(October 14th, 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 the 2013 PASS Summit in Charlotte, North Carolina where, as this email is being sent to you Monday morning, I'm on stage presenting my pre-con workshop *Practical Disaster Recovery Techniques* and Glenn's presenting his workshop *Scaling SQL Server 2012*. The whole SQLskills team is at the conference this week and we're looking forward to meeting as many of the SQL Server community as we can.

Our two-for-one offer is starting to really take off. As an example, for two people to attend IEDev or IEBI in 2014 would normally be $6,590 plus expenses. Now you can have two people attend for $3,295, and the savings should cover all travel expenses from virtually anywhere in the US. The two-for-one offer is available for the first 5 people to register for any 2014 delivery of IE0, IEBI, IEDEV, IE4, IETS (**NOTE:** about **half** the free seats have gone already).

Kimberly's new Pluralsight course SQL Server: Why Physical Database Design Matters is proving to be very popular and generating lots of discussion. I finished recording my SQL Server: Advanced Corruption Recovery Techniques course just last week and that should be published soon, as well as Jonathan's SQL Server: Maintenance Plans course.

The most recent books I've read are Patrick Leigh Fermor's *A Time of Gifts* and *Between the Woos and the Water*. In 1933 Fermor, then 18 years old, set off from the Hook of Holland to walk across Europe to Constantinople (now Istanbul). The first book covers the journey to the Danube River, and the second covers the journey down the Danube as far as the Iron Gates gorge on the border of Romania. It's an incredibly rich narrative of the people and places in pre-war Europe that Fermor didn't write until the 1970's, and one of my favorite parts is travelling through Hungary and Budapest, which I did myself in late 1998, where the snow-covered countryside and villages still seemed ages old. I strongly recommend these books for anyone who enjoys travelogues. Oh, and part three will be released posthumously in March 2014.

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

Last week I did a remote user group presentation for the Southampton (UK) SSUG and gave them a preview of my index fragmentation deck that I'll be presenting on Friday at PASS. In that deck I discuss the differences between *ALTER INDEX ... REORGANIZE* and *ALTER INDEX ... REBUILD*, and then present a consolidated list of comparison points. I find that many people don't know about these differences so in this newsletter I'm going to spell them out for you.

Space Required

Rebuilding an index requires building a new index before dropping the old index, regardless of the amount of fragmentation present in the old index.

Reorganizing an index first squishes the index rows together to try to deallocate some index pages, and then shuffles the pages in place to make their physical (allocation) order the same as the logical (key) order. This only requires a single 8-KB page, as a temporary storage for pages being moved around.

If you have space constraints, and can't make use of single-partition rebuild (offline until SQL Server 2014), reorganizing is the way to go.

Algorithm Speed

An index rebuild will always build a new index, even if there's no fragmentation. The length of time the rebuild takes is related to the size of the index, not the amount of fragmentation in it.

Reorganizing an index only deals with the fragmentation that exists, and the more fragmentation there is, the longer a reorganize will take.

This means that for a lightly fragmented index (e.g. less than 30% fragmentation), it's generally faster to reorganize the index, but for a more heavily fragmented index, it's generally faster to just rebuild the index.

Transaction Log Generated

In the FULL recovery mode, the index rebuild is fully logged, so the transaction log will have to accommodate the full size of the index in a single transaction. This also means the entire generated transaction log will need to be mirrored, sent to your AG replicas, and scanned by replication, and so on.

In the SIMPLE and BULK_LOGGED recovery modes, the amount of transaction log generated by an index rebuild will be minimal – just the allocations of pages and extents. However, the next log backup performed (either in BULK_LOGGED or after switching to FULL) will contain all the extents changed by the rebuild, and so will be roughly the same size as if the rebuild was done in the FULL recovery mode. The benefits are in time and the fact that the transaction log itself does not have to accommodate the full size of the index during the rebuild and in a single transaction.

In all recovery modes, reorganizing an index is fully logged, but is performed as a series of small transactions so should not cause the transaction log to grow inordinately. And of course, transaction log is only generated for the operations performed, which may be less than for a rebuild.

Locks Required

An offline index rebuild of a clustered index holds an exclusive table lock – no updates or reads. An offline index rebuild of a nonclustered index holds a shared table lock – no updates. An online index rebuild of any index acquire a short-term shared table lock at the start of the operation, holds an intent-shared table lock throughout the operation (which will only block exclusive table locks), and then acquires a short-term schema-modification (i.e. super-exclusive) table lock at the end of the operation. 'Online' is a bit of a misnomer.

An index reorganize holds an intent-exclusive table lock throughout the operation, which will only block shared and exclusive table locks. One of the major reasons I wrote *DBCC INDEXDEFRAG* for SQL Server 2000 was as an online alternative to *DBCC DBREINDEX*.

Interruptible or Not

An index rebuild operation cannot be interrupted – it's atomic.

An index reorganize can be interrupted and the worst that will happen is that a single page move operation is rolled back. This is why staggered index maintenance (like I described a month ago here) is possible.

Progress Reporting or Not

Index rebuilds do not have proper progress reporting. You can hack it for online index operations by looking at the *bigintdata1* column in the *Progress Report: Online Index* Profiler event, which happens to show how many rows of the old index have been scanned. You can also hack it for offline index operations by looking at the number of reads the SPID has done.

Index reorganize operations populate the *percent_complete* column of *sys.dm_exec_requests* so you can easily gauge how much work remains. In fact *DBCC INDEXDEFRAG* also used to do progress reporting, but less elegantly, by printing a progress message to your connection every 30 seconds.

Statistics

An index rebuild will always rebuild the index column statistics with the equivalent of a full scan.

An index reorganize does not see a total view of the index and so cannot update statistics, meaning that manual index statistics maintenance is required.

Summary

As you can see, there are quite a few major differences between rebuilding and reorganizing, but there's no right answer as to which one you should use – that's your choice.

Call to action: If you have an index maintenance routine that always rebuilds and never considers reorganizing, I urge you to reconsider. It's usually better to reorganize a lightly fragmented index and rebuild a more heavily fragmented index – to save time and resources. You'll find that most index maintenance products and freely-available scripts allow you to make that choice.

I'm curious to hear your thoughts on whether and how you choose between rebuilding and reorganizing, so please feel free to drop me a line, always treated confidentially, of course.

Video Demo

Just like in the last newsletter, we've all been busy preparing for the PASS Summit so this time I'd like to give you a video from Kimberly's newly-released Pluralsight course. In this demo, Kimberly explains and demonstrates how to recognize data type inconsistencies in query plans that are causing implicit conversions and affecting performance.

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

- In WMV 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 class:

- November 11-15, 2013: Internals and Performance (**IE1**) in Chicago **(SOLD OUT!)**

2014 Immersion Events:

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: **IEBI**: Immersion Event on Business Intelligence
- February 17-21, 2014: **IE3**: Immersion Event on High Availability and Disaster Recovery
- February 17-21, 2014: **IEDEV**: Immersion Event for Developers

Australia

- Coming soon: Sydney, NSW, Australia – March 10-14, 2014
- Coming soon: Melbourne, VIC, Australia – March 17-21, 2014

Chicago, IL

- April 28 – May 2, 2014: **IE1**: Immersion Event on Internals and Performance
- April 28 – May 2, 2014: **IEBI**: Immersion Event on Business Intelligence
- May 5-6, 2014: **IEHW**: Immersion Event on SQL Server Hardware
- May 5-9, 2014: **IE2**: Immersion Event on Performance Tuning
- 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