

SDB306
DBA Mythbusters

Paul S. Randal
SQLskills.com
Paul@SQLskills.com

MICROSOFT
SQL
SERVER
conference and expo


Author/Instructor: Paul S. Randal

- Consultant/Trainer/Speaker/Author
- CEO, SQLskills.com
 - Email: Paul@SQLskills.com
 - Blog: <http://www.SQLskills.com/blogs/Paul>
 - Twitter: @PaulRandal
- Contributing Editor of TechNet Magazine and SQL Server Magazine, author of the SQL Q&A columns, bi-monthly articles on DBA topics, multiple 2008 whitepapers
- 5 years at DEC responsible for the VMS file-system and check/repair
- Almost 9 years as developer/manager in the SQL Storage Engine team through August 2007, ultimately responsible for Core Storage Engine
 - Wrote DBCC component, other Storage Engine code, DBCC CHECKDB/repair for SQL 2005
- Regular presenter at worldwide TechEds, SQL PASS and other conferences on disaster recovery, HA, maintenance, and internals
- Wrote and presented SQL Server 2008 training for Microsoft
- Course author/instructor of Microsoft Certified Master certs
- Co-Chair of the SQL Connections conference



Microsoft
Most Valuable
Professional

Microsoft
Regional Director
PROGRAM



MICROSOFT
**SQL
SERVER**
conference and expo

Why Is This Important?

- Lots of myths and misconceptions have grown over the years about how SQL Server behaves
- Adherence to these misconceptions can lead to:
 - Bad practices
 - Wasted time and resources
 - Confusion
 - Arguments 😊
- Let's debunk a bunch of them!

Myth #0

BUSTED!!

- Oracle is much better than SQL Server
- This one's obviously untrue! 😊
- On to the real ones...

Myth #1

BUSTED!!

- When a database snapshot is created, it makes a copy of the entire database
- Database snapshots use NTFS sparse files for storage
- Although the sparse files present themselves as the same size as the matching database data files, they start very small and grow as required
- This is why database snapshot creation is usually very fast
- Snapshots do not reserve disk space, so can run out of space and become unusable

Myth #2

BUSTED!!

- A SUSPECT database can be repaired by detaching and attaching the database
- A SUSPECT database is one where crash recovery could not complete
- Attaching a database that was not cleanly shutdown requires completing crash recovery
 - If crash recovery cannot complete then the database will not attach again
- Re-attaching a detached SUSPECT database requires hacking it back into SQL Server

Myth #3

BUSTED!!

- Backups cause blocking
- Backups do not take locks so cannot cause blocking
- Backups are very I/O intensive so might cause I/O contention which **appears** to be blocking as the workload slows down drastically
- Consider:
 - Taking large data backups at off-peak times
 - Taking more frequent transaction log backups
 - Separating database and backup file storage

Myth #4

YES AND NO !!

- Instant file initialization can be controlled from within SQL Server
- IFI allows skipping of zero initialization for data files
- SQL Server service account must be granted a Windows privilege to enable IFI, and then SQL Server restarted
- Use TF 3004 to watch zero initialization happening
- Use TF 1806 to temporarily disable IFI
- Btw, IFI is available in **all** editions of SQL Server

Myth #5

BUSTED!!

- Online index operations do not acquire locks
- There are two locks held by online index operations
- 'Almost Online Index Operations' isn't a very good marketing term 😊
- Shared table lock held at start to allow new index to be created
- Schema-Modification lock held at end to drop old index
- Both locks can cause blocking while they're queued

Myth #6

YES AND NO!!

- You can use WITH STOPAT when restoring a full or differential backup
- The restore syntax allows the use of WITH STOPAT when restoring data backups, but it does nothing
- It's only allowed as a syntactical nicety so you can use WITH STOPAT on all commands in the restore sequence to avoid accidentally going past the time/transaction required
- Data backups only allow a single point-in-time restore – as of the time of the end of reading the data

Myth #7

BUSTED!!

- DBCC CHECKDB causes blocking
- In SQL Server 7.0 and before, it did take blocking locks
- DBCC CHECKDB has been online by default since SQL Server 2000
- It is the most I/O intensive operation in SQL Server, as its readahead mechanism is very aggressive, so the I/O load can make it *seem* like the workload is blocked

Myth #8

BUSTED!!

