

The Database is Slow

SQL Server Performance Tuning Starter Kit

Calgary PASS Chapter, 19 August 2015

Randolph West, Born SQL

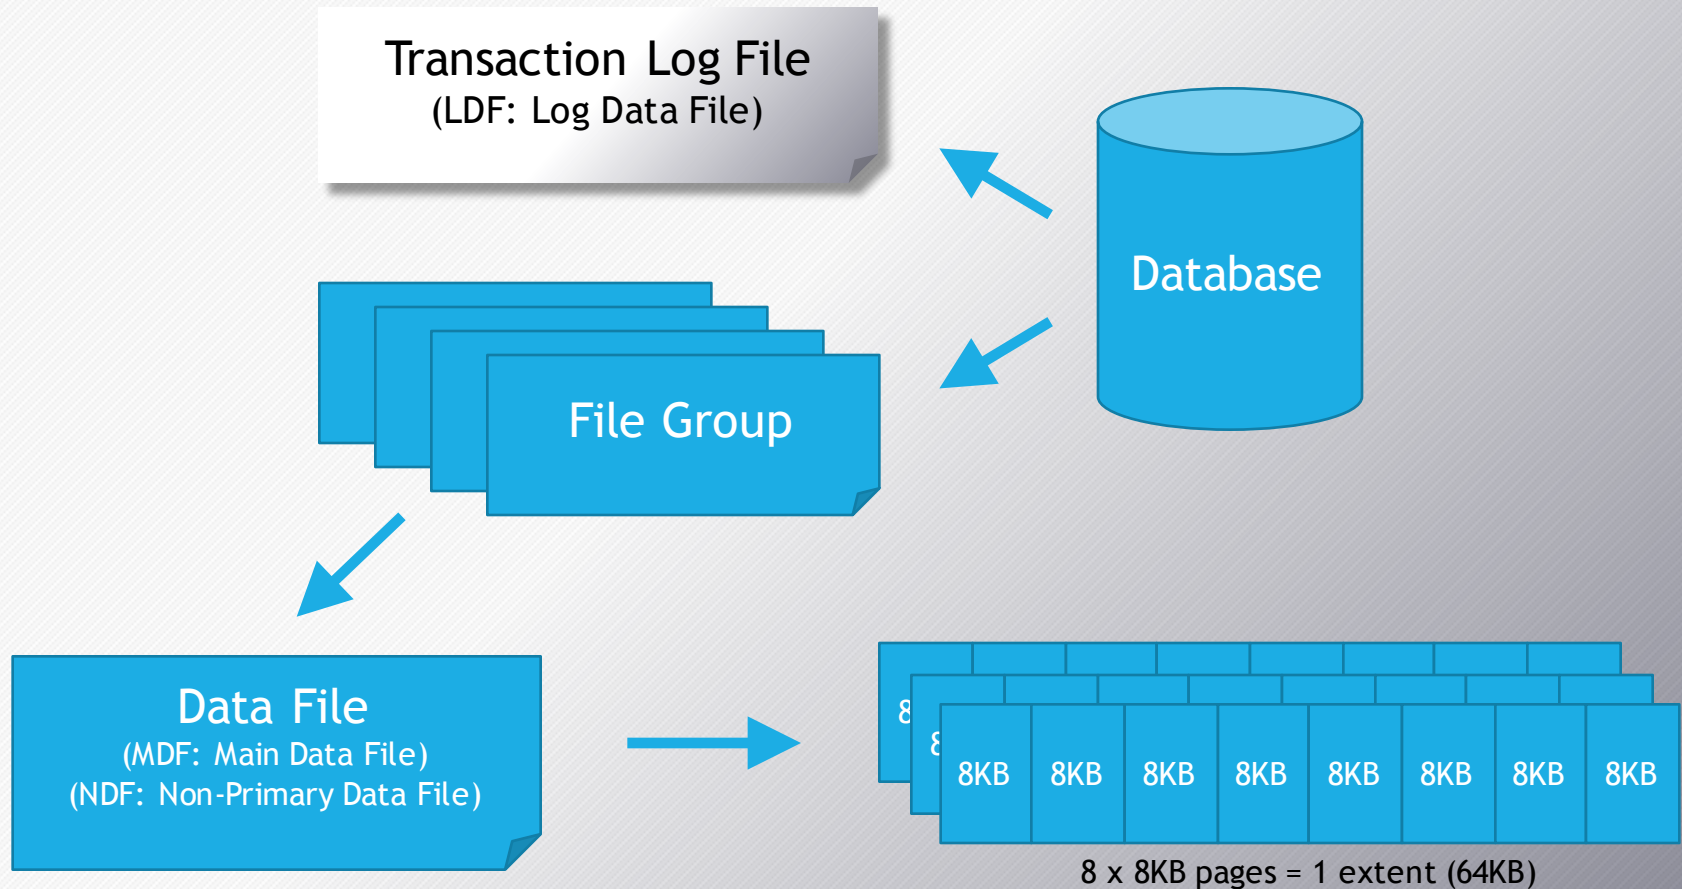
Email: r@ndolph.ca

Twitter: [@rabryst](https://twitter.com/rabryst)

