

(December 5th, 2016)

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 Bali, where we're adjusting to the time zone in Indonesia before embarking on a dive trip around the Komodo area. We'll post some photos and videos later this month and in January.

We hope you'll join us next May in Orlando, FL for our Spring SQLintersection (check out www.SQLintersection.com for more details) and remember that all of our 2017 classes have discounts of up to \$200 for registering before the end of this year (see [here](#) for the schedule).

And even if you can't join us in person, I've put out our annual call to user groups to get remote sessions scheduled for 2017. This year we'll end up having done 94 user groups and PASS Virtual Chapter presentations across the team, up from 86 in 2015, which I think is a tremendous achievement. If you'd like one of us to present for your user group, check out my blog post [here](#).

The latest book I've read is Greg Iles' [The Bone Tree](#). This is the follow-on from the fantastic [Natchez Burning](#), which was my favorite book of 2014, and you really need to read that before this one, as this one continues the story from the next day. This book starts to pull the characters into the Kennedy assassination investigations and links to the extreme racist murders from the 1960s. It's hard to say more without giving away some of the plot and twists. The writing is excellent, with great dialog, and the book is a real page-turner. Highly recommended!

Note: you can get all the prior Insider newsletters [here](#).

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.

I was working with a client recently (actually they were attending our IEPTO1 class in Chicago a couple of weeks back) who was having fragmentation problems for a table that they couldn't figure out. The clustered index key was an identity column and they weren't updating any of the key values.

We discussed the usage pattern of the table and it turned out that there was a LOB column in the table that was being updated to be longer after it had its initial value populated. As the clustered index had a 100% fill factor in use, there was no space available on the data pages for the longer columns, and so page splits were happening.

The fix was to turn on the *large value types out of row* setting for the table, so all new rows had their LOB fields stored outside the data pages, and updates to them would not cause page splits of the data pages.

Bottom line: anything that makes a row longer has the potential to cause a page split, and using a fill factor isn't always the best way to avoid page splits.

Paul's Ponderings

Following on from the SQL Server 2016 SP1 announcement a couple of weeks ago, this week's ponderings are from Glenn as he discusses the hardware ramifications of the changes to Standard Edition. Enjoy!

One thing that did not change with SQL Server 2016 Service Pack 1 was the license-related memory and compute limits for any of the non-Enterprise Edition SKUs. Specifically, SQL Server 2016 Standard Edition SP1 is still limited to 128GB of RAM for the Database Engine. Socket and core limits also remain the same, as the lesser of 4 sockets or 24 physical cores for non-virtualized environments, and the lesser of 4 sockets or 24 logical cores for virtualized environments. These license limits are per instance of SQL Server 2016 and [are fully documented here](#).

One welcome, but somewhat obscure change with SQL Server 2016 Service Pack 1 is that there are now separate per-database memory allocation limits for In-Memory OLTP usage and per-instance allocation limits for columnstore indexes. You can have up to 32GB of memory allocated, per database for In-Memory OLTP objects in Standard Edition, [as documented here](#). You can also have up to 32GB of memory allocated per instance, for columnstore index usage, [as documented here](#). These limits are in addition to the 128GB buffer pool limit for Standard Edition.

If you're going to use SQL Server 2016 Standard Edition, you need to be aware of these limits as you select and size your host hardware, and as you configure your virtual machines. As always, single-threaded CPU performance is critically important for overall database server performance, especially for OLTP workload, where many queries run on a single thread of execution.

Firstly, processor choice...

If you want the most overall CPU capacity possible for a single, non-virtualized instance of SQL Server 2016 Standard Edition, you should be looking at a two-socket server (such as a Dell

PowerEdge R730) with two, [Intel Xeon E5-2687W v4](#), twelve-core processors. This particular processor has a TDP of 160 watts, so it requires upgraded power supplies in most servers. It has a base clock speed of 3.0GHz and a Turbo Boost clock speed of 3.5GHz, making it a far better performance choice than the same-generation twelve-core [Intel Xeon E5-2650 v4](#), that has a base clock speed of 2.2GHz and a Turbo Boost clock speed of 2.9GHz.

