

(September 26<sup>th</sup>, 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 Bellevue, WA where we've just finished two weeks of very successful classes, and tomorrow we're off to Dublin. Incidentally, if you're in Europe, the Dublin IEPTO1 class next week (starting Monday, October 3rd) is the **last time we'll be teaching in Europe until 2018** (note: 2 seats were available at the release of this newsletter) – see [here](#) for details.

SQLintersection session and workshop details have been posted; we have an exciting line-up of speakers scheduled for our show and we're also running our fantastically fun evening event: SQLafterDark! If your team needs architectural advice and sessions from speakers who not only know their technology but know how to convey it – this is the place to be! Be sure to use the 'SQLskills' discount code to save \$50 on registration. Check it out [here](#).

**We've just announced the first half of our 2017 class lineup:** seven classes in three weeks, including a **brand new, 3-day course on PowerShell**. All courses have a discount (US\$120 for 3-day courses; US\$200 for 5-day courses) for registrations and payments received before January 1st, 2017. Check out all the dates and details on our [class schedule page](#).

The latest book I've read is Neal Asher's [The Departure: The Owner: Book One](#), the first in a trilogy. Asher is one of my favorite sci-fi authors and most of his novels are set in his [Polity universe](#). This one is set on Earth a hundred years or so in the future, where there's a single brutal government, and the colony on Mars that's just been abandoned by Earth. The protagonist wakes up inside a sealed box on a conveyor belt leading to an incinerator and has to figure out his previous life and then start working on revenge. Lots of action, cool machines and robots, futuristic technology and all very fast paced. I can't wait to read the next two – 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.*

Jonathan was working with a client recently who was having slowly degrading performance in their entire workload. The server had 64GB of memory installed, with 53GB configured for max server memory, following generally-recommended practices. Looking at the wait statistics, he saw that nearly every query was driving physical reads, which was overloading the I/O subsystem and causing long *PAGEIOLATCH\_SH* waits. The page life expectancy on the server was in 100-200s, showing signs of buffer pool churn and contention on the system.

The next step was to check the contents of the buffer pool and see what was using the memory internally in SQL Server, using a script based on code from [here](#). Jonathan was surprised to see that the distribution database (the server was a replication Distributor as well as serving other roles) was using nearly 44GB of the buffer pool space with the user databases consuming less than 6GB of buffer pool.

He then checked the replication settings and found that someone had set the Distributor to have a retention period of 72 days instead of 72 hours! This was for a single subscriber for reporting that was in the same data center and was only ever offline for routine patching every quarter.

Changing the distribution retention back to hours from days and running the *Distribution Clean Up: Distribution* job on the server reduced the size of the *MSrepl\_commands* table (which is scanned by the Distribution Agents for subscriptions) and allowed the user databases to return to normal buffer pool usage.

It's always interesting to see how changing a setting in one area (albeit accidentally) can have such a profound effect on the overall health of a SQL Server instance. Make sure that you understand all the ramifications of any configuration change that you make.

### **Paul's Ponderings**

One of the topics that I discussed in class last week is why the query optimizer doesn't know (or care) what's in the buffer pool.

Here's a scenario: table T has two nonclustered indexes, A and B, that both cover query Q, and which require a complete index scan in both cases. Index A has 10,000 pages at its leaf level, and index B has 50,000 pages at its leaf level.

Which index should the optimizer use when compiling the query plan?

SQL Server uses a cost-based optimizer, which uses various metrics and statistics to determine the most efficient query plan for the query (given the time limits imposed on its search of the space of all possible query plans). The 'cost' in 'cost-based' means that it considers the CPU cost and I/O cost of the various operators in the query plan, with the I/O cost essentially being relative to the number of physical reads required.

In the scenario above, the optimizer will choose a query plan involving index A. As both indexes cover query Q, the most efficient plan will be the one involving the fewest reads.

Now let's allow the optimizer to base plan choice on what's in the buffer pool...

If index A is mostly not in the buffer pool and index B is mostly in the buffer pool, it would be more efficient to compile the query plan to use index B, for a query running at that instant. Even though index B is larger, and would need more CPU cycles to scan through, physical reads are waaaaay more expensive than CPU cycles so a more efficient query plan is the one that minimizes the number of physical reads.

This argument only holds, and a 'use index B' query plan is only more efficient than a 'use index A' query plan, if index B remains mostly in memory, and index A remains mostly not in memory. As soon as the relative proportions of indexes A and B that are in memory become such that the 'use index A' query plan would be more efficient, the 'use index B' query plan is the wrong choice.

The situations when the compiled 'use index B' plan is less efficient than the cost-based 'use index A' plan are (generalizing):

- Indexes A and B are both memory resident: the compiled plan will use roughly 5 times more CPU than the optimal plan, as there are 5 times more pages to scan.
- Neither index is memory resident: the compiled plan will do 5 times the number of physical reads AND use roughly 5 times more CPU.
- Index A is memory resident and index B isn't: all physical reads performed by the plan are extraneous, AND it will use roughly 5 times more CPU.

So although the optimizer can make use of buffer pool contents knowledge to compile a query that is the most efficient at a single instant, it would be a very dangerous way to drive plan compilation because of the potential volatility of the buffer pool contents, making the future efficiency of the cached plan highly unreliable.

**Call to action:** Although it doesn't always get it right, the optimizer strives to produce the most efficient plan, assuming nothing is in the buffer pool. Understanding how the query optimizer comes to plan choice decisions is extremely useful for understanding query plans themselves and relating them to the code driving the plan. Not really a call to action, more a statement that reinforces a saying that applies strongly to SQL Server performance tuning: the more you know, the further you'll go.

### **Video Demo**

Getting started with Azure SQL Database doesn't have to be complicated. If your database passes the Azure SQL Database validation check, you can easily migrate it to Azure using a

bacpac. In this video, Tim steps through the process of using the built in option “Deploy Database to Microsoft Azure SQL Database” and the “Export Data-tier Application” option.

The video is just over 6 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 just announced the first half of our 2017 class lineup, with discounts available on all classes for payments received before January 1st, 2017. And, don't forget we still have three classes coming up in Chicago in November this year, plus SQLintersection in Las Vegas in October. We hope to see you at a class or a conference soon!

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**

Dublin, Ireland (*returning to Europe in 2018 NOT 2017*)

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

Chicago, IL, November 2016

- **IE0:** Immersion Event for Junior/Accidental DBAs
  - November 7-9
- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - November 7-11
- **IEPDS:** Immersion Event on Practical Data Science
  - November 7-11

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