

(October 12th, 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 Dublin, where Erin, Jon, Kimberly, and I are in the middle of teaching the first day of our IEPTO2 class on performance tuning. Ireland is such a beautiful place and we took in a lot of the sights over the weekend while acclimatizing to the time zone.

We've just announced our 2016 classes in Chicago (April/May), London (June), and Bellevue (September), including **two brand new classes on SSIS**, written and presented by our new and exciting training partner *Linchpin People*. We have extremely high standards for our Immersion Event instructors and we're really excited about these new classes that we're adding to our incredible Immersion Event lineup.

You can save up to \$200 off all our 2016 classes by registering before the end of the year and as an alumnus, you can always get the lowest (frequent-flyer) discount! Schedule details are [here](#) and you can read more on the new SSIS classes [here](#).

You've only got two weeks left to register for SQLintersection at the end of October, and there are just a couple of seats remaining in our IE0 and IEPDS Immersion Events in Chicago in November. Check out more information online at www.SQLintersection.com and don't forget to register with the discount code of SQLskills to save \$50 off your registration.

While we were diving last month I read a bunch of books and I want to tell you about three of them: two travel books by William Dalrymple – [The Age of Kali: Indian Travels and Encounters](#) and [In Xanadu](#), and one novel by Helene Wecker – [The Golem and the Jinni](#).

The Age of Kali is a collection of Dalrymple's travel essays covering India, Sri Lanka, and Pakistan, casting a penetrating look at the problems affecting those societies in the late 1990s. *In Xanadu* is a travelogue following Marco Polo's journey along the Silk Road to Xanadu. They travel through Israel, Syria, Turkey, Iran, Pakistan, and China in the late '80s, with all kinds of interesting encounters along the way. Both books are excellent and highly recommended!

Wecker's book is an excellent debut novel set in early 1900s New York, following the story of a golem (a creature made from clay and brought to life with Kabbalistic magic) and a jinni (a natural, elemental creature made of fire) that was trapped in a copper flask by a wizard a thousand years ago. It covers their problems integrating into the populace of New York, their

eventual meeting, and problems when their true nature starts to be discovered. Very well written and highly engaging – definitely recommended!

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

Paul's Ponderings

When people design indexes, they primarily consider how useful the index will be for the query workload. Secondly, they're likely to consider maintenance requirements and potential problems around index fragmentation that may occur. However, one facet of index design that's usually neglected is considering the *fanout* or *fan-out* of the index structures themselves.

The fanout of an index measures, for a page at level x in an index, how many pages it references in the level below (nearer the leaf level). The higher the fanout is, the fewer the number of levels in the index.

The index key size impacts the size of the structure needed to reference it. Specifically, the index key is pushed up to all entries (and all levels) in the index as it's used to allow navigation through the index from the root page down to the leaf level.

The larger the index key size, the fewer index records can be stored in an index page and so the lower the fanout. The lower the fanout is the more levels are required in the index, depending on the number of pages at the leaf level.

For instance, if the fanout is 10 in an index, that means each index page can hold 10 index records, referencing 10 pages at the level below in the index. If the index has 10,000 pages at the leaf level, there needs to be 1,000 pages in the level above, then 100 pages, then 10 pages, and finally the root page. That's a total of 5 levels.

For the same data, if the index fanout is changed to 100, and the index has 10,000 pages at the leaf level, the next level needs 100 pages, and then there's the root page. That's a total of only three levels.

(I'll do a more complete example on my blog this week or next.)

So why do you care about the number of levels in the index, determined by the fanout?

Well, when traversing an index, at each level in the Storage Engine needs to decide how to navigate further down the index. It does this using the records in the page, to figure out where the key value is that it's searching for.

The more levels there are in the index, the more CPU each search through the index takes, and so the slower each search is. Even though the index itself may be wholly contained in memory, this operation per-level in the index requires performing a search (it's an efficient binary search, but it's still a search that uses CPU resources).

So the reason you need to care about fanout is that low fanout increases CPU usage during searches.

