IAM chains, and allocation units*
  - *Inside The Storage Engine: GAM, SGAM, PFS and other allocation maps*
- Paul's blog post categories:
  - *Inside the Storage Engine*
  - *On-disk Structures*

Here are a few of the blog posts we've published since the last newsletter:

- Paul: *Setting mentoring ground rules*
- Paul: *More live online training from Kimberly in March!*
- Erin: *Query Store and "in memory"*
- Tim: *Three Years with SQLskills*
- Glenn: *SQLskills SQL101: How You Can Make Your Database Backups More Reliable*

I hope you find these useful and interesting!

Video Demo

This is the first in a 3-part series on statistics enhancements in SQL Server 2016/2017. Historically the only option for viewing a statistic's histogram was *DBCC SHOW_STATISTICS*. This changes in SQL Server 2016, where the ability to view the histogram using *sys.dm_db_stats_histogram* was introduced. In this week's video Erin walks through viewing all the stats data, and shows how to persist it for further analysis.

The video is just under 8 minutes long and you can get it:

- In WMV format here.
- In MOV format here.

And the demo code is here.

Enjoy!

Upcoming SQLskills Events

We have lots of events coming up in 2018 – from our online IEVLT course AND SQLintersection in March to our own LIVE, in-person Immersion Events in both the U.S. and London; all events are open for registration. Every event has a different focus and different benefits – from deep-technical training in our online courses and in-person IEs to wide-ranging topics at SQLintersection where you can learn more effectively how to keep moving forward in both your database and your career! And, of course, one benefit all our in-person events provide is networking!

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
- So why do you want to come to our training? And the winners are…
- Community blog posts about our classes
- Immersion Event FAQ

Online, March 2018

- **IEVLT**: Immersion Event on Very Large Tables: Optimizing Performance and Availability through Partitioning **(Final planned delivery in 2018)**
    - March 20-22

Orlando, FL, March 25-28, 2018

- SQLintersection co-located with the DEVintersection conferences (register with the 'sqlskills' discount to get save $50 on registration). See here for details.

Chicago, IL, April/May 2018

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
    - April 23-27
- **IE0**: Immersion Event for the Junior/Accidental DBA
    - April 23-25

- **IEUpgrade**: Immersion Event on Upgrading SQL Server
  - April 23-25
- **IECAG**: Immersion Event on Clustering and Availability Groups
  - April 26-27
- **IEAzure**: Immersion Event on Azure SQL Database and Azure VMs
  - April 26-27
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - April 30-May 4
- **IEBIStrat**: Immersion Event on Developing a BI and Analytics Strategy **(** NEW **)**
  - April 30-May 2
- **IEBISec**: Immersion Event on Securing Your BI Platform **(** NEW **)**
  - May 3-4
- **IEPS**: Immersion Event on PowerShell for SQL Server DBAs
  - April 30-May 2
- **IESSIS1**: Immersion Event on Learning SQL Server Integration Services
  - May 7-11
- **IEPML**: Immersion Event on Practical Machine Learning
  - May 7-11 **(** NEW **)**

Bellevue, WA, June 2018

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - June 18-22

London, UK, September 2018

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - September 10-14
- **IEAzure**: Immersion Event on Azure SQL Database and Azure VMs
  - September 10-11
- **IECAG**: Immersion Event on Clustering and Availability Groups
  - September 12-13
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - September 17-21

Click here for the main Immersion Event Calendar page that allows you to drill through to each class for more details and registration links.

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