(October 9th, 2017)

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 had a great week of classes in Chicago last week, where we launched our new IECAG class on Clustering and Availability Groups and delivered IEPTO + IESSIS1 and IE0. Now, this week, the whole team is here where we're launching our new IEAzure and IEUpgrade classes as well as delivering IEPTO2 and IEPS. It's a packed two weeks here in Chicago with a total of 8 classes running where 3 are brand new!

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

SQLskills News

Erin's latest Pluralsight course has been published! It's called ***SQL Server: Analyzing Query Performance for Developers***. Check out the details here.

Kimberly's latest Pluralsight course has been published! It's called ***SQL Server: Indexing for Performance*** and is a 7+ hour behemoth packed with her indexing wisdom. Check out the details here.

The first batch of 2018 classes is open for registration! We have all our usual classes and we'll be adding some more new classes over the coming weeks around security, data mining and AI, data due diligence/BI. See here for the 2018 Immersion Event class schedule. **Note: register in 2017 for discount prices!**

If not in Chicago, our team is presenting a number of workshops at Fall conferences, we hope you'll join us:

- Paul is presenting a workshop at the Fall SQLintersection conference in Las Vegas on Sunday, October 29th, titled *Performance Troubleshooting using Waits and Latches*. Check out the details here.
- Erin is presenting a workshop at the Fall PASS Summit in Seattle on Monday, October 30th, titled *Solving Common Performance Problems Using Query Store*. Check out the details here.
- Jonathan is presenting a workshop at the Fall SQLintersection conference in Las Vegas on Monday, October 30th, titled *Extended Events: WTF or FTW!* Check out the details here.

- Kimberly is presenting a workshop at the Fall SQLintersection conference in Las Vegas on Friday, November 3rd, titled *Very Large Tables: Optimizing Performance and Availability through Partitioning*. Check out the details here.
- Tim is presenting a workshop at the Fall SQLintersection conference in Las Vegas on Friday, November 3rd, titled *Common SQL Server Mistakes and How to Correct Them*. Check out the details here.

And all of us are presenting sessions at SQLintersection or the PASS Summit (Erin and Glenn). You can read more about our SQLintersection lineup in my blog post here.

Don't forget to check out our SQL101 posts… hopefully they'll help refresh or reinforce topics for even the more seasoned DBAs in the community. The blog posts are automatically collected here.

Finally, even if you can't join us in person, we've renewed our call for remote user group sessions for the second half of this year. We have almost 100 scheduled and completed so far; if you'd like one of us to present for your user group, check out my blog post here.

Book Review

One of the recent books I've read is Ed Yong's *I Contain Multitudes: The Microbes Within Us and a Grander View of Life*. What an excellent book! Thanks to Buck Woody for the recommendation. The book explains all kinds of fascinating things about bacteria and our microbiomes, the history of their understanding by science, and how the quest to rid ourselves of them is futile, and in quite a few cases, actually harming people that live in developed nations.

More interestingly, it also introduces many scientists who have identified beneficial bacteria and used them to help solve medical or environmental problems – e.g. preventing the dreaded Bd fungus killing a group of frogs, inoculating a group of mosquitoes with a bacteria that prevents them carrying dengue fever, and fecal-matter transplants to help sufferers of persistent C-diff infections - poo pill anyone? Hugely interesting, and written in an entertaining manner, this book is for anyone. Highly recommended!

The Curious Case of…

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

Another corruption story this time… Erin and I were working on a client system recently where they had a database which would fail a backup with page checksum errors, but passed a *DBCC CHECKDB*.

The only time this can happen is when an extent is a mixed extent (where the 8 pages in the extent can be allocated to potentially 8 different allocations units – see here) *and* some pages are erroneously marked as allocated by the relevant PFS page.

When that happens, *DBCC CHECKDB* will not attempt to read those pages, as it drives what pages to read from an allocation unit's IAM pages (the first of which lists the pages allocated from a mixed extent). This case is a gap in *DBCC CHECKDB*'s corruption-detection logic (there are more than you would imagine, given the infinite number of combinations of schema elements and possible corruptions…)

Anyway, as *DBCC CHECKDB* couldn't detect the corruption, it wasn't possible for it to make the repairs needed to fix them. So using *DBCC WRITEPAGE*, I worked out the changes needed in the allocation status for the erroneously-allocated pages, directly in the PFS page, and it worked!

