



Sydney SQL Server
Enterprise User Group

SQL Server Myths and Misconceptions

Victor Isakov

victor@sqlserversolutions.com.au

Abstract

As a DBA you have heard of plenty of myths and misconceptions about SQL Server. From how "you can improve transaction log performance with instant file initialisation" to whether "non-logged operations still exist in SQL Server". In this session Victor Isakov will go through over 30 different myths and misconceptions about SQL Server and explain what actually happens in the background. This should be a great session to both test your SQL Server knowledge and to learn how to best administer and configure SQL Server.

Speaker

- Victor Isakov is a Database Architect / Trainer / SQL Ranger who provides consulting and training services to various organizations in the public, private and NGO sectors globally, and been involved in different capacities at various international events and conferences. Victor specializes in:
 - Still does “high-end” SQL Server training
 - Performance tuning and optimization
 - “Health-checks” / “Risk Assessments” / review of SQL Server infrastructure
 - Architecting / re-factoring database solutions
 - Assessing the effectiveness of your outsourced services / licensing
 - Consolidating / upgrading SQL Server infrastructure
- Blog: www.victorisakov.com
- Email: victor@sqlserversolutions.com.au
- Website: www.sqlserversolutions.com.au



Training

- Victor Isakov is only *Microsoft Certified Trainer* in the world who is also a *Microsoft Certified Master* and a *Microsoft Certified Architect* and a *Microsoft Valuable Professional*
- Differentiation:
 - Trained SQL Server since version 6.0
 - Has been consulting in SQL Server technology stack
 - Various organisations in public, private and NGO sectors
 - From short-term projects (weeks) to long-terms project (2 years)
 - Wrote / designed some of the Microsoft Official Curriculum
 - Wrote / designed / validated the SQL Server exams
 - Wrote a number of books on SQL Server
- What does this mean?
 - Offers what we strongly believe is the best SQL training in Asia / Pacific
 - Leverage over a decade of experience
 - Leverage in-depth technical knowledge
 - Knows Microsoft certification and exams in depth
- Running the DBA course in late-February / early-March

Agenda

- Why?
- Myths / Misconception
 - Vote
 - Facts
 - Background
 - Tips
- Be careful of SQL Server versions
 - We are concentrating on SQL Server 2005+

Why?

- Lots of myths / misconceptions “out there”
- Lead to bad practices
- Wasted time and resources
- Confusion
- Water cooler arguments

Database Snapshots

- When creating a database snapshot the database engine creates a copy of the entire database
 - **FALSE**
 - Database snapshots use NTFS sparse file technology
 - Space is not reserved
 - Tip: Can cause DBCC CHECKDB to fail
 - Use copy-on-write technology

Suspect Databases

- You can repair a SUSPECT database by detaching and attaching it
 - **FALSE**
 - In SQL Server 2008+ you cannot detach a suspect database
 - In SQL Server 2000 there was one state
 - Suspect
 - Crash recovery can't run
 - In SQL Server 2005+ there are 2 states
 - Recovery pending
 - Database was not cleanly shut down
 - Suspect
 - Crash recovery started but could not finish

Database Backups

- When performing a backup you can cause blocking
 - **FALSE**
 - Backups do not take locks
 - Database performance can degrade due to I/O contention

Instant File Initialization

- Instant file initialization is available in Enterprise Edition only
 - **FALSE**
 - Available in all SKUs of SQL Server

Instant File Initialization

- Instant file initialization applies to both database and transaction log files
 - **FALSE**
 - Instant file initialization does not apply to transaction log files
 - Database engine needs to zero-out the transaction log file for integrity purposes

Instant File Initialization Configuration

- Instant file initialization (IFI) can be configured (configured) from SQL Server
 - **FALSE** (NOT REALLY)
 - Windows right SE_MANAGE_VOLUME_NAME
 - Through xp_cmdshell
 - Can control behaviour
 - Trace flag 1806 disables IFI temporarily
 - Transparent Database Encryption (TDE)
 - Tip: Trace Flag 3004 (+3605) writes to errorlog

