(March 31st, 2014)

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 at home for a few weeks before heading off to the Spring SQLintersection conference in Orlando (I blogged details here) and then our Chicago classes right afterward. I've been slowly replying to everyone who sent me wait stats for my big survey (thanks!) and we're working on more Pluralsight content.

Speaking of Pluralsight, we've published two new courses since the last newsletter:

- Erin's course on *Supporting SQL Server ISV Applications* (see here)
- Glenn's course on *Scaling SQL Server 2012 – Part 1* (see here)

The most recent book I've read is David Hoffman's *The Dead Hand: The Untold Story of the Cold War Arms Race and Its Dangerous Legacy*. This is a really excellent book I picked up while reading *Command and Control* earlier this year. This book details the hidden Soviet bioweapons programs that existed all the way up to the '90s (anthrax, plague, tularemia, and even smallpox, after it had been eradicate by WHO, for goodness sake!) during the Cold War. The Soviets lied to the world, and Gorbachev and Yeltsin continued to lie after the fall of the USSR in '91. It also explains the 'dead hand' semi-automatic mechanism to allow retaliation by the Soviets after a decapitating nuclear strike against them.

Most alarming is the telling of how the Soviet Union fell apart and left thousands of tons of bioweapons, enriched uranium and plutonium, and nuclear weapons spread across the various republics and Russia without adequate storage, security, and safeguards to stop profiteering and proliferation to rogue states like Iran. Interesting to read about how some parts of the US Government stepped in to buy uranium from Kazakhstan to stop it falling into the wrong hands, and built secure storage for Russia. Not an alarmist, sensational book at all, but an insightful and level-headed description of what went on. Well worth reading and strongly 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

One of the many bits of arcane programming terminology is the phrase ‘magic numbers’. A magic number is one that’s been arbitrarily chosen to mean something and the reasoning is not obvious or explained adequately.

The reason I’m writing about it is that several SQL Server magic numbers cropped up in questions I answered during the last week and I thought you’d all be interested to hear about them.

The first one is around database IDs. Some database IDs are well-known, for instance *master*, *tempdb*, *model*, and *msdb* have database IDs 1, 2, 3, and 4, respectively. These make sense as they’re the first four databases created on any SQL Server instance. Other databases get IDs in the order in which they were created/attached/restored, and ID ‘gaps’ can be reused when new databases are added.

If you look at the contents of the buffer pool using the DMV *sys.dm_os_buffer_descriptors* (see my Buffer Pool blog category for a lot more detail), you’ll see some of the database pages are from database ID 32767. If you try to figure out what database that is by doing *SELECT DB_NAME (32767)*, you get back the helpful answer *NULL*. It’s the resource database, *mssqlsystemresource*, whose database ID is a magic number, 32767. Why? Because a developer just chose that as an unlikely number to ever be needed in the space of database IDs.

A weirder case is selecting *DB_NAME (0)*. The question I had during the week was ‘why does *SELECT DB_NAME (0)* always return master, when *master*’s database ID is 1?’

The answer is that it doesn’t. The 0 is a magic number that means ‘the current database’. The returned value from the *DB_NAME ()* function will change depending on your database context.

Another example of a magic number is to do with using *DBCC PAGE* (see here for more info), and this is a question that came up on Twitter on the #sqlhelp alias: ‘what is object ID 99 in the context of DBCC PAGE?’

There is no such thing as object ID 99 – you’ll get *NULL* as the answer from *SELECT OBJECT_NAME (99)*. Object ID 99 is a magic number that means ‘the allocation system’. Any pages who are listed as owned by object ID as 99 are database-wide allocation maps (GAM, SGAM, DIFF_MAP, ML_MAP, PFS) or things like file header pages (page 0 in every file) or the boot page (page 9 in file 1).

Object ID 99 was chosen to represent the allocation system as there had to be some object ID listed, and is the highest non-user object ID, with all system tables having an object ID less than that. Look for yourself by doing *SELECT * FROM sys.objects WHERE [type] = N'S'*.

Magic numbers lose their mystery once you know what they mean.

Call to action: This is another case where learning little bits of idiosyncratic knowledge can help you out understanding what's going on in SQL Server and avoid confusion. It can also help you win arguments with your co-workers ☺. Any time you come across a magic number that doesn't make sense, do a little research, ask on #sqlhelp, or drop me an email. It always pays to spend a little time figuring stuff like this out.

I'm curious to hear your thoughts on magic numbers, so please feel free to drop me a line, always treated confidentially, of course.

Video Demo

In this newsletter's demo video Glenn explains how to create and use a simple ServerMonitor database to record some key instance-level performance metrics using a SQL Server Agent job. He also shows how to query the database to gain insight into what's happening on your server, and explains a bit more background in his recent blog posts here.

The video is just under five minutes long and you can get it:

- In WMV format here
- In MOV format here

You can get the demo code here.

Enjoy!

SQLskills Offerings

Please know that all of our classes will run and their dates will not change. Additionally, most of our public training courses will be held in the first half of this year. We will add a couple of other classes in the second half of the year, but not all that many (maybe one or two IE1 deliveries and one IE2 and, if added, they'll be held in the US). Please plan accordingly.

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

2014 Immersion Events

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 **(SOLD OUT!)**
- May 12-16, 2014: **IE3**: Immersion Event on High Availability and Disaster Recovery
- May 12-16, 2014: **IEDEV**: Immersion Event for Developers
- 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.

SQLintersection is only TWO weeks away!

Our Spring SQLintersection conference has an amazing line-up with 50 SQL sessions, 3 SQL keynotes, and 7 full-day workshops (2 pre-cons on Saturday [April 12], 3 on Sunday [April 13], and 2 post-cons on Thursday [April 17]). The main conference runs from Sunday evening, April 13 through Wednesday, April 16, 2014 in Orlando, FL at the JW Marriott.

Our conference is focused on best practices, architectural/design decisions, platform choices, and new features coming in SQL Server 2014 (including a full day pre-conference workshop, sessions presented by SQL Server team engineers, a booth run by SQLCAT, and multiple sessions presented by industry-experts during the conference).

The specific sessions, titles, and abstracts are on the SQLintersection site and we still have a few incentive packages remaining (including your choice of an XBOX One, Surface 2, or a gift card).

Our confirmed speaker line-up includes: Kimberly L. **Tripp**, Paul S. **Randal**, Brent **Ozar**, Aaron **Bertrand**, Andrew **Kelly**, Bob **Beauchemin**, Bob **Ward**, David **Pless**, Erin **Stellato**, Evgeny **Krivosheev**, Glenn **Berry**, Grant **Fritchey**, Jeremiah **Peschka**, Jonathan **Kehayias**, Jos **de Bruijn**, Kendra **Little**, Kevin **Kline**, Mike **Weiner**, Mike **Zwilling**, Steve **Jones**, and Tim **Chapman**!

Register with the SQLskills discount code and save $50 on your registration!

At this event it's possible to meet with experts so that you can get direct problem-solving help! Speakers are there to work with you and chat – even between their sessions. This is THE *get-your-questions-answered* and *architect-the-RIGHT-solutions* event to go to and it's running in

just under two weeks! And, it’s at the perfect time to get a jumpstart with SQL Server 2014 so you can be ready when SP1 comes out. ☺

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