This was an extremely rare case – it's much more common that a *DBCC CHECKDB* fails but a backup would succeed. In fact, this case came up on a forum only last week. The problem there is that a page gets corrupted in memory somehow, and is then written to disk with a valid page checksum. This means a backup will succeed, but *DBCC CHECKDB* will find the corruption.

Bottom line: You should always use page checksums for all your databases, perform regular consistency checks with *DBCC CHECKDB*, and use the *WITH CHECKSUM* option for your backups. Doing all of these means you will catch page corruptions, no matter how they occur.

Paul's Ponderings

I was teaching all last week so in this newsletter I'm re-running a Paul's Ponderings from a several years ago around how to ask questions, as it still applies today. Enjoy!

Last week I was complaining on Twitter about some of the ridiculous questions I'd received one day over email (I get a bunch of SQL Server questions every day from random people across the world), where it was clear that the senders had done no research at all before sending me an email. It was pretty frustrating and it bugs me when people are lazy.

I ended up tweeting:

> *Here's the rule: if my 14-yr old daughter could find the answer on Google for your question, you shouldn't ask or send in email.*

It brings to mind a training course from my days at Microsoft entitled Precision Questioning. The course taught how to hone your questioning technique for maximum efficacy, saving time for the questioner as well as the person/group being questioned. It also stressed the need to do due diligence in trying to find the answer to your question before taking up someone else's time to

help. I think something similar would be of great benefit to technical communities on the Internet where people ask questions.

Here are some quick guidelines for asking SQL Server questions I came up with while writing this:

- State your problem as unambiguously as possible.
- State what research you've done before asking the question (this show's you've done due diligence). You could even say what you searched for on Google/Bing.
- State the SQL Server version/SP/CU/build you're using.
- State what you've done to try to alleviate your problem so far, if anything.
- Ideally, include all Transact-SQL code to reproduce the phenomena (or, at least the details of the schema, indexes, query, query plan etc.) – just use your best judgment and include these as an annotated/commented script.
- Include whatever results are pertinent.
- Try to anticipate what people would ask to get more information, and provide answers in your question.
- Say please and thank you.
- Don't demand urgent or immediate help – you're not entitled to anything.
- *Edit 2017: Don't send it individually to multiple people (especially at the same company) as they may all end up doing the same work; it's best to include them on the same email so that they know if something's already been answered and/or want to add anything new.*

Edit 2017: the last three in the list above are the ones that frustrate me the most. I've had people send email asking for help, and then reply again several hours later demanding a reply. One guy even sent me an all-caps follow-up the same day saying PLEASE REPLY!

Now, if you're using #sqlhelp on Twitter, most of this isn't possible. But asking a question over two tweets, using 1/2 and 2/2 is perfectly acceptable, and putting 'Looked on Google already' shows due diligence. You might be asked to post a longer question on a forum, in which case you should do so.

Whatever your chosen method for asking, make sure you respond to those who respond to you. Nothing puts people off from answering questions more than a complete lack of any thanks, or response to a clarifying question. And it's just common courtesy (sadly lacking these days, especially in the relative anonymity of the internet).

Edit 2017: And if you're asking a question on #sqlhelp, and someone replies with a clarifying question, at least have the courtesy to reply, even if to say 'Solved the problem – thanks'

And don't forget, this isn't just for asking questions online, these guidelines should help within your company as well.

Now, where can you ask your SQL Server question online?

- Twitter, using the #sqlhelp hash tag
- Forums like on SQL Server Central, and their dedicated Ask SQL Server Central site
- MSDN's SQL Server forums
- SQL Sentry's performance Q&A forum
- Stack Overflow for general programming questions
- Server Fault for general admin questions
- Direct email to people (usually as a last resort)

And a whole bunch of other places, but these are the main ones as far as I'm concerned.

Edit 2017: contrary to popular myth, although I do reply to everyone that send me a question in email, I often direct people to post their question in a forum, as their question is too in-depth for an email reply, or suggest that they really need some dedicated consulting assistance. I love helping people, but there's only so much time in the day…

The more information you can provide, show you've tried on your own to find an answer, and are polite and grateful, the more likely you'll get useful answers quickly. And if you're sending direct email, you'd better do all of this if you want a reply.

To finish off, I leave you with a blog post I wrote three years ago today: RTFM. No seriously, R.T.F.M. Then ask your question.

