

(May 11th, 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 Chicago, where Jonathan's in the middle of day 1 of our 5-day IEHADR class, teaching with Kimberly, Glenn, and me. The last two weeks here have been a blur, with 80 students coming through four classes!

We've just added a brand new class to our roster for Chicago in November: **IEPDS: Immersion event on Practical Data Science**, taught by a great friend Rafal Luckawiecki. You can get all the details on the class [here](#). Also, the Bellevue (June), London (August), and Dublin (October) classes are filling up fast so don't delay registering to avoid disappointment.

If you need a warm getaway and an infusion of great technical content (including SQL Server 2016), consider joining us for our popular SQLIntersection and DEVIntersection conferences, in Scottsdale, AZ next week – see [here](#) for details. We have Microsoft SQL Server team members joining us and we have a few special SQL Server 2016 sessions planned. We're also working on our trivia questions for our second SQLafterDark event; we're really looking forward to a fantastic week – interacting with SQL professionals and helping everyone get their problems solved!

As for the content, SQLIntersection focuses on the troubleshooting and performance problems that we know you're facing today but we also highlight best practices in architecture, design, SQL Server 2014, and some insight into SQL Server 2016. For a total of 6 full days – you can immerse yourself into great content with speakers that you know will deliver! Also, use the discount code “SQLskills” and you can save \$50 off registration.

I changed reading genres for the last few weeks and the last two books I've read are contemporary thrillers: Stieg Larsson's [The Girl Who Played with Fire](#), and Gillian Flynn's [Gone Girl](#).

I read Larsson's first book in the series ([The Girl with the Dragon Tattoo](#)) back in 2011 and loved the movie last year (the new one, not the older Swedish one, which I haven't seen but I've been told is excellent). The second book is even better than the first one I think - it turned into a real page turner for me over the last couple of hundred pages. It's hard to talk about the plot without giving twists away but it's a great thriller, with the same characters as the first book, and I highly recommend it.

Gone Girl was even better. I could hardly put it down over the three days I was reading it, and I'm really looking forward to watching the movie. The basic premise is that there's a couple whose marriage is on the rocks, and the wife goes missing on their fifth anniversary. The husband gradually becomes the prime suspect of her murder, and things unravel from there. It's an excellent thriller with some cool plot twists, and again I highly recommend it.

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

Paul's Ponderings

There was a discussion on the distribution list for MCMs last week about index maintenance on databases stored on flash drives, where there's repeated advice from people on the Internet that you can ignore index fragmentation on flash drives.

This is a hot button of mine, and I strongly disagree with this advice. Don't get me wrong – I'm not saying that index fragmentation is always a problem on flash drives, just that you can't blindly ignore it just because other people say so.

Flash drives do not stop index fragmentation happening. Sure, the effect of logical fragmentation is reduced because of the raw speed of the reads, even if the readahead isn't as efficient, but logical fragmentation goes hand-in-hand with low page density (from the page splits).

Flash drives don't stop the effects of low page density - wasted disk space, wasted buffer pool, and extra I/Os. You don't want to have your expensive flash drives having a ton of pages with a large amount of free space on them.

More importantly, and what most people who make that claim completely ignore, is the *cause* of fragmentation: page splits. Flash drives don't stop page splits happening. Page splits cause lots of extra transaction log to be generated - which has to be backed up, restored, log shipped, mirrored, scanned by replication/CDC, and so on. Page splits also cause page latch contention and slow down the operations (inserts and updates) that require the extra space.

Also, before you counter-argue that rebuilds and defrags are logged too, sure they are. But you should only be doing index maintenance on this indexes that need it. Taking index fragmentation into account means you make sure that fragmentation is mitigated or prevented so it doesn't affect your workload performance, not just blindly removing and/or preventing fragmentation just because it's there. That doesn't change with flash drives. By preventing page splits from happening when your workload is running, you can trade off the logging cost of an index reorganize every so often using a fillfactor. There are all kinds of strategies, depending on workload, downtime available, and so on.

Jonathan made the point that whether index fragmentation affects performance on flash arrays and whether to defrag indexes or not on flash storage can easily become a philosophical

argument, but with some basic tests, you can see that the internal fragmentation of indexes significantly affects the ability to do read-ahead reads, and forces more small block I/O against the flash storage. Jon blogged about these tests he performed [here](#).

It may be that for your workload in your environment, having it on flash drives means that index fragmentation really doesn't affect your workload performance at all (taking into account the low page density and extra logging from page splits). That's fantastic. But you need to do some testing to come to that conclusion rather than just accepting random advice that index fragmentation doesn't matter on flash drives. This is yet another case where advice may be based solely on narrow-focused testing. It's rare to see absolutely no effect from fragmentation across all operations, resources, and features.

Flash drives do make I/Os faster. That doesn't mean you should do more of them though. Ultimately, it means that you **cannot** ignore index fragmentation.

Call to action: Very simple – no matter what your I/O subsystem is, you still have to consider index fragmentation and its effects on your workload performance. Don't buy in to random Internet advice that lays down a blanket rule on anything about SQL Server (except regular data or log file shrinking – never do that. ☺)

Video Demo

With the shipment of SQL Server 2014, you can now natively encrypt your backups and still benefit from backup compression. Having this feature built in provides you the ability to more securely store your database backups, as sensitive database backups are often stored on devices that other administrators have access to, so these administrators could easily make copies of the backups and restore them to another server to gain access to privileged information. If you haven't been using backup encryption, now is the time to start considering taking a very important next step in securing your company's data assets.

In this video, Tim will walk you through the steps to create the required key and certificate to start utilizing backup encryption as well as demonstrate the importance of your server certificate and the password you encrypted it with.

The video is 2.5 minutes long and you can get it:

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

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

Enjoy!

SQLskills Offerings

We've released all of our classes for 2015, listed below. We'll release the first portion of our 2016 schedule around September/October.

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

- May 11-15, 2015: **IEHADR**: Immersion Event on High Availability and Disaster Recovery (formerly IE3) *Running now!*
- November 16-18, 2015: **IE0**: Immersion Event for the Accidental/Junior DBA
- November 16-19, 2015: **IEPDS**: Immersion Event on Practical Data Science
- November 16-20, 2015: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)

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.

Spring SQLintersection

We're really looking forward to our Spring SQLintersection NEXT week in Scottsdale, AZ. There's still time to join us! See [here](#) for details. Don't forget to use the discount code

“SQLskills” (without the quotes and it isn’t case-sensitive) and you can save \$50 off registration. We hope to see you there!

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