

(December 14<sup>th</sup>, 2015)

If you know someone who you think 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're winding down and looking forward to heading away for Christmas week, with the kids and Kimberly's Mom, to the Olympic Peninsula for some relaxation in beautiful natural surroundings.

So, that's it for this year – the next newsletter will be on Monday, January 4<sup>th</sup>. Kimberly and I would like to take this opportunity to thank all of you who read the newsletter, read our blogs, watch our videos, come to our classes, or do anything related to SQLskills – we do it all for you! (Btw – if you've never heard it before, [check out Kimberly and I singing](#) our Christmas messages for you back in 2010). Happy holidays!

**Our 2016 classes** are filling up in Chicago (April/May), London (June), and Bellevue (September), including two brand new classes on SSIS, and we're running an expanded (to five days) IEPDS class on Practical Data Science. **You can save up to \$200 off all our 2016 classes by registering before the end of the year** and as an alumnus, you can always get the lowest (frequent-flyer) discount! Schedule details are [here](#) and you can read more on the new SSIS classes [here](#). **Note that IEPTO1 in Chicago only has 5 seats remaining!**

**Also, if you're a user group leader (or know one)**, check out my [blog post about 2016 remote presentations](#) – we want to present to your user group!

The two latest books I've read are Daniel Silva's [The English Assassin](#), and Stephen Tanner's [Afghanistan: A Military History from Alexander the Great to the War against the Taliban](#).

Silva's book is the second in his long spy series about an Israeli intelligence agent and is just as good as the first one, [The Kill Artist](#), which I profiled back in [August](#). This book sees the protagonist investigating art looted by the Nazis in WWII and stored in Swiss bank vaults. Fast paced and exciting.

Tanner's book is a very interesting history, which feeds into my fascination with that area of the world and all the strife it's endured. It's not overly deep in any one area (I have other books to read later on things like The Great Game between Britain and Russia in the 1800s) but provides a great general coverage of all the major conflicts through to post-9/11.

Both books are highly recommended. Don't forget to check out my end-of-year book summary that I'll be posting on December 31<sup>st</sup>!

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

## **Paul's Ponderings**

The last Ponderings of the year!

Last week I did my 29<sup>th</sup> user group presentation of the year! I presented to the PASS DBA Fundamentals Down Under virtual chapter on index fragmentation, and there were some cool questions so I thought I'd present a few of them here with some extended answers.

Q1: Given that GUIDs are generated per server, if a system is using peer-to-peer replication and the GUID is the cluster key of a replicated table, how do we avoid GUID collisions with the same GUID being generated by different servers?

A1: (Not a fragmentation question, but interesting!) Although a collision is *\*highly\** unlikely, it's still a possibility I suppose. I wouldn't use a GUID as the cluster key in a peer-to-peer replicated table where all servers can create new rows. Instead I'd use something like a composite key comprised of a number or small char column that identifies the server that created the row, and then an identity column. This guarantees that the rows generated on each server are unique across all the servers, and allows the generating server to be identified.

Q2: Can the `sys.dm_db_index_physical_stats` DMV be run on a database snapshot?

A2: Yes, but don't do that. Database snapshots have a nasty inefficiency around memory usage. Basically, every page in the buffer pool has to be owned by a database. The source database and snapshot databases have different database IDs, so a query in the context of the snapshot database can only read pages owned by the snapshot database. This means that even if the pages required by the query haven't changed in the source database, they need to be read from disk again, into the buffer pool, so they can be in buffers owned by the snapshot database. So running the DMV in the snapshot database may end up taking a lot of buffer pool space, and be less efficient than just running it in the source database itself. You can always restore a copy of the database somewhere else and run the DMV on another server, using the results to drive index maintenance on the main server.

Q3: How can you find progress reports for index maintenance operations?

A3: For a reorganize, look in the `percent_complete` column in `sys.dm_exec_requests`. For a rebuild, you can find the number of pages in the index and then monitor how many page reads the rebuild operation has done. For an online index operation, you can also look at the trace event Progress Report: Online Index Operation where the `bigintdata1` column is the number of

rows process by the internal version scan of the old index. There's also an extended event in 2014 that shows this data too.

Q4: Can you use an index rebuild on a 24x7 system if you have a mid-tier cache?

A4: Yes, I suppose you could if the cache can satisfy all queries that come in while the rebuild is going on, noting that it won't be able to satisfy insert/update requests. I haven't heard of anyone doing this though, as you can reorganize an index without it taking blocking locks (schema-modification locks, that both offline and online index rebuild operations require).

**Call to action:** Make sure you're taking note of index fragmentation on your systems, in terms of the performance hit from page splits causing it, the potential performance hit on scan performance, and the potential performance hit of having lots of wasted space on pages (see [this blog post](#) for details).

### **Video Demo**

From Erin: At the end of November, Service Pack 3 for SQL Server 2012 was released, and it included several improvements in query plan diagnostics, including the ability to see how many rows were actually read when predicate pushdown occurs. In this newsletter's video I'll cover why this information is important, and take a look at where to find it in the query plan.

The video is just under 10 minutes long and you can get it:

- In WMV format [here](#).
- In MOV format [here](#).

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

Enjoy!

### **SQLskills Offerings**

We've released the majority of our 2016 classes for registration, listed below. It's possible that we might add one or two classes in Chicago in the Fall, but that will depend on the Fall conference schedule as well as demand. **Note that you can save up to \$200 by registering for these classes before the end of this year.**

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, 2016

- **IE0:** Immersion Event for Junior/Accidental DBAs
  - April 25-27 (**US\$120 discount for registering in 2015**)
- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - April 25-29 (**US\$200 discount for registering in 2015**)
  - **NOTE: Only five seats remaining!**
- **IEBI:** Immersion Event on Business Intelligence
  - April 25-29 (**US\$200 discount for registering in 2015**)
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
  - May 2-6 (**US\$200 discount for registering in 2015**)
- **\*\*NEW\*\* IESSIS1:** Immersion Event on Learning SQL Server Integration Services
  - May 2-6 (**US\$200 discount for registering in 2015**)
- **IEPDS:** Immersion Event on Practical Data Science
  - May 9-13 (**US\$200 discount for registering in 2015**)
  - **NOTE: This is now five days in length.**
- **IEHADR:** Immersion Event on High Availability and Disaster Recovery
  - May 9-13 (**US\$200 discount for registering in 2015**)

London, UK, 2016

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - June 13-17 (**US\$200 discount for registering in 2015**)
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
  - June 20-24 (**US\$200 discount for registering in 2015**)

Dublin, Ireland, 2016

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - Dates and details coming soon; we're targeting October 2016.

Bellevue, WA, 2016

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - September 12-16 (**US\$200 discount for registering in 2015**)
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
  - September 19-23 (**US\$200 discount for registering in 2015**)
- **\*\*NEW\*\* IESSIS2:** Immersion Event on Advanced SQL Server Integration Services
  - September 19-22 (**US\$160 discount for registering in 2015**)

See [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](mailto:Paul@SQLskills.com) and [Kimberly@SQLskills.com](mailto:Kimberly@SQLskills.com)