(January 15th, 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,

This newsletter is coming to you from Redmond, where we've just finished our first ever online Immersion Event – Kimberly presented her widely-requested course *IEVLT: Immersion Event on Very Large Tables: Optimizing Performance and Availability through Partitioning*. It was a HUGE success and we're planning more online courses through the year from early April onwards. Watch the next newsletter for the announcement of the next one, another run of Kimberly's IEVLT, with an exclusive discount for subscribers.

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: Automatic Tuning in SQL Server 2017 and Azure SQL Database***. Check out the details here.

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 40 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 a Haynes Manual by David Baker: *Nuclear Weapons: 1945 Onwards (Strategic and Tactical Delivery Systems)*. Haynes is a British publishing house that started out in the 1960s creating in-depth manuals for do-it-yourself car maintenance and repairs for cars available in the U.K. (I used to use them in the 1990s and my Dad has always used them extensively for work on his cars) and has since branched out around the world. They also produce 'manuals' for things like spacecraft, Star Wars/Star Trek ships, and other interesting technical items, such as this one.

I have a bunch of their non-car manuals and this is the first one I've read, as I'm on a 'nuclear weapons/strategy' kick at the moment. It's a very interesting book, giving an account of the development and history of each nuclear power's atomic weapons. It's not super in-depth, so it's very readable if you're just casually interested, and it has a ton of photos, which really made it worthwhile for me to read. My Dad was a nuclear reactor engineer in the Royal Navy (serving many years on Polaris missile submarines, eventually as 'Chief of the Boat') and helped develop the Royal Navy nuclear reactors in the 1960s, so it was particularly interesting to read the section about the British nuclear program. 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.

I was working with a new client last week, and as part of a health check I wanted to look at the instance's SQL Server error log. It was truly enormous, more than 20GB! Once I managed to open it, I saw that it was full of backup success messages and stretched back to November 2017.

We see this quite often, and there are two very simple ways to prevent unwieldy error logs:

- Enable trace flag 3226. This stops SQL Server from printing backup success message in the error log. If you have several databases that have very frequent log backups, these success messages can flood your error log, leading to very large error log files and difficulty in spotting really useful error log messages. SQL Server will always report when a backup fails, so we don't need the success messages. Every SQL Server instance in the world should have this trace flag enabled at startup. Erin explains how to do this here.
- Increase the number of error logs that SQL Server can keep to the maximum of 99. You can do this easily in SSMS by right-clicking on the error log and selecting configure. See KB article 196909 for more details.
- Cycle the error log regularly – for example, every day. You can do this easily by creating an Agent job that runs *sp_cycle_errorlog* every day (e.g. at midnight). See KB article

2199578 for more details (which actually links back to the post above as well as a blog post of mine!)

Bottom line: Huge error log files are hard to manage, hard to search, and unnecessary to keep for extended periods, so take these easy steps to make your error logs more manageable.

Paul's Ponderings

One of the questions I was asked in email last week was from someone who'd deleted 500GB of data from a table (in a single-file database) and wanted to shrink out that space so another database can use it. They'd tried just running *DBCC SHRINKDATABASE* without any options and running it with target sizes, to try to make it runs faster, but nothing worked. It just ran very, very slowly and they wondered why.

The unfortunate answer is that the data-file shrink code performs very poorly under many circumstances, and there was never any love given to the code after it was re-written for SQL Server 2000. At one point I had a design to make it move up to 32 pages at a time, instead of just one, but it wasn't implemented. It just wasn't considered worth the engineering effort as you shouldn't regularly be shrinking a database. However, when you really need it…

The basic shrink operation (all methods of running a shrink use the same code) is:

- Find the page that's allocated nearest the end of the data file
- Move it as far towards the start of the data file as possible
- Fully log the page move, and all linkage updates

There are some algorithmic reasons why this page move operation may be very slow:

1. Moving any data file page in shrink requires acquiring an exclusive lock on the page. Shrink will wait forever for that lock to be acquired, which can obviously lead to a long-running operation.
2. Moving a text page requires a table scan. This is because there is nothing in a text page which indicates which data/index/text-tree records point to the text records on the text page being moved, so a scan is required to find the records that need to be updated with the new location of the moved text page. This has always been the case. I investigated adding a back-link added to text records, as that would make shrink faster as well as *DBCC CHECKDB*, under some circumstances, but that wasn't done because it would have been detrimental to the performance of all other text-related operations.
3. Moving a heap data page requires the query processor to update the nonclustered index records that link to the data records on the page being moved. This used to be done under the covers by the Storage Engine before SQL Server 2005, and was much more efficient – now that's not possible any more.

Issue #1 is easy to diagnose though – look in the output of *sys.dm_exec_requests* for the shrink task (the *command* column will say *DbccFileCompact*) and if it's *SUSPENDED* and the *wait_type* column is *LCK_M_X* then you can see who's blocking shrink in the *blocking_session_id* column.

For issues #2 and #3 you need to know what's in your database to be able to tell if these may be affecting the run time of a shrink operation. You can do this by running a simple query like the following:

```
SELECT
      [f].[name] AS 'Filegroup',
      [o].[name] AS 'Object',
      [p].[index_id],
      [a].[type_desc] AS 'Type',
      [a].[total_pages] AS 'Pages'
FROM [sys].[system_internals_allocation_units] AS [a]
      JOIN [sys].[partitions] AS [p]
            ON [p].[partition_id] = [a].[container_id]
      JOIN [sys].[filegroups] AS [f]
            ON [a].[filegroup_id] = [f].[data_space_id]
      JOIN [sys].[objects] AS [o]
            ON [o].[object_id] = [p].[object_id]
WHERE
      ([p].[index_id] = 0
      OR [a].[type_desc] = N'LOB_DATA'
      OR [a].[type_desc] = N'ROW_OVERFLOW_DATA')
      AND [a].[total_pages] > 0;
GO
```

This will show you how many heap data pages and text pages you have in your database, and in which filegroups.

