(February 15th, 2016)

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 back from our dive trip and heads-down creating new Pluralsight courses – there are three new ones in the pipeline for release in March!

Our 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 (the next big thing!). Schedule details are here and you can read more on the new SSIS classes here.

We've actually **added another Chicago IEPTO1 during the week of May 16th** as the one in April has sold out (along with IE0 and IEBI), and the new IEPTO1 is already half full. We've also added an **IEPTO1 course in Dublin** in October – see here for details.

Our full Spring SQLintersection conference line-up has been published and it's looking to be a phenomenal show – one of our best for sure! We have keynotes from Buck Woody and Bob Ward from Microsoft, and another on tools and productivity from SQL Sentry. And with 9 workshops over 4 pre-conference and post-conference days, deeper dives are possible. Best practices, troubleshooting, tips and tricks are our mainstay yet this show also adds a lot of new features for SQL Server 2014 and 2016. The show is in Orlando, FL over April 16-22, 2016; we hope to see you there! Be sure to use the discount code 'SQLskills' when registering to receive $50 off. See here for more details.

The latest book I've read is James Gleick's *The Information: A History, A Theory, A Flood*. This is an excellent history of the methods of dissemination of information (think printing press, visual telegraph, Morse code), and on the creation and development of the various facets of information theory, including quantum computation and genetics. It's a dense book, but very readable and immensely interesting. Highly recommended!

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

Paul's Ponderings

This time it's another topic drawn from several email questions I received. They boiled down to this (paraphrasing):

I've got a large heap table where the space isn't being given up when I delete a large number of records, but then when I shrink the database the heap is reduced in size. Can you explain?

The behavior you're seeing is how SQL Server works, but it's pretty non-intuitive. When a page in an index becomes empty, it's always deallocated (as an empty page isn't permitted in a SQL Server index structure from SQL Server 2005 onwards). However, the structure of a heap is different and as a result, the behavior is too.

Whenever a row is deleted in a heap, it's most likely that the page does not become empty. However, if the page that the row is stored on becomes empty as a result of the delete, the page cannot be deallocated from the table unless an exclusive table lock is held (to remove the page from "tracking"). This is usually not the case, unless lock escalation has occurred because you're deleting enough rows to trigger escalation, or if you specifically use the *TABLOCK* hint on the delete statement, for instance. But, because all of these things are unlikely, the page usually cannot be deallocated.

There is a Knowledge Base article that describes this phenomenon: KB 913399. However, the KB article only references up to and including SQL Server 2005 but this behavior exists in more recent releases too and is very easy to reproduce if you want to prove it to yourself.

The empty pages will be reused by subsequent inserts, but if the space isn't going to be reused following a large delete in a heap, you might consider using the *TABLOCK* hint to allow the empty pages to be deallocated and the space made available for other objects in the database to use.

Another alternative is to just use a clustered index instead, or if a heap is necessary, you could rebuild the heap using *ALTER TABLE ... REBUILD* (that was added in SQL Server 2008 to support enabling compression on a heap), with the caveat that this will cause all the table's nonclustered indexes to be rebuilt.

On the extreme end (in my opinion), you could reclaim the empty heap space using a shrink operation. Shrink won't free up space inside pages as it moves them (with the exception of compacting LOB pages as it goes – somewhat unsuccessfully depending on which version and build you're on – see KB 2967240), but it will remove empty pages rather than moving them. This will effectively shrink the heap after a large delete, but with the usual caveats about shrink causing index fragmentation and generally being an expensive, slow operation to perform.

Call to action: Not really a call to action, but one more thing to be aware of around how SQL Server works and understanding the behavior you're seeing when using heaps instead of clustered indexes.

Video Demo

In this week's insider video, Glenn demonstrates how to use Microsoft DiskSpd to test your I/O subsystem, and explains what the basic command line switches do. You can also read an article Glenn wrote about DiskSpd over on SQLPerformance.com here.

The video is about 3.5 minutes long and you can get it:

- In WMV format here.
- In MOV format here.

No demo code this time.

Enjoy!

SQLskills Offerings

We've released all of our 2016 classes for registration, listed below. It's possible that we might add one or two classes in Chicago in November, but that will depend on the Fall conference schedule as well as demand.

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

- **IE0**: Immersion Event for Junior/Accidental DBAs
  - April 25-27 **SOLD OUT!**
- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - April 25-29 **SOLD OUT!**
- **IEBI**: Immersion Event on Business Intelligence
  - April 25-29 **SOLD OUT!**
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - May 2-6
- **IESSIS1**: Immersion Event on Learning SQL Server Integration Services
  - May 2-6 **NEW!!**
- **IEPDS**: Immersion Event on Practical Data Science
  - May 9-13 **HOT!!**
  - **NOTE: This is now five days in length.**
- **IEHADR**: Immersion Event on High Availability and Disaster Recovery
  - May 9-13
- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - May 16-20 **NEW!!**

London, UK

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - June 13-17
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
  - June 20-24

Bellevue, WA

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - September 12-16
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
  - September 19-23
- **IESSIS2:** Immersion Event on Advanced SQL Server Integration Services
  - September 19-22 **NEW!!**

Dublin, Ireland

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - October 3-7 **NEW!!**

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 and Kimberly@SQLskills.com