(April 10th, 2017)

If you know someone who 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 is our third letter with our #TBT section (after the demo video section), and we've heard great feedback; make sure to check it out! If you have any suggestions for topics to cover in any portion of the newsletter, drop us a line and we'll see what we can do.

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

SQLskills News

We've started a new initiative where we're all blogging about introductory topics, to help the burgeoning numbers of non-DBAs or junior DBAs who are responsible for SQL Server instances. We're calling it SQL101, and hopefully it'll even help refresh or reinforce topics for the more seasoned DBAs in the community. The blog posts will be automatically collected here.

We're looking forward to seeing a bunch of you this year in our classes in Chicago in April/May and Bellevue in July/August. Note that several classes are sold out and all other classes are 2/3 full. **Also, please note that our May delivery of IEHADR will be the last.** Despite the importance of having a solid disaster recovery plan, many don't plan at all and very few put the resources behind having a well-designed and thoroughly-tested strategy. As a result, we've decided to convert this content into online training instead of in-person training. But, we're still looking at areas where there is more interest; we're particularly pleased to have a brand new **three-day class on using PowerShell** to administer SQL Server, taught by industry-expert, MCM, and MVP, Ben Miller. See here for the complete 2017 class schedule.

Our Spring SQLintersection conference is also fast approaching at the end of May, and we have a phenomenal line up of workshops, sessions, and speakers (including Brent Ozar, Buck Woody, and Bob Ward!). Check out this blog post for all the details, and use the discount code 'sqlskills' when you register to save $50.

Even though we're not teaching any Immersion Events in Europe this year, Kimberly and I will both be **presenting at SQLSaturday #620 in Dublin in June**.

Finally, even if you can't join us in person, we're still taking requests for remote sessions for this year. We have 57 scheduled so far; if you'd like one of us to present for your user group, check out my blog post here.

Book Review

The latest book I've read is Robert Heinlein's classic *The Moon is a Harsh Mistress*. Although I'm a huge sci-fi fan, this is the only one of Heinlein's books that I've read. It's excellent sci-fi from the mid-1960s that covers the revolt of the penal colony on Luna and its eventual secession from Terran rule and emergence as a sovereign state. The revolt and the following short war with Terra are masterminded and made possible by a self-aware computer, and everything in the book is believable. Great stuff and highly recommended.

The Curious Case of…

This section of the newsletter explains problems we've found on client systems; they might be something you're experiencing too.

A couple of weeks back, someone emailed our request alias asking for help with corruption. They had a SQL Server 2008 instance that had been destroyed and they'd salvaged their main database file.

They installed a SQL Server 2012 instance and tried to attach the database, which failed with a *not a primary database file* error, so then they tried a 3rd-party data recovery tool, which failed and corrupted the database some more (by changing the database version number, from the looks of things).

Their only backup was from a year previously, and it didn't have any problems.

There were two big mistakes made here (apart from the obvious one of not having any proper backup strategy):

- Trying to attach the potentially corrupt database to a more recent instance of SQL Server. Corruptions should always be addressed on the same version of SQL Server, as a corrupt database is likely to fail to upgrade and then can't be used on either version.
- Attempting to recover using a 3rd-party tool on the only copy of a corrupt database. Always, always, use a copy of a database to do something like that, in case things get worse.

I advised them to try replacing the file header page of the MDF completely using my boot-page-swap trick (see here) and then attach to a SQL Server 2008 instance, but I unfortunately didn't hear back from them to know whether that helped them or not or whether they had to go back to the year-old backup.

Bottom line: Keep calm when dealing with corruption and take steps that you've practiced, or ask for help if you have no idea what to do. Check out the entry in out SQL101 series where I discuss this: *SQLskills SQL101: Dealing with SQL Server corruption*.

Paul's Ponderings

As databases get larger and larger and as a lot of historical data is kept longer and longer, the interest in how to deal with VLDBs increases. This is an updated editorial from a past newsletter where I discuss things you can do to save space in your databases apart from using data compression.

I've discussed data compression and its uses many times in the past and I've encouraged those of you with Enterprise Edition to check it out as a way to reduce storage costs and increase buffer pool efficiency. Now with SQL Server 2016 SP1, data compression is available for everyone to use. It's definitely something to evaluate, but beware of potential performance impacts on highly-volatile OLTP databases. This is why the info below comes into play to help you reduce you database size…

The benefits of data compression include:

1. The database and backups are smaller
2. Data density increases for large range scans, reducing the number of read I/Os
3. More data fits in memory for general workloads, reducing the number of read I/Os

However, just because data compression is available, that does not mean it's the right solution for reducing the amount of disk space and memory used by your data. There are trade-offs that must be evaluated and even if you *are* using data compression, there are extra things you should do to further reduce wasted space in the database.

First, check to make sure that you don't have duplicate or unused nonclustered indexes. Any extra indexes take up disk space, take up memory in the buffer pool, and their maintenance (during inserts/updates/deletes and periodic fragmentation removal) uses CPU and generates extra transaction log. They're pure overhead and by removing them you can realize benefits #1-3 above.

Check out the following posts for details of what to do:

- Removing duplicate indexes
- Indexes From Every Angle: How can you tell if an index is being used?