Call to action: Apart from the obvious consideration of whether you can perform a shrink operation in a more efficient way (see this blog post), if you're going to have to run a shrink operation then be prepared for it to take a long time. If you can work out in advance that the shrink is going to take a very long time, it might be better to export everything into a new database, or schedule repeated shrink operations over a week or two with a progressively smaller target size, so you can fit the shrink operation in your maintenance window or periods of lighter load. And if you do use shrink, don't forget about the potential for index fragmentation which might have a detrimental effect on performance (detailed in the same blog post referenced here).

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.

Seagate Multi-Actuator Technology (MAT) for Hard Drives

Seagate has introduced its new Multi-Actuator Technology (MAT) for hard drives. This technology will improve conventional magnetic hard drive performance by using two or more sets of actuator arms that operate independently from a single pivot point. This will result in better random IO performance, which will scale with the number of actuator arms.

These new hard drives, which are scheduled to go into production in 2018, will also be able to use SAS, SATA interfaces, and also use the NVMe protocol. These new drives are meant for data center use, to go along with higher capacity points of up to 20TB by 2019. They should provide better performance than current hard drives, and be less expensive per GB than some types of flash storage.

SQL Server 2017 CU3 Released on January 4, 2018

On January 4, 2018, Microsoft released SQL Server 2017 CU3, which is Build 14.0.3015.40. By my count, this CU has sixteen public hotfixes, many of which are for the SQL Engine or SQL performance. There are also some new manageability and programmability features that have been added, such as support for the *MAXDOP* option for *CREATE STATISTICS* and *UPDATE STATISTICS*.

In addition, Microsoft has included the security fixes from the January 3 SQL Server security update in this Cumulative Update. Microsoft has very detailed guidance on how this may affect SQL Server here.

Since SQL Server 2017 won't be using Service Packs as part of its servicing mechanism, you will need to start testing and deploying Cumulative Updates on a schedule that makes sense for your organization.

As always, I think it is a good idea to make an effort to stay current on Cumulative Updates, as does Microsoft.

New WPA3 Wireless Network Security Standard

The Wi-Fi Alliance has announced its next generation WPA3 security standard, which will replace the existing WPA2 standard that was introduced back in 2004.

One of the key improvements in WPA3 will help make open Wi-Fi networks more secure by using individualized data encryption, which scramble the connection between each device on the network and the router. WPA3 will also be more secure against brute-force dictionary attacks designed to guess router passwords.

It will also make it easier to improve security on IoT devices by making it easier for you use a phone to make configuration changes to IoT devices that don't have a display. The new WPA standard is expected to show up in devices later in 2018.

Spectre and Meltdown Exploits

News started breaking on January 2, 2018 about several new ways to access privileged information by taking advantage of how many modern CPUs normally operate in order to improve performance. These exploits are known as Spectre variant 1, Spectre variant 2, and Meltdown.

Since then, I have done quite a bit of research and testing, which have resulted in several blog posts so far:

- Microsoft SQL Server Updates for Meltdown and Spectre Exploits
- Checking Your Meltdown and Spectre Mitigation Status in Windows
- Performance Effects of Meltdown and Partial Spectre Fixes on Intel Core i7-7500U Laptop

Microsoft has published several useful posts on this subject:

- Understanding the performance impact of Spectre and Meltdown mitigations on Windows Systems
- SQL Server Guidance to protect against speculative execution side-channel vulnerabilities
- Windows Server guidance to protect against speculative execution side-channel vulnerabilities
- Windows Client Guidance for IT Pros to protect against speculative execution side-channel vulnerabilities

#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 trace flags, to match the earlier Curious Case:

- Erin's blog posts:
  - *SQLskills SQL101: Trace Flags*
  - *Trace Flag Information in Query Plans*
- My blog posts:
  - *Misconceptions around TF 1118*
  - *The pros and cons of trace flags*
- An interesting discussion on StackExchange where Paul White and I expound on how to find trace flags
- A really cool github repository or trace flags and trace flag information

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

- Paul: *Calling all user group leaders! We want to present for you in 2018!*
- Paul: *SQLintersection Spring 2018*
- Paul: *New Pluralsight course: Automatic Tuning in SQL Server 2017 and Azure SQL Database*
- Glenn: *Checking Your Intel Processor Features Regarding the Meltdown Exploit*
- Glenn: *SQL Server Diagnostic Information Queries for January 2018*
- Glenn: *Performance Effects of Meltdown and Partial Spectre Fixes on Intel Core i7-7500U Laptop*
- Glenn: *SQL Server 2017 CU3 Released on January 4, 2018*
- Glenn: *Checking Your Meltdown and Spectre Mitigation Status in Windows*
- Glenn: *Microsoft SQL Server Updates for Meltdown and Spectre Exploits*

I hope you find these useful and interesting!

Video Demo

In the first recorded demo of 2018, Jonathan shows how to isolate the memory usage of an In-Memory OLTP enabled database for the In-Memory objects to a specific resource governor resource pool for accurate tracking of memory usage at the database level.

The video is just over five minutes long and you can get it:

- In MOV 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 SQLintersection in March to our own in-person Immersion Events in the U.S. and London; all events are open 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
- So why do you want to come to our training? And the winners are…
- Community blog posts about our classes
- Immersion Event FAQ

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