(March 4th, 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 Japan where Kimberly and I are on vacation doing a wildlife photo tour (snow monkeys, eagles, cranes, and more – photos coming soon!).

Don't forget to check out the upcoming SQLIntersection Conference in Las Vegas in a few weeks! We've also added another IE1 in Chicago at the end of May as the first one has sold out.

The most recent book I've read is Peter Hessler's *Oracle Bones: A Journey Between China's Past and Present*. I love China. I've been there three times for a total of seven weeks, and China's history is engrossing to read about. This book examines the human side of the rebirth that China has experienced over the last 20-30 years, interspersed with nuggets of history. It's a great book and I'm looking forward to reading Hessler's *River Town: Two Years on the Yangtze* and *Country Driving: A Chinese Road Trip*. I also recommend Paul Theroux's *Riding the Iron Rooster: By Train Through China*.

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

At the start of IE2 last week we got into a big discussion about creating multiple filegroups in databases. There are a few reasons you may want to do this for larger databases, including:

- For targeted restores during a disaster
- To support partitioning and better manageability
- Potentially to allow isolation of different workloads within the database onto different portions of the I/O subsystem

I want to focus on the first reason this time.

Imagine you have a 1TB database, with most of the space taken up with a Sales table that contains data from 2009 through today. If the database is destroyed in a disaster, what's the smallest amount you can restore? If it's all in one filegroup there aren't any options – you must restore 1TB, including all of the older data you don't need to be online immediately.

What would be much better for disaster recovery is if the database is split into separate filegroups: primary, 2009, 2010, 2011, 2012, and 2013. If a disaster occurs, you want the 2013 data to be online as fast as possible, as that's what supports your OLTP sales system. As long as you have Enterprise Edition, you can leverage partial database availability and begin this process by restoring only the critical files/filegroups:

```
-- Start with the PRIMARY filegroup
RESTORE DATABASE [SalesDB] FILEGROUP = N'PRIMARY'
FROM DISK = N'D:\SQLskills\SalesDBBackup.bak'
WITH PARTIAL, NORECOVERY, STATS;
GO

-- Bring the OLTP filegroup online
RESTORE DATABASE [SalesDB] FILEGROUP = N'SalesDBSales2013'
FROM DISK = N'D:\SQLskills\SalesDBBackup.bak'
WITH NORECOVERY, STATS;
GO

-- Restore transaction log backups
RESTORE LOG [SalesDB]
FROM DISK = N'D:\SQLskills\SalesDBBackup_Log.bak'
WITH NORECOVERY, STATS;
GO

-- Bring the database online
RESTORE DATABASE [SalesDB] WITH RECOVERY;
GO
```

Then you can restore the other filegroups online and at your leisure. Now you're leveraging another Enterprise Edition feature called online piecemeal restore – a fantastic combination of features for VLDBs.

These features reduce your downtime and give you a lot of flexibility when it comes to prioritizing your restore sequence during a disaster. It also means you can perform a piecemeal restore if only a portion of the database is damaged for some reason – further limiting your downtime requirements if a disaster occurs.

Call to action: Consider your databases – do you have any large ones (more than several hundred GB) that only have a single filegroup? Think about whether you can split them into multiple, smaller filegroups to allow a targeted restore in the event of a disaster. And when planning new databases, consider your filegroup design carefully.

I'm curious to hear your thoughts about multiple filegroups, so please feel free to drop me a line, treated confidentially of course.

Video Demo

From Jonathan:

In this demo video I show how the system_health session in Extended Events proactively collects the xml_deadlock_report event, eliminating the need for additional trace flags, Event Notifications, or SQL Traces to perform deadlock analysis.

The video is about 7 minutes long and you can get it:

- In WMV format here
- In MOV format here

You can get a zip of the demo code here.

Enjoy!

SQLskills Offerings

All of our 2013 public classes are filling up fast! Based on requests from people, attendee ratings of the hotels we used last year, and the ease of using hotels we know, we're using the same locations again. This means we cover both sides of the US, central US, and Europe.

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. We are not planning any other locations for 2013.

- April 29-May 3, 2013: Internals and Performance (**IE1**) in Chicago, IL ***SOLD OUT!***
- April 29-May 3, 2013: Immersion Event for Business Intelligence (**IEBI**) in Chicago, IL – USA (co-located but in a different training room. Attendance is for one event or the other; these cannot be combined for one attendee where they move back/forth.)
- May 6-10, 2013: Performance Tuning (**IE2**) in Chicago, IL – USA
- May 13-17, 2013: High Availability & Disaster Recovery (**IE3**) in Chicago, IL – USA
- May 13-17, 2013: Immersion Event for Developers (**IEDev**) in Chicago, IL – USA (co-located but in a different training room. Attendance is for one event or the other; these cannot be combined for one attendee where they move back/forth.)
- ***NEW*** May 20-24, 2013: Internals and Performance (**IE1**) in Chicago, IL – USA
- May 20-24, 2013: Development Support (**IE4**) in Chicago, IL – USA
- June 3-7, 2013: Internals and Performance (**IE1**) in London – UK
- June 10-14, 2013: Performance Tuning (**IE2**) in London – UK
- June 17-21, 2013: High Availability & Disaster Recovery (**IE3**) in London – UK
- June 24-28, 2013: Development Support (**IE4**) in London – UK
- September 16-20, 2013: Internals and Performance (**IE1**) in Bellevue, WA – USA
- September 23-27, 2013: Performance Tuning (**IE2**) in Bellevue, WA – USA

One thing to note is that the course prices have increased slightly for 2013, reflecting increasing food, logistics, travel, and accommodation costs. We kept our prices the same for the last three years but now we have to raise them a little.

For US classes, the new early-bird price is US$3,295 and the full-price is US$3,795. However, for all past attendees in the 12 months prior to registration, we will only charge the 2012 early bird price of US$2,995. As an alumnus, send us an email prior to registering and we'll give you a code to use to access this special rate. And **be sure to get your registrations in early!**

For UK classes, the new early-bird price is US$3,795 and the full-price is US$4,295. There is a similar past-attendee price equal to the 2012 UK early bird price of US$3,495. **Again, be sure to get your registrations in early!**

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

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