(November 6th, 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 fantastic time at our SQLintersection show last week in Las Vegas, and I was a lot more successful at the craps table than last year! Now we're done teaching until 2018, so for the rest of the year we'll be heads-down producing more Pluralsight courses for you, working with our clients, and planning our 2018 events.

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

SQLskills News

Glenn's latest Pluralsight course has been published! It's called ***SQL Server: Understanding, Configuring, and Troubleshooting Log Shipping***. Check out the details here.

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, including PowerShell, Azure, Availability Groups, and more. See here for the 2018 Immersion Event class schedule. **Note: Class prices are going up a bit in 2018 so register in 2017 for 2017 prices/discounts!**

Also, Kimberly and Stacia have been discussing the confusion around "the future of data" and that a lot of companies just aren't sure about where (and how) they should move forward to best utilize the data they have. There's Relational Data Warehousing, there's Data Science, there's the Microsoft BI stack – all of this is about your businesses intelligence. And now we have a class to help you figure out one of the most important questions today; ***what's your business intelligence strategy***? **As a result, we've *just* added two new classes around Business Intelligence:**

- **IEBIStrat on Developing a BI and Analytics Strategy**
- **IEBISec on Securing Your BI Platform**

Click the links for more details.

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 more than 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 Judith Schalansky's *Atlas of Remote Islands: Fifty Islands I Have Never Set Foot On and Never Will*. I've always been fascinated by maps and cartography, and this is one of the more unusual atlases I've read. It describes (usually with an historical narrative linked to the island) and illustrates 50 of the most remote islands in the world. What I found particularly interesting was using Google Maps to enter the latitude and longitude of each island and then look at them closely using the available satellite imagery. Great little book and 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.

Last week when I was at SQLintersection I was talking to an attendee who wanted some advice on making part of their workload go faster. (Actually, I had many such fruitful conversations, but I picked this one for the newsletter.)

The portion of the workload in question was a nightly process that would delete all the data from the current month from a reporting table and then rebuild it. They had a valid reason for doing this, but that's not the interesting part of the story.

As the month progressed, and more and more data accumulated, the delete portion of the process was taking longer and longer. They wanted to know how to speed up the deletes.

I had two questions for them:

1. Do you delete *all* the current month data every night?
2. Are you on Enterprise Edition, or moving to 2016/2017 Standard or Enterprise?

Answer 1: Yes. Answer 2: 2016 Standard by the end of the year.

Perfect. My suggestion to them was to consider moving the reporting table to be partitioned, by month. This means the nightly delete can be simplified into one of two optional processes:

1. If they have isolated the table to a filegroup and they want to be able to back it up and archive it, then they can switch out the current month's partition and then backup the filegroup.
2. If you really only want to "switch-out for drop," then you can also take advantage of the new SQL Server 2016+ feature to truncate just a partition. See TRUNCATE TABLE … WITH (PARTITIONS).

They were ecstatic and told me that idea alone was worth the cost of the conference. That made me pretty happy too. ☺

Bottom line: Sometimes a performance boost isn't best accomplished from tuning the existing solution but from considering whether there's an entirely alternative strategy/feature that can provide the boost with a minimum of fuss and risk.

Paul's Ponderings

Last week I was helping someone try to recover data from a corrupt database, from an online forum question. They did not have any up-to-date backups without the corruption in, so fixing their backup strategy was a piece of advice they were given by a few people.

The output from DBCC CHECKDB on the database was:

Msg 8921, Level 16, State 1, Line 1

Check terminated. A failure was detected while collecting facts. Possibly tempdb out of space or a system table is inconsistent. Check previous errors.

Msg 824, Level 24, State 2, Line 1

SQL Server detected a logical consistency-based I/O error: torn page (expected signature: 0x0; actual signature: 0x5555300). It occurred during a read of page (1:58) in database ID 10 at offset 0x00000000074000 in file 'D:\dbname.mdf:MSSQL_DBCC10'. Additional messages in the SQL Server error log or system event log may provide more detail. This is a severe error condition that threatens database integrity and must be corrected immediately. Complete a full database consistency check (DBCC CHECKDB). This error can be caused by many factors; for more information, see SQL Server Books Online.

They'd tried running repair, but of course if *DBCC CHECKDB* says that it has to stop (i.e. error message 8921), then it can't run repair.

I explained this, and how page 1:58 is a system table page and unrepairable, and so they'd have to script out as much of the database schema as possible, create a new database, and extract as much data as possible from the broken database.

I also explained that the page is part of the *sys.syscolpars* table, which is the equivalent of the old *syscolumns*, so that approach might not work.

Unfortunately my suspicions were correct, and the script/extract approach did indeed fail.

On a whim, I suggested trying something radical. A few years ago I blogged about a way to 'fix' broken boot pages using a hex editor to overwrite a broken boot page with one from an older copy of the database (see here). I'd never tried it on a system table page before, but I figured that the page ID was low enough that the page likely hadn't changed for a while.

What do I mean by that? Well, the *sys.syscolpars* clustered index is ordered by object ID, so the first few pages in the clustered index (of which page 1:58 is one), have the columns from the system tables, with very low object IDs. There's never going to be the case where a new user table gets created and causes an insert into one of these low tables.

This means that an older backup of the database would have the current state of page 1:58 in it. So I suggested using the boot page hack on page 1:58 from the person's older backup.