Second, make sure you're managing index fragmentation correctly. This means setting appropriate fill factors to help avoid costly page splits that produce low page densities, logical fragmentation, and lots of extra transaction log when they occur.

See the following blog posts for details:

- How expensive are page splits in terms of transaction log?
- Performance issues from wasted buffer pool memory

- How to choose a good index fillfactor

Third, and this one's a lot harder to fix, you may be wasting lots of space in your nonclustered indexes from having large clustered index keys. Every nonclustered index record must include the cluster key from the clustered index record on which it is based. The longer the cluster key is, the more space is taken up in aggregate across the nonclustered indexes.

As an example, consider a cluster key made up of a 16-byte GUID and three 4-byte integers, for a total of 28 bytes. If the table has 1.5 billion rows and 12 nonclustered indexes, the space taken up in the nonclustered indexes for the cluster key is 1.5 billion x 12 x 28 = ~469GB. If the cluster key were changed to an 8-byte bigint, the space taken up in the nonclustered indexes drops to only 1.5 billion x 12 x 8 = ~134GB. That's a pretty big space saving!

See the following blog posts for details:

- How cluster key size can lead to GBs of wasted space
- Code to list potential cluster key space savings per table

Finally, don't discount backup compression. If you only need to reduce the size of your backups, you may find that data compression is overkill – just use backup compression. Even if you're using data compression for the benefits apart from backup size reduction, backup compression can also provide further compression of backups. Just make sure you test the backup compression ratio to make sure you're getting a decent amount of compression that makes it worthwhile to use the extra CPU used when backing up and restoring. See SQL Server 2008: Backup Compression.

Call to action: Whether you're using data compression or not, make sure you're doing everything you can to reduce the size of your databases and make your buffer pool usage as efficient as possible. You'll likely be surprised at just how much space is being wasted in your databases.

Video Demo

In this demo video, Jonathan shows how to add a new replica to an existing Availability Group and configure the replica to automatically seed the databases without manually performing a full and transaction log backup and restore to synchronize the data.

The video is about 6 minutes long and you can get it:

- In WMV format here
- In MOV format here

And the demo code is here.

Enjoy!

#TBT

(Turn Back Time…) This section of the newsletter highlights some older resources we've referred to recently that you may find useful, plus select blog posts we've published since the previous newsletter.

One of the areas that we've been working with quite heavily over the last few months has been replication. Even though it hasn't had any love from the SQL Server team for several releases, it's still very popular. Here are some replication resources that may be useful to you:

- Joe's Pluralsight course on *SQL Server: Transactional Replication Fundamentals* that has 2 hours explaining what replication is, how to set it up, how to monitor and troubleshoot it, and how to combine it with other HA technologies
- Whitepaper on *SQL Server Replication: Providing High-Availability Using Database Mirroring* that I wrote for Microsoft back in 2008, but is still completely relevant today and can be extrapolated to working with Availability Groups too.
- Old whitepaper on *Proven SQL Server Architectures for High Availability and Disaster Recovery* that I wrote for Microsoft in 2010, and isn't listed on their site any more, but has an interesting set of architectures, including one using peer-to-peer replication.
- And a bunch of blog posts:
  - *In defense of transactional replication as an HA technology*
  - *REPLICATION preventing log reuse but no replication configured*
  - *The Transactional Replication Multiplier Effect*
  - *When is the Publication Access List required?*
- And check out the Repltalk blog on MSDN that's been a fount of replication knowledge since 2010.

Here are a few of the blog posts we've published since the last newsletter:

- Paul: *SQLskills SQL101: Switching recovery models*
- Kimberly: *SQLskills SQL101: Partitioning*
- Kimberly: *SQLskills SQL101: Indexes on Foreign Keys*
- Erin: *SQLskills SQL101: Updating SQL Server Statistics Part I – Automatic Updates*
- Erin: *SQLskills SQL101: Updating SQL Server Statistics Part II – Scheduled Updates*
- Glenn: *SQL Server 2016 Service Pack 1 CU2 Released*
- Glenn: *Performance and Stability Related Fixes in Post-SQL Server 2014 SP2 Builds*
- Glenn: *SQL Server Diagnostic Information Queries for March 2017*

I hope you find these useful and interesting!

Upcoming Immersion Events

The classes for both Spring and Summer 2017 are available for registration!

To help your boss understand the importance of focused, technical training, we've also added a few 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

Chicago, IL, April/May 2017

- **IE0**: Immersion Event for Junior/Accidental DBAs
  - April 24-26 ****SOLD OUT****
- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - April 24-28 ****SOLD OUT****
- **IESSIS1**: Immersion Event on Learning SQL Server Integration Services
  - April 24-28 ****SOLD OUT****
- **IEBI**: Immersion Event on Business Intelligence
  - May 1-5 **** Only 4 seats remaining! ****
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - May 1-5 **** Only 2 seats remaining! ****
- **IEPS**: Immersion Event on Powershell
  - May 8-10 ****Only 4 seats remaining! ****
- **IEPDS**: Immersion Event on Practical Data Science
  - May 8-12 **** Only 2 seats remaining! ****
- **IEHADR**: Immersion Event on High Availability and Disaster Recovery
  - May 8-12

Bellevue, WA, July/August 2017

- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - July 31-August 4
- **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2
  - August 7-11

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

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