Call to action: The next time you're stumped for an answer and you'd like some free help, post your question using one of the mechanisms I've described above. But make sure you've done due diligence in trying to find the answer first, phrase the questions coherently, and if you have space, explain how you've searched and give as many pertinent details as possible. You'll be amazed at the great help you can get online!

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.

NVMe RAID Support for AMD Ryzen Threadripper Platform

AMD has released a set of free NVMe RAID drivers for Windows 10 running on AMD Ryzen Threadripper platform machines. This will let you use up to six NVMe SSDs in a RAID array in

a system with a single GPU. In order to take advantage of this you will also need a BIOS update from your motherboard vendor.

One you have done this, you can create a bootable NVMe RAID array for your OS using the motherboard BIOS, and/or you can create and manage a non-bootable NVMe RAID array using AMDs RAIDXpert2 Management software for Windows 10. This gives you the flexibility to have either even more performance or more redundancy than you can get with a single NVMe SSD.

Intel has plans for a similar feature called Virtual RAID on CPU (VROC) with their X299 HEDT platform, but they are planning on charging extra for the feature and requiring a hardware dongle. PCPer has an interesting test of a pre-release version of X299 VROC setup here.

In case you are wondering why this could be relevant to you, this will let you have previously unheard-of levels of storage performance for a desktop workstation. For example, AMD claims they saw 21.2GB/s reading from six 512GB Samsung 960 PRO NVMe drives in a RAID-0 array (you also have the choice of using RAID-1 or RAID-10.

Using Adaptive Query Processing in SQL Server 2017

Microsoft's Joe Sack has a great blog post about how the new Adaptive Query Processing feature in SQL Server 2017 can give you better and more predictable query performance in some scenarios, including batch mode adaptive joins, batch mode memory grant feedback, and interleaved execution for multi-statement table valued functions. This feature is also available in Azure SQL Database.

You will get this new behavior for a particular database by enabling compatibility level 140. Books Online goes into much more detail with the entry Adaptive Query Processing in SQL databases.

SpaceX Mars Colonization Architecture

On September 29, Elon Musk presented SpaceX's new Mars Colonization plan at the 68th International Astronautical Congress (IAC) in Adelaide, Australia. The gist of this new plan is to develop a new "BFR" launch system (I'll leave it to you to figure out what BFR stands for) with 31 Raptor engines that will soon replace the existing Falcon 9 and upcoming Falcon Heavy rockets, which will be used for SpaceX's existing satellite launch and ISS resupply business and also for their Mars Colonization plans.

The BFR system is designed to be 100% reusable, which will dramatically reduce launch costs. SpaceX plans to launch at least two uncrewed ships to Mars in 2022, with crewed ships leaving in 2024. You can watch the complete presentation here.

#TBT

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

We had some good discussions on using DMVs for finding performance data in our IEPTO1 class last week, so that's the theme for this TBT:

- Glenn's course: *SQL Server DMV Diagnostic Queries: Part 1*
- Glenn's course: *SQL Server DMV Diagnostic Queries: Part 2*
- Glenn's course: *SQL Server DMV Diagnostic Queries: Part 3*
- Glenn's monthly-updated DMV queries

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

- Paul: *SQLskills SQL101: Running out of ints and bigints*
- Glenn: *Modern SQL Server Servicing Model*
- Erin: *Changes to query text and the effects in Query Store*
- Erin: *Query Store Pre-Con at PASS Summit: More Details*

I hope you find these useful and interesting!

Video Demo

The demo video this time is taken from Erin's new Pluralsight course *SQL Server: Analyzing Query Performance for Developers*. In the demo, Erin shows how to examine execution statistics for query plan operators.

The video is around 6.5 minutes long and you can get it:

- In MP4/MOV format here

The demo code is here.

Enjoy!

Upcoming Immersion Events

We've released the first set of 2018 classes for registration. Over the coming weeks we'll be adding a few more classes to the Spring line-up in Chicago, including classes on data mining and AI, security, and data due diligence/BI. We'll also be adding one or more classes in Europe, in London and/or Dublin, in the second half of the 2018.

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

Chicago, IL, April/May 2018 **(all classes have discounts for 2017 registrations!)**

- **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
- **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

Bellevue, WA, June 2018 **(all classes have discounts for 2017 registrations!)**

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

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