

(January 19th, 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 I just got home from a week of diving with Jonathan in the Bahamas on Jon's first liveaboard dive trip (he's hooked!). Kimberly wasn't able to make this trip so that let me catch up 25 dives (to a total of 425) on her almost 300 dive head-start on me when I started diving in 2006.

Today, the big news is that we celebrate expanding our team with [Tim Radney](#) officially joining us!

Last week Erin and Glenn did the first of 55(!) scheduled remote user group presentations this year, and Kimberly and I are doing four more this week. We hope that you can join us along the way!

The most recent book I've read is Andrew Chaikin's [A Man on the Moon: The Voyages of the Apollo Astronauts](#). I've been a fan of all things to do with space and astronomy since I could read but I'd never read a comprehensive history of the Apollo program. This book does a great job of covering each Apollo mission (in varying amounts of detail) and makes their stories readable rather than just being a dry recounting of events. It also covers the astronauts themselves, their often quirky personalities, and how they got into the space program in the first place. I found it exciting reading about the events of around 50 years ago and I've always thought it such a shame that the space program was scaled back and allowed to languish once the "beat the Russians to the moon" goal was met. Highly recommended.

Please [let us know](#) if you liked what you read/saw here and/or have any suggestions for future Quick Tips.

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

Paul's Ponderings

One of the things that DBAs and developers don't think about very often (from experience) is the amount of data that's in a table, tables, or database. In general the data volume is just accepted – of course it's supposed to be there, it's in the database, right?

Well, no. If you step back and think about it, there are a couple of important questions that you should be asking:

1. Does all this data really *need* to still be in the production database?
2. Is all that data *supposed* to be in the production database?

If a production database is bloated with unnecessary data, there are a bunch of problems that can result, such as:

- Storage requirements are unnecessarily large (and don't just think of the database itself, also consider AG secondary replicas, a mirror database, replication subscription databases, log shipping secondary databases, test and development database copies)
- Database backups back up older data making them unnecessarily large and longer to execute
- Index fragmentation (e.g. index rebuilds) generates more transaction log than necessary (which can have a host of knock-on effects around log performance, log backup size, replication, and so on) and takes longer
- Statistics updates take longer
- Queries may process more data than necessary, affecting query plans, performance, buffer pool space, I/O performance, and more

The list goes on and on.

If the answer to those two questions is 'no' then they can be mitigated by, respectively:

1. Careful archiving of unneeded data on a regular basis, either through a partitioning process like the 'sliding window', or through data deletion that's architected to happen in manageable chunks (rather than one enormous delete statement that takes forever to complete, probably blows out the log, and causes bigger problems than those it's trying to solve)
2. Analyzing code to make sure that extra data isn't erroneously being inserted, imported, copied, duplicated, or whatever. This may mean discussing the data in the database with the developers/architects first to figure out where the extra data is.

I've seen plenty of clients have problems from just having too much data that's not needed, and the easiest solution is to get rid of it. Addressing the symptom by doing something like enabling data compression everywhere isn't the answer – you need to address the root cause, just like with any other performance problem.

Call to action: If you're finding that your data volume is just growing and growing and it's starting to cause problems, consider whether all the data is needed or should be there. If you've never thought about these two questions, or checked on your production database, you might be surprised to find problem data that you can get rid of to ease your problems.

I'm curious to hear your thoughts on having too much data, so please feel free to [drop me a line](#), always treated confidentially, of course.

Video Demo

Incremental statistics were a highly-anticipated feature in SQL Server 2014 (ok, maybe only high-anticipated if you love statistics or use partitioned tables!). In this Insider video Erin will take a first look at incremental statistics, show how to enable them for an existing index, and discuss what information you can see about them within SQL Server (spoiler alert: not much).

The video is just over 8 minutes long and you can get it:

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

And you can get the demo code [here](#).

Enjoy!

SQLskills Offerings

We've released most of our classes for 2015. We might add one more domestic delivery of IEPTO1 in the 2nd half of the year and possibly IEPTO2 / IEBI but these will be partially based on instructor and conference schedules; it's important to note that 2nd half courses will be very limited.

Finally, to help your boss understand the importance of focused, technical training, we've added a few new 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

Sydney, NSW, Australia

- February 23-27, 2015: **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2 (formerly IE2)

Chicago, IL

- April 27-29, 2015: **IE0**: Immersion Event for the Accidental/Junior DBA
- April 27-May 1, 2015: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)
- May 4-8, 2015: **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2 (formerly IE2)
- May 4-8, 2015: **IEBI**: Immersion Event on Business Intelligence

- May 11-15, 2015: **IEHADR**: Immersion Event on High Availability and Disaster Recovery (formerly IE3)

Bellevue, WA

- June 8-12, 2015: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)
- June 15-19, 2015: **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2 (formerly IE2)

London, UK

- August 24-28, 2015: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)

Dublin, Ireland

- October 12-16, 2015: **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2 (formerly IE2)

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