(August 20th, 2012)

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 bi-weekly Quick Tips is coming to you from Redmond, where we're just wrapping up IE2, our Immersion Event on Performance Tuning, before going straight into IE3 (High Availability & Disaster Recovery) next week. We've got the entire SQLskills team in town this week and we've have a great time hanging out together!

The most recent book I've read is Hilary Mantel's *Bring Up The Bodies***,** the sequel to her excellent *Wolf Hall*. The book covers the proceedings leading up to the downfall of Anne Boleyn and the Boleyn family, as told through the eyes and thoughts of Thomas Cromwell. I recommend reading Wolf Hall first, as that covers the downfall of Cardinal Wolsey, instigated by Anne Boleyn. Both books are really excellent and I strongly recommend them for history fans.

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

This week's class is all about performance tuning. Basically, we discuss the common question: what's eating all of the resources on my SQL Server and preventing my workload from running as fast as I want it to?

And, as a part of our discussions we tackle a commonly overlooked cause of poor performance which is having duplicate nonclustered indexes. SQL Server will let you create as many nonclustered indexes as you want (ok, technically 999), that are exactly the same. *Exactly* the same, down to the INCLUDEd columns and key order. Why? The reason is backward compatibility and so that they don't break horribly-written applications that use hard-coded index hints. In actuality, there's absolutely NO good reason to allow them and often people end up with duplicates that don't even look similar (because of how SQL Server internally changes the indexes based on the clustering key). The end result, you might have duplicates and not know it!

You might think that duplicate nonclustered indexes won't be used by SQL Server but they ALL can be (and not at the same intervals). They are not just wasted space.

Every INSERT, UPDATE, and DELETE operation (which I'll just call 'DML' from now on) will affect every nonclustered index (except possibly filtered indexes), as each nonclustered

index has to have a row for each row in the table. This means that each DML operation generates transaction log for the commensurate changes to each nonclustered index, and the index pages for each index have to be in the buffer pool.

So each duplicate index is taking up space in the data files, taking up space in your precious buffer pool, taking up space in the transaction log, needing locks, needing latches, taking up processing time, and so on.

And then what about index fragmentation? Each of these duplicate nonclustered indexes will be getting fragmented, so they'll get picked up by your index maintenance. And your statistics maintenance. And will have to be checked by DBCC CHECKDB.

You get the idea. Duplicate nonclustered indexes are really bad. And, all you have to do is just remove them and you'll reduce your overhead and increase performance.

That's the catch though – you have to know internals and how SQL Server stores data in the leaf and non-leaf levels of an index. Luckily Kimberly worked out how to find them and she wrote two posts:

- How can you tell if an index is REALLY a duplicate?
- Removing duplicate indexes

The second post has all the code to use to figure out which indexes are duplicates on 2005 onwards. One of our class attendees from Tampa earlier this year also ported the code to SQL Server 2000 – see here.

Call to action: Read the first post and then download the code from the second post and check out some of your databases to see whether you have any duplicate indexes. You'll be happy if you find them as every one of them is a leech sucking resources out of your SQL Server. One DBA found 77 duplicates in one of their databases just last week!

I'm curious to hear your thoughts about duplicate indexes, especially how many you find, so please feel free to drop me a line, treated confidentially of course.

Video Demo

As we're all in class this week, we didn't record a new video for you. Instead, I'd like to point you at a demo I recorded for the MCM training videos we produced for Microsoft back in 2010. This one explores the RECOVERY, NORECOVERY, and STANDBY options of restore, combined with restoring to a point in time. You can get the video in a variety of formats here.

Enjoy!

SQLskills Offerings

We've now got two more courses online with Pluralsight, one by Joe (SQL Server: Transact-SQL Basic Data Retrieval) and one by Jonathan (SQL Server: Collecting and Analyzing Trace Data) – check them out! That takes our online course total to four, with 5 more coming online by the end of September. I'll be blogging more about our plans with Pluralsight in the next week or two.

And, finally, we're planning to release our complete 2013 class calendar in a week or two – watch this space!

As a note for those of you who really want us to come to *your* city… we'd LOVE to. However, we just can't be everywhere (we wish we could!). Having said that though, these events are truly ***immersion*** events and if there isn't one in your city, it might be a good thing to consider the travel and time away. Being in a hotel for the event allows you to participate in evening events, network with your colleagues and focus on the class without the usual requirements of being in your hometown.

So, that's it for now. We hope to see you soon!

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