- An index FILLFACTOR is used by regular DML operations
- FILLFACTOR only applies during index build and rebuild operations
- Same for the PAD_INDEX setting
- (Btw, a fill factor of 0 is the same as 100 – no space)

Myth #9

BUSTED!!

- Lock escalation goes row to page, then page to table
- Lock escalation always goes directly to a table lock
- SQL Server 2008 introduced a partition-level lock
 - Default is still table-level lock escalation as some customers experienced deadlocks when partition-level escalation was tried as the default during early 2008 CTPs

Myth #10

YES AND NO!!

- Rebuilding a clustered index always rebuilds all nonclustered indexes
- In SQL Server 2000 and before, this could happen...
 - A non-unique clustered index has a 'uniquifier' column that serves to allow a unique index key to be generated
 - Non-clustered indexes include the clustered index key
 - SQL Server 2000 would regenerate the clustered index uniquifiers on a rebuild, invalidating all the nonclustered indexes
 - SQL Server 2000 SP1 also had a bug that cause nonclustered indexes to always be rebuilt
- Doesn't happen ever in SQL Server 2005 onwards!

Myth #11

BUSTED!!

- Failure detection is instantaneous in database mirroring
- It depends on what the failure is:
 - Fast: SQL Server instance crashes
 - Slow: Windows crashes
 - Slower: Log drive failure
 - Maybe: Corrupt page

Myth #12

BUSTED!!

- Failover is instantaneous in database mirroring
- Automatic failover only happens if:
 - Synchronous database mirroring is being used
 - A witness server is configured
 - The mirror is in the SYNCHRONIZED state
- Even then, mirror database will not come online until all its REDO queue of transaction log has been recovered
 - A SYNCHRONIZED mirror could still have a lot of transaction log to process – meaning seconds, minutes, or hours of delay!
- Monitor the REDO queue on the mirror!

Myth #13

BUSTED!!

- A nonclustered index should be created for each column in the table
- Nonclustered indexes are only used in query plans if:
 - They are selective enough AND
 - They are the most efficient choice for getting the query results
- It's unlikely that an index on each column will be useful
- Use `sys.dm_db_index_usage_stats` to check whether an index is being used or not
 - Make sure to check over an entire business cycle

Myth #14

IT DEPENDS!!

- Failover clustering is the best HA technology to use
- Failover clustering protects against a failed server but has a big problem: no redundancy at the I/O subsystem level
- You must provide some additional protection for the data like SAN replication (a.k.a. SAN mirroring)
- Failover clustering is often not the most appropriate HA technology, based on *requirements*

Myth #15 **YES AND NO !!**

- Concurrent data and log backups are not possible
- This was true in SQL Server 2000 and before
 - The log backup would block, waiting for data backup completion
- No longer true for SQL Server 2005 onwards
 - Log backups can run concurrently with data backups
 - Log truncation/clearing is deferred until the data backup completes (see next Session for details)
- This gives rise to the next myth...

Myth #16

BUSTED!!

- A full backup allows the log to clear/truncate
- In the FULL or BULK_LOGGED recovery models, the only thing that clears the log is a log backup
- In the SIMPLE recovery model, the only thing that clears the log is a checkpoint
- Nothing else.

Myth #17 **YES AND NO !!**

- You cannot run DMVs against a database in the 80 compat-mode
- You cannot do it if you're in the context of that database
- You can do it if you run the DMV in the context of a database in the 90 (or higher) compat mode

Myth #18

BUSTED!!

- You can cause corruption using a T-SQL statement, or interrupting a T-SQL statement
- Common statements cited: shrink and index rebuild
- You can't.
- Unless there's a bug in SQL Server 😊

Myth #19

BUSTED!!

- You can restore parts of a database to different points in time
- In any piecemeal restore operation, the piece of the database being restored must be restored to the same point in time as the primary filegroup
- And you can't mix-n-match filegroups from different databases either.

Myth #20

BUSTED!!

