

(April 28<sup>th</sup>, 2014)

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 as it hits your inbox, I'm teaching the first module of our IE1 course on performance tuning through design and indexing. We're here for four weeks altogether, with 8 classes running over the month.

We've also just wrapped up SQLIntersection in Orlando, FL and we received rave reviews. Folks loved the ability to intersect with speakers and get their questions answered. And, as for the content, it was right on!

*"I have to say that this conference is the best I have ever been to; I really enjoyed it and met some new friends. If you want to learn from industry experts and you enjoy discussion, go to SQLIntersection. The SQL community is the only one I know that the experts are so willing to share their knowledge and help you become a better DBA. That is what the SQLIntersection is all about. I cannot wait until the November conference."*

Jim Pecora, NCH Marketing Services, Inc.

We're already planning our fall lineup for SQLIntersection in Las Vegas at the MGM Grand so SAVE THE DATE for November 11-12-13, 2014. Workshop, session, and keynote details will be posted shortly.

The most recent book I've read is Sarah Key Penman's [The Sunne in Splendour](#). Penman's novels are excellent historical fiction, meticulously based on actual history. In this book she follows the life of Richard III, steeped in betrayal and treason against both himself and his brother, the earlier Edward IV (with Edward V only reigning for a few months until pronounced illegitimate). In more than 900 pages of small-font, the book is huge and a long, engrossing read. Penman puts the death of the infamous 'princes in the Tower' at the hands of the Duke of Buckingham, which makes a lot of sense given that neither Richard nor Henry Tudor (later Henry VII) had both motivation \*and\* opportunity to do the deed. Altogether a splendid book, and I strongly recommend it for all fans of history.

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 discussions I've had over the last two weeks was around why SQL Server doesn't have a way to automatically remove index fragmentation. I've had similar discussions over the years, both inside and outside Microsoft and the SQL team, on SQL Server's lack of automatic facilities to do things that involve a bunch of resource consumption (e.g. index fragmentation removal, consistency checking, index creation/changes, de/normalization).

While I was on the SQL team, some of these were seriously considered, but I was always very cautious about them. Here's a (non-exhaustive) list of questions to consider, using index fragmentation removal as an example:

- What time of day/week should the process run?
- What thresholds to use for 'do nothing', 'reorganize', 'rebuild'?
- What if there's no disk space to do a rebuild? How to check that? What if a rebuild fails because of lack of space: should a reorg be done instead?
- What if the database has to autogrow to accommodate the new index? Should some kind of shrink be run afterward?
- How much parallelism should an index rebuild use?
- Who manages the transaction log size if the table is really large? Should we automatically take log backups? Where should they be backed up to and how should they be named? What if there's no space for the log backups? What about space on the log shipping secondaries? And the time required to replay the log on the log-shipping secondaries?
- What about the impact of extra transaction log on database mirroring or availability groups? Both in terms of the huge SEND queue preventing log truncation, and the huge REDO queue preventing the mirror/AG replica coming online quickly?
- What if a cluster failover occurs during the operation and the rollback prevents the database coming online within the company's RTO?
- What if there's no Enterprise Edition? Then there's no online index rebuild and so the application/users are blocked while the table is SCH\_M locked while the offline index rebuild occurs. Does that imply always do a reorg in Standard Edition?

My point is that it's a lot harder than you think to just put some automatic behavior into SQL Server. Given the capacity for resource usage, if it's wrong once and causes a performance problem, or excessively large backups, or something else negative, it will get a bad reputation and people won't use it.

**Call to action:** It's very tempting to blame Microsoft for not having an automatic method of doing X or Y. However, when you stop to consider the complexity of the problem and the sheer diversity of environments that the automatic feature would have to work in, it should come as no surprise that Microsoft tends to \*not\* invest engineering time in such features. Very often there is a third-party solution to the problem you'd like to see solved – for instance, [Ola Hallengren's wonderful database maintenance solution](#) that you can customize to achieve much of what I've

considered above. Use a little bit of ingenuity and you can solve the problem yourself, for your environment, and in a way that best suits your application.

I'm curious to hear your thoughts on automatic features in SQL Server, so please feel free to [drop me a line](#), always treated confidentially, of course.

### **Video Demo**

From Erin: When you add ORDER BY to a query, what happens to a query plan? You probably expect to see a sort operator in the plan, but sometimes you may not. In this video I'll review the basics of the sort operator, discuss why it appears in the plan, and look at a query plan warning that's worth further investigation.

The video is just over twelve 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**

Please know that all of our classes will run and their dates will not change. Additionally, most of our public training courses will be held in the first half of this year. We will add a couple of other classes in the second half of the year, but not all that many (specifically, one IE1 and one IE2 in the US and two IE1 deliveries in Australia)). Please plan accordingly. And, if you're in Australia and want direct notice when our classes are added, please email [training@SQLskills.com](mailto:training@SQLskills.com).

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](#)

### **2014 Immersion Events**

Chicago, IL

- May 5-6, 2014: **IEHW**: Immersion Event on SQL Server Hardware
- May 5-9, 2014: **IE2**: Immersion Event on Performance Tuning (**SOLD OUT!**)
- May 12-16, 2014: **IE3**: Immersion Event on High Availability and Disaster Recovery

- May 12-16, 2014: **IEDEV**: Immersion Event for Developers
- May 19-23, 2014: **IE4**: Immersion Event on Security, PowerShell, and Developer Support
- May 19-21, 2014: **IE0**: Immersion Event for the Accidental/Junior DBA

Bellevue, WA

- June 9-13, 2014: **IE1**: Immersion Event on Internals and Performance
- June 16-20, 2014: **IE2**: Immersion Event on Performance Tuning

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

### **Fall SQLintersection**

As mentioned above, we're already planning for our Fall SQLintersection event. This year our Fall SQLintersection conference will be the week of November 10<sup>th</sup> in Las Vegas. We'll have a lot more details in a month or so when we've worked out the speaker lineup, workshops, and sessions. 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](mailto:Paul@SQLskills.com) and [Kimberly@SQLskills.com](mailto:Kimberly@SQLskills.com)