Online Index Operations

- Online index operations do not acquire locks
 - **FALSE**
 - Two locks held by online index operations
 - Share table lock acquired at start
 - New empty index (in-build) created
 - Schema version (minor) modified to invalidate existing query plans
 - Schema-Modification lock held at end
 - Drop old index
 - Schema version (major) modified
 - Both locks can cause blocking
 - No locks are acquired / held during index creation
 - Tip: For true 24/7 shops might have to use REORGANIZE

STOPAT Restore Option

- You can use WITH STOPAT when restoring a full or differential backup
 - FALSE | TRUE
 - Can use it
 - Does nothing

DBCC CHECKDB

- DBCC CHECKDB causes blocking within database
 - **FALSE**
 - DBCC CHECKDB uses database snapshots in SQL Server 2005+

FILLFACTOR

- FILLFACTOR is used by DML operations (INSERTs, UPDATEs, DELETEs)
 - **FALSE**
 - FILLFACTOR only applies during index builds and index rebuilds
 - Tip: FILLFACTOR 0 == FILLFACTOR 100

Lock Escalation

- Lock escalation goes from row-level to page-level to table-level
 - **FALSE**
 - Lock escalation always goes to table-level
 - Row-level -> table-level
 - Page-level -> table-level
 - SQL Server 2008 introduced partition-level lock
 - Tip: can control lock type at table / index level

Rebuilding Clustered Indexes

- Rebuilding a clustered index always rebuilds all nonclustered indexes
 - **FALSE**
 - SQL Server 2005+ retains the same uniquifier value for non-unique clustered index keys

Database Mirroring Failure

- A failure is instantaneously detected in database mirroring
 - **FALSE**
 - Mirror pings Principal every 1 second
 - Mirroring partner timeout defaults 10 seconds
 - I/O completion failure
 - After 20 seconds write to ERRORLOG
 - After 40 seconds take database offline
 - Depends on failure:
 - SQL Server instance crashes
 - Windows crashes
 - Log drive
 - Corrupt page

Database Mirroring Failover

- A database mirroring failover occurs instantaneously
 - **FALSE**
 - Automatic failover only happens the mirroring partnership is SYNCHRONIZED and quorum is formed
 - Depends
 - Failover detection
 - Mirror's redo queue
 - Tip: Don't forget to monitor the REDO queue

Nonclustered Indexes

- A nonclustered index should be created on every column in a table
 - **FALSE**
 - Check `sys.dm_db_index_usage_stats` to see whether an index is being used
 - Seeks
 - Scans

Failover Clustering

- Failover clustering is the best High-Availability technology to use
 - **FALSE**
 - Failover clustering has no redundancy at disk I/O subsystem
 - Depends on business
 - RPO
 - RTO
 - Performance overhead
 - Budget 😊

Database Backups

- A full database backup clears the transaction log
 - **FALSE**
 - Only a log backup clears the transaction log when in FULL or BULK_LOGGED recovery model
 - When in SIMPLE recovery model the CHECKPOINT clears the log

Log Backups

- You cannot take a concurrent transaction log and database backup
 - **FALSE**
 - You can perform a log backup when a FULL or DIFFERENTIAL backup is being performed in SQL Server 2005+

DMVs

- You cannot run DMVs against a database in 80 compatibility level
 - **FALSE**
 - Do not support calling functions as a DMV parameter
 - OBJECT_ID(...)
 - Might need to run it from a 90 compatibility level database
 - Use three-part name

Database Corruption

- You can cause database corruption using a T-SQL statement, or interrupting a T-SQL statement (such as shrinking a database or rebuilding an index)
 - **FALSE**

Point-in-Time Recovery

- You can restore parts of a database to different points in time
 - **FALSE**
 - Always need to restore to same point in time with piecemeal restores
 - `READ_ONLY` file groups work differently

Database Versions

- You can restore or attach a database to an earlier version of SQL Server
 - **FALSE**
 - Physical structure of database different
 - Different transaction log “architecture”
 - Only solution is to export and import data and code

