

(February 29<sup>th</sup>, 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 heads down working on Pluralsight content and playing with SQL Server 2016 – lots of cool functionality in there! In fact, there are lots of great resources starting to come online; Kimberly did a blog post with a few of them [here](#). Also check out the [series of 2016 blog posts](#) that our friend Bob Dorr has just started on the CSS Escalation Engineers blog where the first post discusses how 2016 runs faster because of scalability improvements they've made!

**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 for your career!). Schedule details are [here](#) and you can read more on the new SSIS classes [here](#).

With a few courses sold out and others filling fast, be sure to book early so that you're not disappointed. Our newly-added **IEPTO1 course in Chicago, running the week of May 16th**, is 2/3 full. Plus we've also added an **IEPTO1 course in Dublin** in October – see [here](#) for details.

Our 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 Steven Erikson's [Gardens of the Moon](#). This is the start of a 10-volume, fantasy epic. I picked up the first few to try out and have been enjoying it since I started the first one on the flights home from diving earlier this month. The story is very involved, with magic, immortals, empires, and lots of intrigue and the book throws you right in from page one. Highly recommended!

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

## Paul's Ponderings

This time I want to talk about log file corruption and what you can do to work around it – in this case, corruption that prevents a log backup running.

Sometimes you might see a log backup fail with an error, but a full backup succeeds, and so does a *DBCC CHECKDB*. What's going on?

Let's say there's a corruption in a portion of the transaction log in a VLF, and that portion of log hasn't been backed up since the last log backup.

A log backup has to back up the transaction log records generated since the last log backup – so it will try to back up the corrupt portion and fail.

A full database backup, on the other hand, only has to back up enough of the log to allow the database to be restored and recovered to a transactionally-consistent point. In other words, it only requires the transaction log back to the beginning of the oldest active transaction at the point that the data-reading section of the full backup completes. So it might not have to back up the corrupted portion of the log and it succeeds.

A *DBCC CHECKDB* operation uses a database snapshot to get a transactionally-consistent view of the database on which to run consistency checks. When the database snapshot is created, crash recovery is run on it to make it transactionally consistent. That requires the same amount of log as if a full backup was taken – back to the beginning of the oldest active transaction at the time the database snapshot is created. So if the corrupt portion of the log isn't needed, the *DBCC CHECKDB* will succeed without any problem

The question then becomes: how to recover from this situation?

Assuming that the database files are intact, there is a simple solution. While this solution isn't recommended under normal working conditions, it works here because the log backup chain is *\*already\** broken. This solution is to throw away part of the log chain and essentially begin a new log backup chain. Here's what to do:

1. Stop all user activity in the database
2. Switch to the SIMPLE recovery model (breaking the log backup chain)
3. Perform a manual checkpoint of the database (which should remove the requirement that the damaged portion of log must be backed up)
4. Switch back to the FULL recovery model
5. Take a full database backup, or a differential database backup (thus starting a new log backup chain)
6. Start taking log backups

You might want to manually shrink and grow the log file in between steps 2 and 3 too – in case the log file is on a damaged portion of disk. You also will need to do some root-cause analysis to determine why the corruption occurred in the first place, and to take preventative measures to stop it happening again.

**Call to action:** If you see weird log corruption messages, try to work around it this way rather than doing something drastic like detaching the database and deleting the log file, or immediately jumping to running *EMERGENCY*-mode repair.

### **Video Demo**

We've been doing a lot with Azure for clients over the last few months, so it's time to start sharing some tips and tricks around Azure. To kick things off, in this week's video, Tim demonstrates how to use the DTU Calculator to determine a good starting point for which pricing tier you will need for your Azure SQL Database.

The video is about 4.5 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!!**
- **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](mailto:Paul@SQLskills.com) and [Kimberly@SQLskills.com](mailto:Kimberly@SQLskills.com)