(May 29th, 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,

If you're in the US, I hope you enjoyed the long weekend, which is why this newsletter is coming out on Tuesday instead of Monday.

We're putting together a bunch of new live, online courses which we'll be announcing in the next newsletter – stay tuned!

In the meantime, our Bellevue IEPTO1 class is fast approaching, and there are still a few seats remaining in our London classes in September…

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

SQLskills News

In-person US classes: Our single remaining in-person class for 2018 is IEPTO1 in Bellevue, WA in June. **We don't have any additional classes planned in the US in 2018;** see here for our 2018 Immersion Event class schedule.

In-person London classes: We're bringing four of our Immersion Events to London in September: IEPTO1 and IEPTO2, plus our new classes: **IEAzure** (on Azure, Azure VMs, and azure Managed Instance) and **IECAG** (on clustering and availability groups). See here for details.

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 70 for 2018 already! If you'd like one of us to present for your user group, check out my blog post here. **Note: Tim has a new user group session on Azure Managed Instance** that he's happy to present to your group – see here for details.

Book Review

The latest book I've read is Donald Matthew's *King Stephen*. I have a biography of each English/British monarch from Edward the Confessor through Edward VII, and I'm (very) slowly working through them – this is #5. Stephen reigned from 1135 to 1154, and is popularly blamed for the 'anarchy' of his reign between Henry I and Henry II. This book tries to redress that a little, while also presenting the goings on of the reign. It's a pretty dry book, as Stephen wasn't that

exciting a king, but interesting for serious history fans. A much more readable (and fictionalized) account of his reign is the excellent *When Christ and His Saints Slept* by Sharon Kay Penman, which I highly recommend.

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.

Last week I was helping someone recover data from a corrupt database and they didn't have any current backups (a familiar story in this section of the newsletter now…)

The corruption was confined to a single clustered index and the repair was going to delete a couple of data pages at different points in the leaf-level of the index, resulting in some data loss (and those pages were unreadable).

I suggested to them that they check what nonclustered indexes they had, and as luck would have it their nonclustered indexes contained all the table columns. I then told them about a technique I'd demonstrated several times publicly that would allow them to recover the deleted data with zero data loss by extracting all the columns from the nonclustered indexes using a backup of the pre-repair, corrupt database.

They proceeded to successfully follow the technique and recovered without any data loss! I'll explain how to do that in the *Ponderings* section below.

Bottom line: Any time you're looking at data loss while recovering from corruption, stop to think whether there's any way of getting the deleted/lost data back. If you wait until after you've done the repair, it might be too late. More details below…

Paul's Ponderings

Whenever I'm involved in a case of corruption where there's going to be data loss I always check to see if there's some way of getting the data back (assuming that no backups are available). This means thinking a little outside the box to identify where copies of some of the data might be found.

Let's say you're going to lose a bunch of data from some clustered indexes in an OLTP environment. Here are some examples of where this data may be duplicated, and so salvageable:

- In nonclustered indexes (which I'll explain after the list)
- In a data warehouse (albeit maybe up to a day or two behind the OLTP system, depending on how often the warehouse is updated/refreshed)
- In a test/dev system (with the same caveats as above)

- In a replication subscription database (if the affected tables are part of a publication, taking into account filtering and column selection)
- In a Change Data Capture table (if the affected tables have CDC enabled, and you're using Enterprise Edition prior to 2016 SP1, or Standard/Enterprise in 2016 SP1+)
- In hard-copy printouts (I've seen databases recreated by hand-entering printed data over many days)
- In any other places the data could be (one case I was involved in had an insurance company extracting data from insurance policy scanner memory…)
- And then there's the case of the guy who wrote some code to piece together a broken database programmatically by figuring out where to get duplicate blocks from the broken RAID array (see here) because their most recent backup was SIX YEARS OLD!

Nonclustered indexes are the easiest way to salvage data. By their very nature, they are copies of the table data, so the more columns in the table that are in the nonclustered indexes, the more complete your data recovery can be, as long as the nonclustered indexes aren't damaged.

The technique is relatively simple:

- Back up the damaged database and restore it somewhere else
- In the restored copy:
    - Figure out the missing range of records (i.e. which cluster keys are missing)
    - Figure out what columns you have in all the nonclustered indexes (remembering INCLUDEd columns)
    - Build a dummy table with the same set of columns as the damaged table
    - Select the first set of columns from the first nonclustered index for just the missing cluster keys, inserting NULLs for the other columns,
    - Select the second set of columns from the second nonclustered index for just the missing cluster keys, joining with the dummy table
    - And so on, until you've recovered as many columns as you can
- In the real copy:
    - Run repair
    - Select from the dummy table in the restored copy into the real table with the deleted rows, using *IDENTITY_INSERT* if necessary
- Success!