Database Versions

- You can restore or attach a database to an earlier version of SQL Server
 - **FALSE**
 - Physical structure of database different
 - Different transaction log “architecture”
 - Only solution is to export and import data and code

Nested Transactions

- SQL Server supports nested transactions
 - **FALSE**
 - @@TRANCOUNT increments each time a new BEGIN TRAN is encountered
 - Need to COMMIT TRAN until @@TRANCOUNT reaches 0
 - Any ROLLBACK TRAN rolls back all “nested transactions”

FILESTREAM

- FILESTREAM data can be stored on a network share remote from SQL Server
 - **FALSE**
 - FILESTREAM data container must be local to the SQL Server
 - FILESTREAM data can be accessed remotely

Page Checksums

- Turning on page checksum immediately “protects” all pages
 - **FALSE**
 - A page checksum is not generated until it needs to be written back to disk.
 - No tool exists that “touches” every page in your database
 - Otherwise need to rebuild all indexes
 - But that will not fix allocation pages, system tables, etc
 - You could turn Transparent Database Encryption (TDE) on then off

Avoiding DBCC CHECKDB

- You can avoid running DBCC CHECKDB as long as you use BACKUP ... WITH CHECKSUM
 - FALSE
 - A page with corrupt data but a valid checksum can occur
 - The page could be corrupted in memory (Error 832)
 - Rogue process
 - SQL Server bug
 - Faulty memory

Multiple Database Mirrors

- With database mirroring it is possible to have multiple mirrors of a mirrored database
 - **FALSE**
 - Only one mirror per principal
 - Cannot chain mirrored database
 - Log shipping / replication is the only such solution
 - Tip: Wait for SQL Server 2011 (“Denali”)

DDL Triggers

- DDL Triggers are INSTEAD OF triggers
 - **FALSE**
 - Operation is performed and then rolled back

COUNT(*)

- SELECT COUNT(*) will always use a table scan
 - FALSE
 - The query processor needs only to scan the smallest nonclustered index
 - Non-filtered index in SQL Server 2008

Tempdb Size

- Tempdb should be x% the size of your largest database
 - **FALSE**
 - Depends on many factors
 - Needs to support the largest transaction that occurs on the SQL Server instance

Tempdb Files

- Tempdb should always have one data file per processor core
 - **FALSE**
 - Trying to alleviate PAGELATCH waits
 - Trace Flag –T1118
 - SQL Server 2005+ has been further optimized
 - Microsoft initially said 1:1
 - Then seen 1:4 to 1:2 ratio
 - Too many files can also degrade round robin allocation performance

In-flight Transactions

- In-flight transactions are preserved after a failover in a clustered environment
 - **FALSE**
 - In-flight transaction are not preserved
 - Crash recovery is always involved
 - Uncommitted transaction will be rolled back
 - Trace Flag –T1118
 - However...
 - Consider using live-migration / V-motion in a virtual environment

Transaction Log

- It is not possible to analyse the transaction log using SQL Server
 - **FALSE**
 - DBCC LOG is “old school”
 - fn_dblog

Transaction Log

- It is not possible to analyse the transaction log using SQL Server
 - **FALSE**
 - DBCC LOG is “old school”
 - Use fn_dblog instead
 - Also be aware of fn_dump_dblog
 - Allows you to dump transaction log from a dump backup

Restoring Database Backups

- You can restore a database backup to an earlier version of SQL Server
 - **FALSE**
 - All depends on physical version of the database
 - Database engine only understands a certain range of physical versions

Fixing Database Corruptions

- Restarting SQL Server will fix database corruptions
 - **FALSE**
 - Actually that has solved the problem in extremely, extremely, extremely rare and extremely, extremely, extremely luck circumstances
 - Due to “stuck” I/Os at a SQL Server and I/O subsystem level
 - Not a recommended course of action

Resource Governor

- Resource Governor allows you to govern I/O
 - **FALSE**
 - Governs memory and processor
 - Query execution memory grants
 - Not buffer pool memory

Q & A

- Questions?
- Email: victor@sqlserversolutions.com.au
- Blog: www.victorisakov.com
- Website: www.sqlserversolutions.com.au