

(June 20th, 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 London where we're starting IEPTO2, after a great IEPTO1 last week. We've done some interesting sightseeing while we've been here (Kensington Palace, National Portrait Gallery, and a Houses of Parliament tour) and yesterday we saw an excellent live performance of the hit comedy show *Whose Line Is It Anyway?*

Our Fall classes are also starting to fill up, so be sure to book your Immersion Event in advance to avoid the class you want selling out – all the schedule details are [here](#). After another incredibly successful delivery, we've added an IEPDS class on data science, R, and data mining to our Chicago lineup in November. This is one of the next big things in data and a must for those of you looking to expand your knowledge/resume!

We have a special offer running through June 30, where any client – new and existing – can get a **single-instance health check for only US\$2,500**, which is more than 1/3 off the normal price! Details about why our health checks are so cost-effective are [here](#).

And some **exciting news** – I've finally released my comprehensive wait stats and latch classes library! Check out the details [here](#).

The latest book I've read is Jeremy Scahill's [*Dirty Wars: The World Is a Battlefield*](#). This is a really interesting and disturbing book that investigates the ongoing, autonomous U.S. military operations around the world that have been undertaken as part of the war on terror. It exposes the extrajudicial capture/torture/kill program that was created in secret by the Bush administration (and continued by the Obama administration), and operated by the CIA and various Special Forces in Yemen, Iraq, Pakistan, Afghanistan, and Somalia. It's a huge book that's full of factual information without engaging in partisan criticism – we're left to draw our own conclusions. Highly recommended!

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

The Curious Case of...

We've added this new section to the newsletter to explain problems we've found on client systems; it might be something you're experiencing too.

Tim was working with a new client who had recently experienced corruption and found that their log backups were only happening every 45 minutes or so, which meant they lost more than their planned 15 minutes of work (according to the their 15-minute data loss Service Level Agreement).

After some investigation, Tim found out that the server involved was the result of a poorly-planned consolidation project, and had more than 500 databases on it. The log backup job was a permanently-running process, which did the following:

- Start a timer
- Perform log backups on all databases
- Wait until the timer reaches 15 minutes
- Loop

The problem was that it was taking more than 40 minutes for each loop of the job, so the data loss SLA was being broken for all the databases.

After discussing with the client, all but 40 of the databases were discovered to have static data so could be switched to the simple recovery model. This allowed the log backup job to run in less than 15 minutes. Arguably a better solution might have been to individually schedule some or all of the log backups, but the client wanted to stay with the existing job.

The lesson from this story is that whenever an environmental change occurs, the HA/DR strategy needs to be re-evaluated to make sure that downtime and data loss SLAs can still be met.

Paul's Ponderings

As we've been teaching all last week, the newsletter this time is a mish-mash of interesting tidbits for you to check out...

First up is an interesting discussion I was involved in over on Stack Exchange, where someone was interested in ways to find new trace flags that are added to SQL Server. There's really no good way to do it, as SQL Server doesn't have a central list that it checks against so there's nothing you can search for in a DLL, but Paul White and I provided some insight. Check it out [here](#).

Next, we had a long discussion in class around using the *fn_dump_dblog* function to examine the transaction log from log backups, both for analyzing fragmentation-causing page splits and for finding the point in time to restore to after someone's inadvertently dropped a table. The code to do these operations is at [here](#) and [here](#), respectively.

We also had an interesting discussion about the inefficiencies of the allocation bitmaps in SQL Server data files, and how you can have problems with multi-terabyte databases. One case in particular that we discussed in class was from a prior client of ours (who had a couple of

attendees in the class last week) where they occasionally were seeing spikes in the elapsed time for inserts to happen. It turned out that this was on a very large table, where the IAM chain was many hundreds of pages long, and occasionally most of the IAM pages would be aged out of the buffer pool. The next insert that needed an allocation would cause all the IAM pages to be read back into memory again, leading to the long insert time.

Their fix is to have an Agent job that periodically walks the IAM chain using *DBCC PAGE*, to ensure that the IAM pages' LRU values in the buffer pool are always recent enough that those pages are not dropped from the buffer pool when it's under memory pressure. This is a great example of a practical application of understanding SQL Server internals.

Finally, I've been hearing from people who've been kicking the tires of SQL Server 2016 with a variety of workloads and found that many operations are noticeably faster. It was very pleasing to see that the SQL did a lot of work on fixing performance bottlenecks in the code, rather than expending 100% of developer energy on new features. Bob Dorr and others have been blogging a series on the PSSSQL blog covering many of the scalability and performance enhancements and it makes for some very interesting reading – check it out [here](#).

Call to action: Nothing this time – lots of stuff for you to read!

Video Demo

In this video, Glenn demonstrates the performance benchmark feature of the Samsung Magician utility and compares the results to CrystalDiskMark. Vendor storage management utilities can be very useful for maintaining and monitoring the health and performance of that vendor's products. They can also be useful for running a quick performance benchmark test to confirm the performance levels reported by other benchmark tools.

The video just under 3 minutes long and you can get it:

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

No demo code this time.

Enjoy!

SQLskills Offerings

Our 2016 classes are all open for registration (listed below), including three new classes in Chicago in November added because of popularity. We hope to see you at a class or a conference 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

Bellevue, WA

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
 - September 12-16
- **IEBI:** Immersion Event on Business Intelligence
 - September 19-22 **Just added!!**
- **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 course!!**

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

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

Chicago, IL

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

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