(May 8th, 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,

Today we're starting the last week of our Spring Immersion Events in Chicago, with IEPDS, IEPS, and IEHADR running simultaneously.

We're expanding the newsletter again with this issue, adding a new section called Glenn's Tech Insights, where Glenn shares news and views from the hardware and Windows world. The pace of change is so fast these days!

Speaking of Glenn, we're excited to share that Glenn reached his five-year anniversary of being part of the SQLskills team on May 1st – congratulations Glenn! We're looking forward to many more!

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

SQLskills News

We've just released our October line-up of classes in Chicago, including IE0, IEPTO1, IEPTO2, IEPDS, IESSIS1, plus two new courses. We're debuting a **new two-day class on Azure SQL Database and Azure VMs** and a **new three-day class on Upgrading SQL Server**. See here for the complete 2017 SQLskills Immersion Event class schedule.

Our Spring SQLintersection conference is also fast approaching at the end of May, and we have a phenomenal line up of workshops, sessions, and speakers (including Joe Sack, Buck Woody, Brent Ozar, Bob Ward, and many from SQLskills). Check out this blog post for all the details, and use the discount code 'sqlskills' when you register to save $50.

We've started a new initiative where we're all blogging about introductory topics, to help the burgeoning numbers of non-DBAs or junior DBAs who are responsible for SQL Server instances. We're calling it SQL101, and hopefully it'll even help refresh or reinforce topics for the more seasoned DBAs in the community. The blog posts will be automatically collected here.

Even though we're not teaching any Immersion Events in Europe this year, Kimberly and I will both be **presenting at SQLSaturday #620 in Dublin in June**.

Finally, even if you can't join us in person, we're still taking requests for remote sessions for this year. We have over 70 scheduled so far; if you'd like one of us to present for your user group, check out my blog post here.

Book Review

The latest book I've read is a Paul Theroux's *Fresh Air Fiend: Travel Writings*. I love all of Theroux's travel books and this one was no different. Instead of being about a single journey, this is a collection of long essays and short articles about travel writing. It covers how he approaches a subject, inspirations for some of his books, reviews and eulogies of travel writing friends, and some plain old travel writing itself. Highly recommended for all Theroux fans.

Also check out his classics *The Great Railway Bazaar* and *Riding the Iron Rooster: By Train Through China*.

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.

Jon was working with a client recently who contacted us to troubleshoot terrible performance and 100% CPU on a production virtual machine. He looked at SQL Server and could quickly see that there didn't seem to be anything out of the ordinary.

He then stepped back and looked at the VM configuration. The VM was configured to have 8 CPUs with a total of 7GHz of clock speed, and with such a low amount for each CPU, the workload was causing the VM to be throttled. The physical host had 8 CPUs running at 3GHz each.

The problem was that the VM had originally been a test instance but had recently moved to production, and no-one had remembered to reconfigure it for the production workload.

Bottom line: Windows sees what the hypervisor tells it, so in this case, the 100% CPU was real as far as Windows was concerned. When troubleshooting performance in a VM, it's always worth checking the VM configuration to make sure someone hasn't misconfigured it.

Paul's Ponderings

In a newsletter back in November 2014, I discussed the delayed durability feature that debuted in SQL Server 2014 and how it can help the logging throughput of certain workloads. (As a quick refresher, it allows a transaction to commit without having to do a log block flush – hence delaying the durability of the transaction – and can be configured for everything in a database, or just for transactions that request it. It's available in all editions.)

There are some big drawbacks with delayed durability though:

- It only helps workloads that are bottlenecked on log block flushing (usually because a transaction cannot release any locks it is holding until the log flush has completed).
- It has some undesirable side effects:
    - The transaction may not be there after a crash, if the crash occurs before the actual log block flush happens
    - The transaction may also not be on a synchronous AG replica after a failover, if the failover occurs before the actual log block flush happens, as the copy of the log block to the AG replica is triggered by the local log block flush

Then there's In-Memory OLTP (a.k.a. 'Hekaton') that was also introduced in SQL Server 2014, and is much more usable in SQL Server 2016. This can remove all logging overhead completely; leading to incredible performance boosts (e.g. 1 million one-row-at-a-time inserts in 7 seconds compared with 6 minutes with normal logging). It's available in all editions from SQL Server 2016 SP1.

Unfortunately there are restrictions on where it can be used and it can involve a lot of work to change a workload to use it. We've seen our clients using it successfully in the following situations:

- Speeding up an ETL process by making a staging database completely in-memory and non-logged
- Removing the bottleneck of ASP.NET session-state tracking
- Changing temp tables into in-memory tables

Now there's a new kid on the block that works for *all* scenarios to remove the logging overhead with zero application changes and no downsides: the ability to store the tail of the log on an NV-DIMM (an NV-DIMM is non-volatile memory, which means its contents survive a crash).

