

(September 16th, 2013)

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 Bellevue, WA where we've just started the first day of our Immersion Event on Internals and Performance (IE1). Over the last few weeks Kimberly and I have been recording new courses for Pluralsight and they should be live soon – watch my blog for announcements!

We've just released 17 new 2014 classes for registration in Tampa, Chicago, and Bellevue, plus we're planning classes in Australia in March. All classes have alumnus discounts and many have buy-1-get-1-free offers for the first five registrants in 2013 – see my blog post [HERE](#) for details!

The most recent book I've read is Neil Gaiman's [American Gods](#). I've been meaning to read this for several years since I bought it, and finally needed a break from the heavy-duty classics I've been reading this summer so I picked it up. I've read all of Gaiman's graphic novels, as I especially enjoyed the fabulous [Sandman series](#), but never any of his prose. And I'm pleased I did! The story is about old gods that were brought to America through being worshipped by the people that came here. They're forgotten and un-worshipped now, living parodies of human lives, but now there's going to be a war between them and the new gods that people worship like Media, Technology, and so on. Great story, quick read, and 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

There are two things I'd like to talk about in this newsletter: staggered index maintenance, and more on Microsoft's sudden cancellation of the entire MCM/MCSM program.

Staggered Index Maintenance

Last week I did a remote user group presentation for the Wisconsin SQL Server User Group on index fragmentation – a preview of one of my PASS Summit sessions.

One of the things I mentioned briefly during the talk was how people who are using synchronous database mirroring or synchronous availability group replicas sometimes give up on index maintenance. This happens because both these technologies require the full recovery model, and index rebuilds in the full recovery model generate a *lot* of transaction log in a single transaction (sometimes an amount equal to the size of the index being rebuilt).

If someone's been used to using the bulk_logged recovery model to reduce the amount of log generated during index rebuilds, moving to the full recovery model can make the extra log generation prohibitive. There's also the fact that in synchronous mode for both these technologies, all the generated log for a transaction must be sent to the remote server before the transaction can commit – again, likely to be prohibitive for rebuilding large indexes.

There are two solutions, which can be combined.

The first is to make judicious use of index fillfactors to reduce the amount of page splits that occur, and hence the amount of index fragmentation that occurs. This helps reduce the amount of log being generated (as page splits are always fully logged), reduces low page density that can lead to wasted buffer pool space and extra I/Os, and reduce the need for index maintenance. Kimberly discusses fillfactors in more detail in [this blog post](#).

The second is to switch from using index rebuilds (or not removing problematic fragmentation at all) to using index reorganize through the *ALTER INDEX ... REORGANIZE* command (that is essentially the *DBCC INDEXDEFRAG* I wrote for SQL Server 2000). An index reorganize only takes care of fragmentation that exists and can be interrupted without loss of work – which makes it perfect for a staggered index maintenance strategy.

Consider a 200GB clustered index. If rebuilding is out of the question for you, create a maintenance routine that runs reorganize against the index for, say, an hour each day at the time with the lowest level of workload activity. After an hour, kill the reorganize. Each day the reorganize should quickly get back to where it left off the night before and do a bit more defragging. Rinse and repeat, over and over to keep fragmentation under control with the minimum of hassle and log generation.

From Kimberly: I would also add that schema and/or index changes might significantly reduce the need for maintenance. Partitioning the data into separate tables (possibly by year) where only the current year needs regular maintenance and the older years are read only (data dependent) might give you both maintenance gains as well as performance (better, more granular statistics), etc. And, looking into the indexes themselves might be beneficial. There are some cases where just changing the definition of the clustering key might significantly reduce or eliminate some maintenance. But, this is a more challenging discussion. And, any schema changes might be very difficult to make (all of the related code would need to [at a minimum] be tested and in the worst case, changed).

