(March 16th, 2015)

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 Redmond, where we're home for seven weeks – a rarity for us! We're concentrating on producing more Pluralsight courses, revamping our classes for this year's deliveries, and getting some time to work on our hobbies.

Registrations are coming in thick and fast for the next of our popular SQLintersection conferences, in Scottsdale, AZ this May – see here for details. Our show will mainly focus on the troubleshooting and performance problems that we know you're facing today but will also highlight best practices in architecture, design, and SQL Server 2014. For a total of 6 full days – you can immerse yourself into great content with speakers that you know will deliver! Also, use the discount code "SQLskills" and you can save $50 off registration!

Check out this hardware analysis case study that Glenn published in February – he's been doing quite a few of these recently for clients who are upgrading to SQL Server 2014 and replacing old server hardware at the same time.

The two most recent books I've read are Dan Brown's *Inferno*, and Charles Petzold's *The Annotated Turing: A Guided Tour Through Alan Turing's Historic Paper on Computability and the Turing Machine*, and I highly recommend both of them.

I picked up *Inferno* in Sydney airport thinking I'd read it all on the flight home, but then got sucked into watching Interstellar instead (great movie IMHO!) Inferno is a similar to Brown's other books in that it's a page turning thriller with plenty of history thrown in, and involves the same Professor Langdon (art history and symbology) as the other books. I'd say *Inferno* isn't as good as *The Da Vinci Code* but it's clever, with lots of twists and is very entertaining overall, so definitely recommended.

In contrast, Petzold's book is one of the hardest books I've ever read, because of the amount of sheer concentration and thought necessary to understand Turing's paper and the explanation of it. We watched the brilliant movie *The Imitation Game* back in February and I'd always wanted to read Turing's paper so this seemed the ideal opportunity. Don't think that my saying it was hard meant I didn't enjoy it; Petzold did a fantastic job of explaining Turing's paper and the book is excellent. It's not for the faint-hearted though – significant amounts of complicated math and logic are explained and need to be understood. If you want insight into the foundations of computer science, and you're a geek like me, this book is for you.

Note: you can get all the prior Insider newsletters here.

Paul's Ponderings

It's been more than 18 months since I last wrote about my baby – *DBCC CHECKDB* – for the newsletter, so I'd like to give you an assortment of known issues, bug fixes, and general advice.

First of all, you should all be running regular consistency checks on your production databases. This includes master, model, msdb, tempdb, distribution, read-only databases, and databases you don't think anyone's using. Most corruptions happen because of the I/O subsystem, and it's what I call an 'equal opportunity corrupter' – it'll corrupt any of your databases and they're all important.

What's regular? I like to say a full *DBCC CHECKDB* once a week. Easy to do yourself using SQL Agent or you can use Ola Hallengren's handy maintenance solution to do it.

If you can't do that on your main system, consider going the backup+copy+restore+*DBCC CHECKDB* route using another instance. Beware that if you include master in that process, you'll get errors from *DBCC CHECKDB* on the restored copy of master (restored to a different name) as there are some structures in master than can only appear in database ID 1.

You could also consider splitting the checks up over time using a combination of *DBCC CHECKTABLE* and *DBCC CHECKFILEGROUP* – we've done that for a few clients.

It is incorrect to think you can offload the consistency checks to secondary replicas / copies of a database. If you run *DBCC CHECKDB* on a secondary AG replica, for instance, that tells you nothing about the consistency state of the primary replica. Data file pages are not shipped using SQL Server HA/DR technologies, just log blocks, so a corruption in the primary database will not be introduced into its copies such that *DBCC CHECKDB* could detect it.

Remember, restoring the database gives you an exact copy of the pages themselves so it is possible to analyze a restored copy for corruption (knowing that it could give false positives but never a false negative). This is NOT true of a replica.

In actuality, you need to run consistency checks on all copies of a database that are in use or may become used after a failover, including all AG replicas. You should even consistency check a database mirror – although technically you're supposed to buy a license for the database snapshot you'll need to create to do that – I'd be very surprised if Microsoft dinged someone for doing that without a license just to make sure the mirror didn't have corruptions.

As far as performance is concerned, there have been various fixes and problems discovered over the last couple of years:

- Scalability of the readahead mechanism it uses – see KB 2634571 and Bob Ward's post here

- Limiting the degree of parallelism it uses – described in more detail with links in the September 2013 newsletter
- Broken cardinality estimation in the query processor – see the newsletter link above (fixed in 2014!)
- Nonclustered indexes on computed columns causing insane slowdowns due to another Query Processor limitation – see the newsletter above again

Probably the most significant issues come from running *DBCC CHECKDB* at too high a degree of parallelism. I recommend 4 as a maximum for *DBCC CHECKDB* (you'll need to use Resource Governor in Enterprise Edition – but you're on Enterprise already if *DBCC CHECKDB* is running in parallel) otherwise internal bottlenecks become problematic (such systems weren't available at the time I wrote it for testing, so deep wait/latch bottleneck analysis on it was not done back in 2001-2003).

And there are some bugs you should know about:

- A bug where if the database snapshot runs out of NTFS fragments, you'll see error 3449 and a server shutdown on 2012 and 2014. Fix coming in the next CU. See my post here for details.
- A bug where an incorrect boot page flag setting in master and model prevent data purity checks occurring, on all versions from 2005 onward. This you can fix yourselves. See my post here for details.
- A bug where *DBCC CHECKDB* will leak memory when run on a database with columnstore indexes in 2014. Fixed in CU6. See KB 3034615.

Finally, remember that the repair facility of *DBCC CHECKDB* is not infallible – there are various kinds of corruption that repair cannot fix, like broken system table indexes where running repair will actually make things much worse. Backups are the best way to recover from corruption with the minimum of data loss – you should never rely on repair being able to fix everything – it can't and isn't designed to.

Call to action: Make sure you're running consistency checks frequently on your databases, that you're running them efficiently, and that you're not exposed to any of the bugs I've listed above.

Video Demo

One of the things we see over and over with new clients is tempdb configured incorrectly – usually with the default number of data files and auto-growth settings, and without trace flag 1118 enabled. In this video, Glenn demonstrates how to add additional data files to the tempdb database and how to check whether trace flag 1118 is enabled.

The video is about 4 minutes long and you can get it:

- In WMV format here
- In MOV format here

And you can get the code here.

Enjoy!

SQLskills Offerings

We've released most of our classes for 2015. We are likely to add one more domestic delivery of IEPTO1 in the 2nd half of the year but that is currently the only course we're considering.

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

Upcoming Immersion Events

Chicago, IL

- April 27-29, 2015: **IE0**: Immersion Event for the Accidental/Junior DBA ***3 seats remaining!***
- April 27-May 1, 2015: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1) ***SOLD OUT!***
- May 4-8, 2015: **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2 (formerly IE2) ***4 seats remaining!***
- May 4-8, 2015: **IEBI**: Immersion Event on Business Intelligence
- May 11-15, 2015: **IEHADR**: Immersion Event on High Availability and Disaster Recovery (formerly IE3)

Bellevue, WA

- June 8-12, 2015: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)
- June 15-19, 2015: **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2 (formerly IE2)

London, UK

- August 24-28, 2015: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)

Dublin, Ireland

- October 12-16, 2015: **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2 (formerly IE2)

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

Spring SQLintersection

This year our Spring SQLintersection conference will be the week of May 17th in Scottsdale, AZ. See here for details. Don't forget to use the discount code "SQLskills" (without the quotes and it isn't case-sensitive) and you can save $50 off registration! 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 and Kimberly@SQLskills.com