Using Windows Server 2016 (or Windows 10), SQL Server 2016 SP1, and a server that supports NV-DIMMs, you can configure a database to use the NV-DIMM as the cache of log blocks that have just been written. This means that when any log block write occurs, there is no waiting at all – it's as if the flush completes instantaneously – completely removing the overhead of having to wait for log block flushes.

When the NV-DIMM fills up, it gets written to the actual log on disk, and if a crash occurs while there's unwritten log in the NV-DIMM, the log tail is reconstructed from the NV-DIMM before crash recovery takes place.

You can read about how to configure this and more on how it works in this blog post.

Call to action: if you find that your log is a bottleneck, and it's already on the fastest portion of the I/O subsystem, and there's nothing you can do to change the pattern of your transactions, one of these three options may work for you. I'm really excited about the NV-DIMM option, as that requires absolutely no changes and has no downsides at all.

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.

Enhancements coming in SQL Server 2017

Microsoft announced some very useful new manageability enhancements for SQL Server 2017, as detailed in this blog post: Community driven Enhancements in SQL Server 2017

Intel's Next-Generation of Processors

According to some leaked Intel documentation, Intel is changing the branding for their next generation of Xeon processors (Skylake-SP). The Xeon E5 line will become Xeon Gold, while the Xeon E7 line will become Xeon Platinum. The Gold line will include the Intel Xeon 5000 and 6000 series, while the Platinum line will include the Intel Xeon 8000 series.

From a Product Change Notification (PCN) temporarily published by Intel, we have the model numbers for 34 of these upcoming Xeon processors. We also have the base clock speed for each model, but no other specifications.

Unfortunately, we don't have the physical core counts for these upcoming processors (which is critical information for SQL Server usage). Current speculation is that these new processors (which will require new model servers) will be generally available in late Q2 or early Q3 of 2017.

They are not electrically or physically compatible with Broadwell-EP/EX, so they will require new model servers from the server vendors.

Some more information about them is here.

Intel DC P4600 and P4500 Series NVMe Flash Storage

These have significantly better performance specifications than the previous DC P3700 Series (even though they use TLC NAND instead of MLC NAND), and they will go up to 4TB in size. They also use NVMe 1.2. More info is here: 4TB DC P4600.

They also will have a lower cost line, with less endurance and lower write performance, see here: 4TB DC P4500.

The only downside is the lower endurance, which is 4 drive writes per day (DWPD) for the P4600 series, but they do have a five-year warranty.

Storage Spaces Direct

I recently read a good blog post from Microsoft's Claus Joergensen about the performance advantages of remote direct memory access (RDMA) networking compared to TCP/IP on the same hardware. The only difference in the testing is whether or not RDMA is enabled or not. The key takeaway is that RDMA reduced latency by 28-36% and reduced CPU consumption on the host by 27% compared to TCP/IP.

Check it out here: To RDMA, or not to RDMA – that is the question.

#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.

Last week in IEPTO2 I taught my favorite module out 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 300 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: Query plans based on what's in memory*
- Paul: *How are default column values stored?*
- Glenn: *SQL Server Diagnostic Information Queries for May 2017*
- Glenn: *SQLskills SQL101: Sequential Throughput*
- Erin: *Endpoints for Mirroring and AGs in SQL Server 2016*

I hope you find these useful and interesting!

Video Demo

Many times when Tim is presenting his courses on Azure, people ask if they still need to perform maintenance tasks on their databases such as index reorgs/rebuilds, updating statistics, and checking for corruption. In this video, Tim answers that question and shows you one of the simple ways of scheduling Ola Hallengren's *Index Optimize* and *Database Integrity Check* procedures using a linked server from an on-premises SQL Server.

Stay tuned for more upcoming videos showing you additional ways of scheduling your jobs against an Azure SQL Database.

The video is about 4.5 minutes long and you can get it:

- In WMV format here
- In MOV format here

And the demo code is here.

Enjoy!

Upcoming Immersion Events

All 2017 classes are available for registration!

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
- Community blog posts about our classes
- Immersion Event FAQ

Bellevue, WA, July/August 2017

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - July 31-August 4
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - August 7-11

Chicago, IL, October 2017

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - October 2-6

- **IE0**: Immersion Event on Performance Tuning and Optimization – Part 2
  - October 2-4
- **IESSIS1**: Immersion Event on Learning SQL Server Integration Services
  - October 2-6
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - October 9-13
- **IEPDS**: Immersion Event on Practical Data Science
  - October 9-13
- **IEAzure**: Immersion Event on Azure SQL Database and Azure VMs
  - October 9-10 **** New class! ****
- **IEUpgrade**: Immersion Event on Upgrading SQL Server
  - October 11-13 **** New class! ****

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