And fanout is controlled by index key size, so careful index key choice entails not just considering fragmentation effects, but also minimizing the key size as much as possible to make the highest possible fanout.

Call to Action

When you're evaluating indexes for efficiency, make sure you consider the index key size and the fanout, with a view to minimizing the number of levels in the index. You can see the number of levels in your indexes using the *sys.dm_db_index_physical_stats* DMV, and also looking at the *IndexDepth* value in the *INDEXPROPERTY* function.

As I mentioned above, I'll do a much more detailed example, with code and workload timings on my blog sometime in the next week or two, but this at least will get you thinking in the meantime.

Additionally, if you want to get more details on *why* physical database design matters – check out [Kimberly's Pluralsight course](#) on exactly this topic, and there's a demo from it below.

Video Demo

For this newsletter I've pulled together two demos from Pluralsight courses.

The first demo from my course on [SQL Server: Index Fragmentation Internals, Analysis, and Solutions](#). The demo focuses on the low-priority lock waits that were introduced in SQL Server 2014 to allow easy control of what happens when an online index operation gets blocked or it block other queries. The video is just under 4 minutes long and you can get it in MP4 format [here](#) with the demo code available [here](#).

The second demo is from Kimberly's course [SQL Server: Why Physical Database Design Matters](#). This demo focuses on the impact of the clustering key on wasted index structures showing both the impact to disk space as well as performance. The video is just over 4 minutes long and you can get it in WMV format [here](#) with the demo code available [here](#).

Enjoy!

SQLskills Offerings

We've released the majority of our 2016 classes for registration, listed below. It's possible that we might add one or two classes in Chicago in the Fall, but that will depend on the Fall conference schedule as well as demand. **Note that you can save up to \$200 by registering for these classes before the end of 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

Chicago, IL, 2015

- **IE0:** Immersion Event for the Accidental/Junior DBA
 - November 16-18 (**2 seats remaining!**)
- **IEPDS:** Immersion Event on Practical Data Science
 - November 16-19 (**2 seats remaining!**)
- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
 - November 16-20 (**SOLD OUT!**)

Chicago, IL, 2016

- **IE0:** Immersion Event for Junior/Accidental DBAs
 - April 18-20 (**US\$120 discount for registering in 2015**)
- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
 - April 18-22 (**US\$200 discount for registering in 2015**)
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
 - April 25-29 (**US\$200 discount for registering in 2015**)
- **IEBI:** Immersion Event on Business Intelligence
 - April 25-29 (**US\$200 discount for registering in 2015**)
- **IEHADR:** Immersion Event on High Availability and Disaster Recovery
 - May 2-6 (**US\$200 discount for registering in 2015**)
- ****NEW** IESSIS1:** Immersion Event on Learning SQL Server Integration Services
 - May 2-6 (**US\$200 discount for registering in 2015**)

London, UK, 2016

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
 - June 13-17 (**US\$200 discount for registering in 2015**)
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
 - June 20-24 (**US\$200 discount for registering in 2015**)

Bellevue, WA, 2016

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

- September 12-16 (**US\$200 discount for registering in 2015**)
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
 - September 19-23 (**US\$200 discount for registering in 2015**)
- ****NEW** IESSIS2:** Immersion Event on Advanced SQL Server Integration Services
 - September 19-22 (**US\$160 discount for registering in 2015**)

See [here](#) for the main Immersion Event Calendar page that allows you to drill through to each class for more details and registration links.

SQLintersection

Our final schedule and line-up has been posted and we're looking forward to another fantastic event! Our Fall SQLintersection is running the week of October 26, 2015 and we're back in Las Vegas. Yes, it's the same week as PASS. No, that's not really a problem! There aren't a lot of folks that attend both PASS and SQLintersection so the overlapping week isn't really a scheduling issue. And, it's a great way to get different perspectives on the tech and have everyone come back and do knowledge-transfer.

Check it out online at www.SQLintersection.com. When you register, 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