- You can restore/attach a database to an earlier version of SQL Server (e.g. 2005 database to a 2000 server)
- This has never been possible and never will be
- Don't confuse database compat mode with database physical version
- Upgrade changes the physical structure of the database
 - Earlier version have no way to understand the newer structures
 - New log records are introduced in new versions too
- Only solution is export + import

Myth #21

BUSTED!!

- Nested transactions exist in SQL Server
- Although the syntax is there, it's meaningless
- Committing a 'nested' transaction does nothing
- Rolling back a 'nested' transaction rolls back the entire outer transaction, regardless of whether you think you've committed 'nested' transactions
- Nothing commits until @@TRANCOUNT = 0
 - This leads to problems with transaction log growth
 - See next Session for details

Myth #22

BUSTED!!

- FILESTREAM data can be stored on a network share remote from the SQL Server instance
- The FILESTREAM data container (the NTFS directory holding all the FILESTREAM data) must be local to the SQL Server instance
 - Just the same rules as for regular SQL Server data files
- The FILESTREAM data can be accessed remotely by SQL Server clients
 - This is where the misconception comes from

Myth #23

BUSTED!!

- Turning on page checksums protects all pages in the database immediately
- Turning it on does nothing
- A page does not get a page checksum until it is read from disk, changed, and then written back to disk
- There is no tool to 'touch' every page and put a page checksum on it

Myth #24

BUSTED!!

- You can avoid running DBCC CHECKDB as long as you use BACKUP ... WITH CHECKSUM
- A page could be corrupt AND have a valid checksum
- The page could be corrupted in memory by:
 - A rogue process scribbling in SQL Server's buffer pool
 - A badly written extended-stored procedure corrupting memory
 - A bad memory chip
 - A SQL Server bug
- And then written to disk with a valid checksum
- Only DBCC CHECKDB can detect this

Myth #25

BUSTED!!

- It is possible to have multiple mirrors of a database with database mirroring
- Only one mirror per principal database
- But can also have multiple log-shipping secondaries for warm-standby
- Also cannot have 'chaining' with database mirroring (i.e. a database is a mirror, and a principal for another mirror)

Myth #26

BUSTED!!

- DDL triggers are INSTEAD OF triggers
- DDL triggers are (unfortunately) AFTER triggers
 - The operation is performed and then rolled-back
- Can be better to use user/schema separation and explicit permissions than relying on DDL triggers

Myth #27

BUSTED!!

- SELECT COUNT (*) will always use a table scan
- The query processor will pick the table structure that has the smallest number of pages
- This will be the smallest nonclustered index (if present)
 - Must be a non-filtered index in SQL Server 2008

Myth #28

BUSTED!!

- Tempdb should be X% the size of the largest database
- Variations:
 - The transaction log should be X% the size of the database
 - Auto-growth should be X% the size of the file
- No validity for any of these
- Size of any of them *depends* big-time
- Tempdb: size of largest memory spill or temp table
- Log: size of largest single transaction, or amount of log between log truncates
- Auto-growth: more complicated for log and data

Myth #29

BUSTED!!

- In-flight transactions are preserved after a failover
- In-flight transaction are **never** preserved when a connection is terminated
- Crash recovery is always involved, which rolls-back uncommitted transactions
- Only exception is VM that undergoes live migration
 - Connections are not dropped, migrated to new physical host transparently

Myth #30

TRUE !!

- I spent more time working on DBCC CHECKDB than I've spent being married to Kimberly Tripp
- But it wasn't as much fun!
- And she's catching up fast 😊

Plenty More Myths Around...

- Database repairs
- Index operations
- Backup & restore
- Transaction log
- Resource governor
- Performance
- Well, anything around SQL Server!!



Summary

- Make sure you corroborate what you read online
- If something sounds fishy, try it yourself!

- Checkout my Debunk-A-DBA-Myth-A-Day blog post series all through April:
 - <http://www.SQLskills.com/blogs/paul/Misconceptions.aspx>

- Don't forget the next three sessions:
 - Understanding Logging and Recovery
 - Building the Right Backup Strategy
 - Follow the Rabbit: Open Q&A



Your Feedback is Important

Please fill out a session evaluation form
drop it off at the conference registration
desk.

Thank you!