And it worked!

Luckily there wasn't any other corruption in the database, so all the person had to do was root-cause analysis and remediation, and fixing the backup strategy so the situation wouldn't arise in future.

Call to action: The moral of my story is that in a disaster situation, when backups aren't available; don't be afraid to try something radical. As long as you try it on a copy of the database, it's not as if you can make the situation any worse. And if you're lucky, you'll be able to make the situation a lot better.

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.

Intel Optane SSD 900P has been released

Intel has released the desktop client focused Intel Optane SSD 900P, which is a new storage device that uses 3D XPoint Technology. This product is initially being released as a half-height, half-length add-in-card (HHHL AIC) or a 2.5" U.2 form factor, both of which using a PCI 3.0 x4 interface and the NVMe protocol.

The AIC will be available in 480GB or 280GB capacities, while the U.2 form factor will only be available in a 280GB capacity. The Intel Optane SSD 900P will be priced at $389.00 (280GB)

and $599.00 (480GB). The Intel Optane SSD 900P is roughly two to three times the cost-per-GB of NAND SSDs.

You might be wondering why this is a significant product that warrants the extra cost? Here are some of the advantages over current NAND-based flash storage devices:

- High random read and write performance
- High performance at low queue depths
- High simultaneous read and write performance
- High read and write performance at small capacity points
- High performance maintained as the drive fills with data
- Higher endurance than current NAND technology

Traditional NAND-based SSDs excel at very high queue depths that are not usually encountered outside of synthetic benchmarks (especially for random read performance). Optane-based SSDs perform extremely well for random reads at low queue depths. This gives you outstanding responsiveness and performance where it is going to be most noticeable in daily usage. Here are several reviews of this new product:

- PC Perspective
- Tom's Hardware
- Anandtech

Windows Server, version 1709

Microsoft has announced some changes to the release schedule and servicing model for Windows Server. The new Semi-Annual Channel is a twice-per-year feature update release with an 18-month servicing timeline (meaning that Mainstream support ends 18 months after that Semi-Annual Channel release becomes available).

The current release in this channel is Windows Server, version 1709, which became available on October 17, 2017. This release will fall out of Mainstream support on April 18, 2019, and there is no Extended support period. In this new model, Windows Server releases are identified by the year and month of release: for example, in 2017, a release in the 9th month (September) would be identified as **version 1709**.

If you want to use the Semi-Annual Channel, you will have to be comfortable running SQL Server on Server Core (with no integrated GUI). You can either start using Powershell or you can use tools like Project Honolulu, which is a locally deployed, browser-based, management tool set that enables on-premises administration of Windows Servers with no Azure or cloud dependency. Project Honolulu requires Windows 10 or Windows Server 2016.

The two most interesting new features for SQL Server in Windows Server, version 1709 are Storage-Class memory support for Hyper-V VMs and Virtualized Persistent Memory (vPMEM) for Hyper-V VMs.

Storage-class memory support for VMs enables NTFS-formatted direct access volumes to be created on non-volatile DIMMs and exposed to Hyper-V VMs. This enables Hyper-V VMs to leverage the low-latency performance benefits of storage-class memory devices. Virtualized Persistent Memory (vPMEM) is enabled by creating a VHD file (.vhdpmem) on a direct access volume on a host, adding a vPMEM Controller to a VM, and adding the created device (.vhdpmem) to a VM. Using vhdpmem files on direct access volumes on a host to back vPMEM enables allocation flexibility and leverages a familiar management model for adding disks to VMs.

Virtualized Persistent Memory (vPMEM) is enabled by creating a VHD file (.vhdpmem) on a direct access volume on a host, adding a vPMEM Controller to a VM, and adding the created device (.vhdpmem) to a VM. Using vhdpmem files on direct access volumes on a host to back vPMEM enables allocation flexibility and leverages a familiar management model for adding disks to VMs.

Using Storage-class memory in a VM will let you use the Persisted Log Buffer feature (aka "tail of the log caching") that was introduced in SQL Server 2016 SP1, as described 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.

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 new features rather than scenarios: *Partitioned Table and Index Strategies Using SQL Server 2008*
- Kimberly's recorded partitioning training for the MCM certification

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

- Paul: *New class: IEBISec on Securing Your BI Platform*
- Paul: *New class: IEBIStrat on Developing a BI and Analytics Strategy*
- Erin: *PASS Summit 2017: Day 1* (Day 1 keynote notes and summary)

- Erin: *PASS Summit 2017: Day 2* (Day 2 keynote notes and summary)
- Glenn: *SQL Server 2017 CU1 is Available*
- Glenn: *SQL Server Diagnostic Information Queries for November 2017*

I hope you find these useful and interesting!

Video Demo

In a departure from our normal demo videos, this time I'd like to present you with two videos from my Pluralsight course *Communications: How to Talk, Write, Present, and Get Ahead!* In these two videos, I explain how to properly express yourself when communicating with other people.

The videos are 8 minutes long in total and they're in WMV format: part 1 and part 2.

No demo code this time! ☺

Enjoy!

Upcoming Immersion Events

We've released the first set of 2018 classes for registration. And, we've just added some new ones on security and data due diligence/BI. Over the coming weeks we're still planning to add a few more classes to the Spring line-up in Chicago, including classes on data mining and AI. We'll also be adding a few classes in London in September 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
- **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

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