(October 3rd, 2011)

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 home – Redmond, WA – where we have some time with local customers (as well as presenting at PASS in Seattle) between 3-week trips around the US to clients and conferences (including our Chicago Immersion Event and SQL Connections).

Today we have a new addition to the SQLskills team – Joseph Sack. He's a long-time SQL expert who joins us from Microsoft where he used to be a Premier Field Engineer and latterly managed the Microsoft Certified Master program (he's also an MCM for 2005 and 2008). You can read my announcement here. This is really exciting for us as Joe brings an amazing breadth and depth to our team and years of consulting experience.

The most recent book I've read is Stacy Schiff's *Cleopatra: A Life*. I haven't read much about Egyptian history and so this caught my eye last time I was browsing in SeaTac airport. Schiff does an excellent job of describing the travails of Cleopatra's life and her relationships with Julius Caesar and Mark Anthony during the 1st century BC.

Check out the final part of the newsletter for details of our **new remote DBA service**, plus our projected class schedule through August 2012.

Please let us know if you liked what you read/saw here and/or have any suggestions for future Quick Tips.

Paul's Ponderings

Over the last three weeks Kimberly and I have been on-site with customers in NYC and Omaha (where we presented at their local user group) and each week I was asked the same question: why doesn't SQL Server take care of fragmentation automatically?

There are two choices for automatically correcting fragmentation: at the time it occurs and in the background. Real-time fragmentation removal would slow operations down too much so isn't even worth considering – I don't know of any RDBMS or file system (having worked extensively on both technologies) that attempts to do it.

So that leaves the possibility of a background process. Here are a few gnarly questions that would need to be answered (or have configuration options) to make this work:

- What fragmentation threshold should it use to trigger action?

- How to automatically determine the fragmentation level of an index without causing a bunch of I/O overhead? Theoretically this could be done by the range-scan readahead mechanism.
- How to determine which indexes to act on, even for determining fragmentation?
- How often should the background task fire up to do some work? How long should it run for?
- Which algorithm should it use – reorganize (slower, less space, and logging) or rebuild (faster, lots of space and logging)?
- What about automatically raising or lowering index fill factors? How to take into account the amount of free space available in the database?
- What kind of I/O, CPU, and memory limit should be placed on the background task?
- Who manages the transaction log size if the index is really large? Should it automatically take log backups? What if there's no space for the log backups? What about space on the log shipping secondaries? And the time required to roll forward the log on the log-shipping secondaries?
- What about the impact of extra transaction log on database mirroring or Denali Availability Groups? Both in terms of the huge SEND queue preventing log truncation, and the huge REDO queue preventing the mirror coming online quickly?
- What about the impact of extra transaction log on the transactional replication log reader Agent job? Should there be an automated tracer-token mechanism to throttle back the defrag task if replication latency gets too high?
- What about the locking required to do the defrag? What if it causes blocking?

And then imagine if you had auto-shrink and auto-defrag enabled at the same time? It would be the perfect storm of activity on the server for absolutely no gain! ☺

Summary: I highly doubt we'll ever see things like automated fragmentation removal, automated missing-index creation, or automated unused index removal. From being involved in several design explorations back at Microsoft, it's a lot harder than you may think to put automated behavior into a general purpose RDBMS.

Call to action: anything that you think should be automated is likely to be something you should be doing regularly on your servers. Make sure you've got the right Agent jobs and alerts set up so your servers run smoothly.

I'm really interested to know your thoughts on automated features—feel free to drop me a line, confidentially as always.

Video Demo

Several times over the last few weeks I've seen people being confused about the number of log records that will be generated by fully-logged bulk operations so I thought I'd do a quick demo to show you how things work and that expected behavior is not always what you get. The video is about 10 minutes long.

I produced the video in WMV and MOV formats so everyone can watch. You can get the videos:

- For WMV: here
- For MOV: here

I recommend downloading before watching. And you can get the demo code here.

SQLskills Offerings

We've just launched a new "remote DBA" service where we provide repeated mini-health checks on regular intervals (once a server has been initially health-checked) plus automated monitoring of SQL Server for problems – and what to do when a problem inevitably *does* arise. This is an excellent way to gain access to our team's incomparable expertise and experience – especially if you don't have a full-time DBA on staff. However, even if you do, this gives you an additional set of expert eyes to watch over your critical data.

Let me know if you're interested in getting more details.

Registrations are continuing for our remaining Immersion Events this year – a *new* Developer Immersion in Chicago, October 24-28, and Internals and Performance (IE1) plus a *new* BI Immersion in Atlanta, December 5-9. **The IE2 offering in Chicago in October is now sold out!**

For everyone that's been asking – YES – we will be offering all four of our Immersion Events in 2012 with these classes being added to our schedule starting in October. Here's a list of some of the classes and cities we're targeting:

- Available for registration later **this week**:
  - February 2012: Internals and Performance (IE1) in Tampa, FL
  - March 2012: Performance Tuning (IE2) in Tampa, FL
  - March 2012: High Availability/Disaster Recovery (IE3) in Tampa, FL
- Available for registration December/January
  - April 2012: IE1 & IE2 in Chicago, IL
  - May 2012: IE1 & IE2 in London, UK
  - August 2012: IE1-IE2-IE3-IE4 in Redmond/Bellevue (WA) again!

See here for all the details.

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