If you have a smaller workload, and you want even better single-threaded performance (and lower SQL Server 2016 licensing costs), then you should consider a lower core count, “frequency-optimized” processor. The best choices are the [Intel Xeon E5-2637 v4](#), the [Intel Xeon E5-2643 v4](#), and the [Intel Xeon E5-2667 v4](#).

These four “good” processors are compared in the table below.

Processor Model	Estimated TPC-E Score	Estimated Score/Core	Standard Edition License Cost
Xeon E5-2687W v4	3673.00	153.04	\$44,616.00
Xeon E5-2667 v4	2611.91	163.24	\$29,744.00
Xeon E5-2643 v4	2081.36	173.44	\$22,308.00
Xeon E5-2637 v4	1428.39	178.54	\$14,872.00

The estimated TPC-E score and estimated TPC-E score/core illustrate the overall CPU capacity of the system and the single-threaded CPU performance, respectively. These figures come from some proprietary analysis and calculations that I have done, based on actual TPC-E benchmark submissions. The license costs are for a two-socket server, with two of the chosen processors.

Secondly, memory sizing and allocation...

Prior to SQL Server 2016 SP1, I would typically recommend that people have between 192-256GB of RAM in a non-virtualized, two-socket host for SQL Server 2016 Standard Edition (or for SQL Server 2014 Standard Edition). This would allow you to set max server memory to 131072MB (128GB), and still have plenty of RAM left over for the OS and other SQL Server components.

This “excess” RAM recommendation is also due to the fact that current generation two-socket servers have 24 DIMM slots, and how you should populate them for the best performance. For a virtual machine, I would recommend having a little bit less RAM, perhaps in the 144-192GB range, while being careful to not have more memory than will fit in a single NUMA node.

With SQL Server 2016 Standard Edition SP1, I would raise these memory sizes, given the separate memory usage by columnstore indexes (32GB per instance) and by In-Memory-OLTP (32GB per database). Depending on whether you are going to be using columnstore indexes and

In-Memory OLTP, and how many databases you will have per instance, I think a recommendation of 256-512GB is more appropriate.

Obviously, you can go lower on these memory recommendations if you have smaller databases and a less intense workload. On the other hand, server-class DDR4 RAM is very affordable, at about \$7.00/GB, so even 512GB of RAM is going to be a small fraction of your SQL Server 2016 Standard Edition license costs in most scenarios.

Call to action: make sure you're taking into account the SQL Server Edition and features you're going to use when designing a new server for SQL Server 2016.

Video Demo

In this video from my Pluralsight course [SQL Server: Detecting and Correcting Database Corruption](#), I show how to figure out the optimal restore sequence to use based on the Log Sequence Numbers in the various backups you have available.

The video is about ten minutes long and you can get it in WMV format [here](#).

The demo code is available [here](#).

Enjoy!

SQLskills Offerings

We've announced the first half of our 2017 class lineup, with discounts available on all classes for registrations/payments received before January 1st, 2017.

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](#)

Upcoming Immersion Events

Chicago, IL, April/May 2017

- **IE0:** Immersion Event for Junior/Accidental DBAs
 - April 24-26 **US\$120 discount for registering in 2016!**
- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
 - April 24-28 **US\$200 discount for registering in 2016!**
- **IESSIS1:** Immersion Event on Learning SQL Server Integration Services
 - April 24-28 **US\$200 discount for registering in 2016!**

- **IEBI:** Immersion Event on Business Intelligence
 - May 1-5 **US\$200 discount for registering in 2016!**
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
 - May 1-5 **US\$200 discount for registering in 2016!**
- **IESSIS2:** Immersion Event on Advanced SQL Server Integration Services
 - May 1-5 **US\$200 discount for registering in 2016!**
- **IEPS:** Immersion Event on Powershell
 - May 8-10 ****NEW** class, US\$120 discount for registering in 2016!**
- **IEPDS:** Immersion Event on Practical Data Science
 - May 8-12 **US\$120 discount for registering in 2016!**
- **IEHADR:** Immersion Event on High Availability and Disaster Recovery
 - May 8-12 **US\$200 discount for registering in 2016!**

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