MCM Program

Further to everything I said in the last newsletter about the cancellation and my feelings on it, Microsoft Learning (MSL) organized a conference call with the MCM community this week and attempted to explain a bit more about the decision. Basically the reasons for canceling the program boil down to:

- The various Product Groups in Microsoft don't provide enough funding for the MCM program.
- The community of MCMs isn't as large as Microsoft wants it to be.
- They're canceling the MCM program so they can use the resources to figure out what to do next.

And there you have it. And they said that the program is really dead, there's no possibility of a reprieve.

One good thing I've heard is that the remaining time to attempt the MCM exams has been extended until December 31, 2013 – which is great news for those who need to complete the certification.

I published the [results of a survey](#) last week that gave some idea of how people have been affected by the cancellation, and Joe wrote an [excellent blog post](#) about the conference call with MSL.

My final words on the subject are those from the end of my blog post: Don't give up learning – it's always worth it.

Call to action: Again, two calls to action:

- If you're using synchronous database mirroring or synchronous availability group replicas and you've had to vastly reduce the amount of fragmentation removal you can perform, consider moving to a staggered index maintenance scheme as I describe above. Also don't forget to make use of fillfactors to reduce the page splits that occur and the fragmentation they produce.
- If you'd abandoned hopes of (re)taking the MCM exams because of the imminent closure on October 1st, reconsider now that you have until December 31st. Better to prove to yourself that you've got what it takes than to give up after all the work you've put in to get this far.

I'm curious to hear your thoughts on either of these topics, so please feel free to:

- [Drop me a line](#) about staggered index maintenance
- [Drop me a line](#) about the MCM program cancellation

Always treated confidentially, of course.

Video Demo

From Joe:

In this video, I demonstrate how sometimes you can miss the top query resource consumers on a SQL Server instance if you're not careful, and how to use query hashes with the DMV `sys.dm_exec_query_stats`.

The video is three and a half minutes long and you can get it:

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

The demo code is available [here](#).

Enjoy!

SQLskills Offerings

As I mentioned above, we've just opened up our first-half of 2014 classes for registration. All classes have alumnus discounts equivalent to 25%-off the full price, and many have a special offer for the first 5 people who register for each class in 2013 – buy-1-get-1-free!

Please know that these classes are final as the hotel contracts are signed, and the classes will not be cancelled or moved for any reason, nor will the dates change.

Also, we've added a few new items to help you justify your training dollars:

- [Letter to your boss explaining why SQLskills training is worthwhile](#)
- [Community blog posts about our classes](#)
- [Immersion Event FAQ](#)

Remaining 2013 classes:

- September 23-27, 2013: Performance Tuning (**IE2**) in Bellevue, WA
- September 30-October 2, 2013: Accidental DBA (**IE0**) in Bellevue, WA
 - See [here](#) for special pricing offers!
- September 30-October 3, 2013: Advanced T-SQL (**IETS**) in Bellevue, WA
- November 11-15, 2013: Internals and Performance (**IE1**) in Chicago, IL

New 2014 classes:

Tampa, FL

- February 3-7, 2014: **IE1**: Immersion Event on Internals and Performance
- February 3-5, 2014: **IE0**: Immersion Event for the Accidental/Junior DBA
- February 6-7, 2014: **IEHW**: Immersion Event on SQL Server Hardware
- February 10-14, 2014: **IE2**: Immersion Event on Performance Tuning
- February 10-14, 2014: **IEDEV**: Immersion Event for Developers
- February 17-21, 2014: **IE3**: Immersion Event on High Availability and Disaster Recovery
- February 17-21, 2014: **IEBI**: Immersion Event on Business Intelligence

Chicago, IL

- April 28 – May 2, 2014: **IE1**: Immersion Event on Internals and Performance
- May 1-2, 2014: **IEHW**: Immersion Event on SQL Server Hardware
- May 5-9, 2014: **IE2**: Immersion Event on Performance Tuning
- May 5-9, 2014: **IEBI**: Immersion Event on Business Intelligence
- May 12-16, 2014: **IE3**: Immersion Event on High Availability and Disaster Recovery
- May 13-16, 2014: **IETS**: Immersion Event on Advanced Transact-SQL
- 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.

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