In the demo video below, I'll show you exactly what I mean. If you have more than one missing range of records, you need to follow the process for each missing range.

Call to action: Same as for the *Curious Case*; any time you're looking at data loss while recovering from corruption, stop to think whether there's any way of getting the deleted/lost data back. If you wait until after you've done the repair, it might be too late.

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.

Microsoft Guidance for Speculative Store Bypass (SSB)

On May 21, 2018, Microsoft and Google announced that a new subclass of speculative execution side channel vulnerabilities known as *Speculative Store Bypass (SSB)* has been publicly announced and assigned CVE-2018-3639. This vulnerability affects both Intel and AMD processors.

As of the time of publication, Microsoft is not aware of any exploitable code patterns from this vulnerability in their software or cloud service infrastructure. Despite this, they are completing final testing and validation of AMD-specific updates for Windows client and server operating systems. Linux distributors are also developing operating system updates for SSB.

AMD has updated their AMD Processor Security Updates page with more information. They also have a five page whitepaper with code samples on how to enable Speculative Store Bypass Disable to protect against this vulnerability.

Intel Optane 905P Released

Intel has released the new Intel Optane 905P family of SSDs, which is a slightly improved version of their recent Intel Optane 900P family. This new line is available as a 960GB PCIe add-in card (AIC) for $1299.00 and as a 480GB PCIe 2.5" U.2 15mm drive for $599.00.

These drives use Intel's 3D XPoint memory technology that doesn't require a read, modify, and write cycle like NAND flash, and it performs roughly the same no matter how full the drive is (unlike NAND flash drives). They also have extremely low latency and excellent random read I/O performance at low queue depths. PC Perspective has a good review here.

AMD Launches "This is EPYC" Campaign

AMD has launched a marketing campaign around the AMD EPYC server processors that were launched in June 2017. AMD is making a big push to get the message out that their Zen-based EPYC processors have a number of superior technical features compared Intel Xeon processors. These include 2TB of RAM per socket (with an eight-channel memory controller) and 128 PCIe 3.0 lanes in a one-socket server.

AMD has been quite successful making inroads into the cloud space, with Microsoft using AMD EPYC processors in their Azure Lv2 series of storage-optimized VMs. They are also doing well with large server vendors, with, HPE announcing AMD EPYC-based ProLiant DL385 Gen 10

servers in November 2017 and Dell EMC announcing AMD EPYC-based servers (which include the PowerEdge R6415, PowerEdge R7415, and PowerEdge R7425) in February 2018.

Dell Releases PowerEdge R940 Four-Socket Server

Dell has released the new four-socket PowerEdge R940 3U rack-mount server. These servers have 24 2.5" front drive bays (12 of which can be NVMe) and 48 DDR4 DIMM slots (for up to 6TB of RAM). They support the current Intel Xeon Scalable Processor family (Skylake-SP) and they will also support the upcoming Intel Xeon Cascade Lake-SP processor family that is scheduled for release in H2 of 2018.

If you absolutely need more total CPU cores, PCIe lanes, and memory slots than you can get in a two-socket Intel Xeon Scalable Processor-based server, you can move up to a four-socket machine. There is some increased NUMA overhead with a four-socket server, so you **will not** get twice the capacity of a two-socket server with identical processors.

#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.

Two weeks ago I taught an online session with our partners SentryOne on my favorite topic of all the material I teach – wait statistics – so that's the theme for #TBT this time. Here are some wait statistics resources for you:

- My Pluralsight: *SQL Server: Performance Troubleshooting Using Wait Statistics*, that has 4.5 hours covering waits, latches, and spinlocks
- My main post about wait statistics
- My wait stats library with info on more than 360 wait types
- The introductory whitepaper Jonathan and Erin wrote
- Microsoft whitepapers on latches and spinlocks I helped with
- My three blog post series on waits, latches, and spinlocks
    - Our articles on SentryOne's sqlperformance.com about wait statistics

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

- Paul: *SQLskills SQL101: Why do some wait types need to be ignored?*
- Glenn: *SQLskills SQL101: The Importance of Maintaining SQL Server*
- Tim: *Why You Need Baselines*

I hope you find these useful and interesting!

Video Demo

Continuing the theme of out-of-the-box thinking around corruption recovery, the video this time is an example of what happened with the *Curious Case* above. It's a demo from my Pluralsight course *SQL Server: Advanced Corruption Recovery Techniques* and shows how you can recover data deleted by a repair operation using nonclustered indexes in a pre-repair copy of the database.

The video is just over 7 minutes long and you can get it in WMV format here.

And the demo code is here.

Enjoy!

Upcoming SQLskills Events

We have lots of events coming up in 2018 – from our ***live, online*** courses (more coming in the next newsletter!) 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

Bellevue, WA, June 2018

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - June 18-22 **(** Buy 2, get 1 free!, *only 8* seats remaining **)**

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