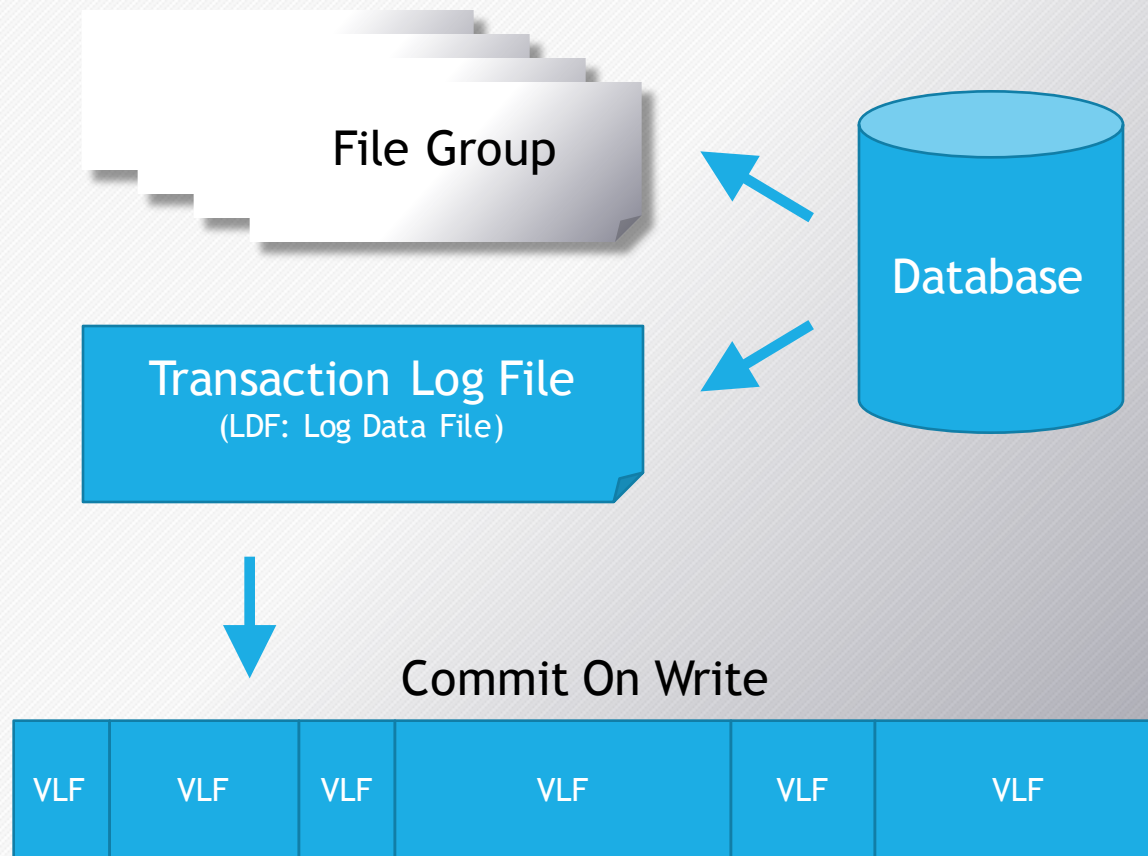
Basic Internals

- Data File
- Transaction Log File
- Data Page
- Indexes
- Statistics

Basic Internals: Data File

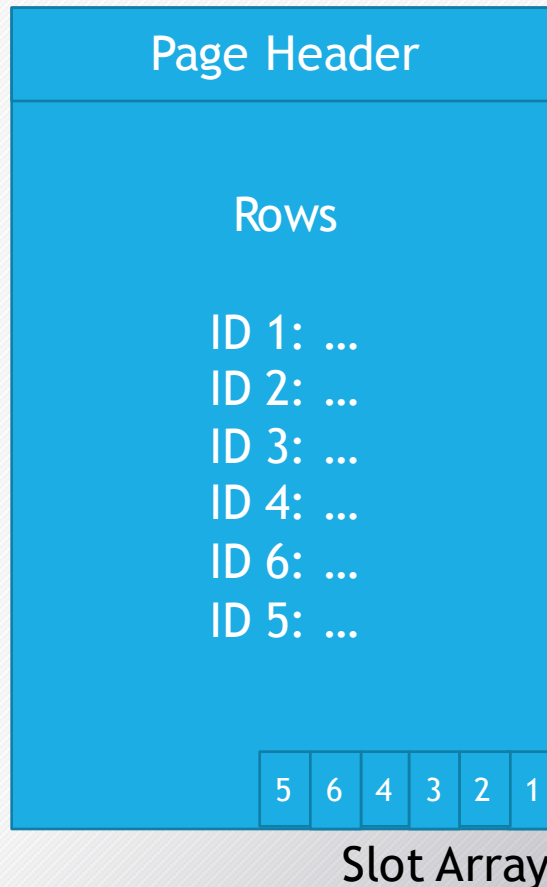


Basic Internals: Transaction Log



Basic Internals: Page

8KB

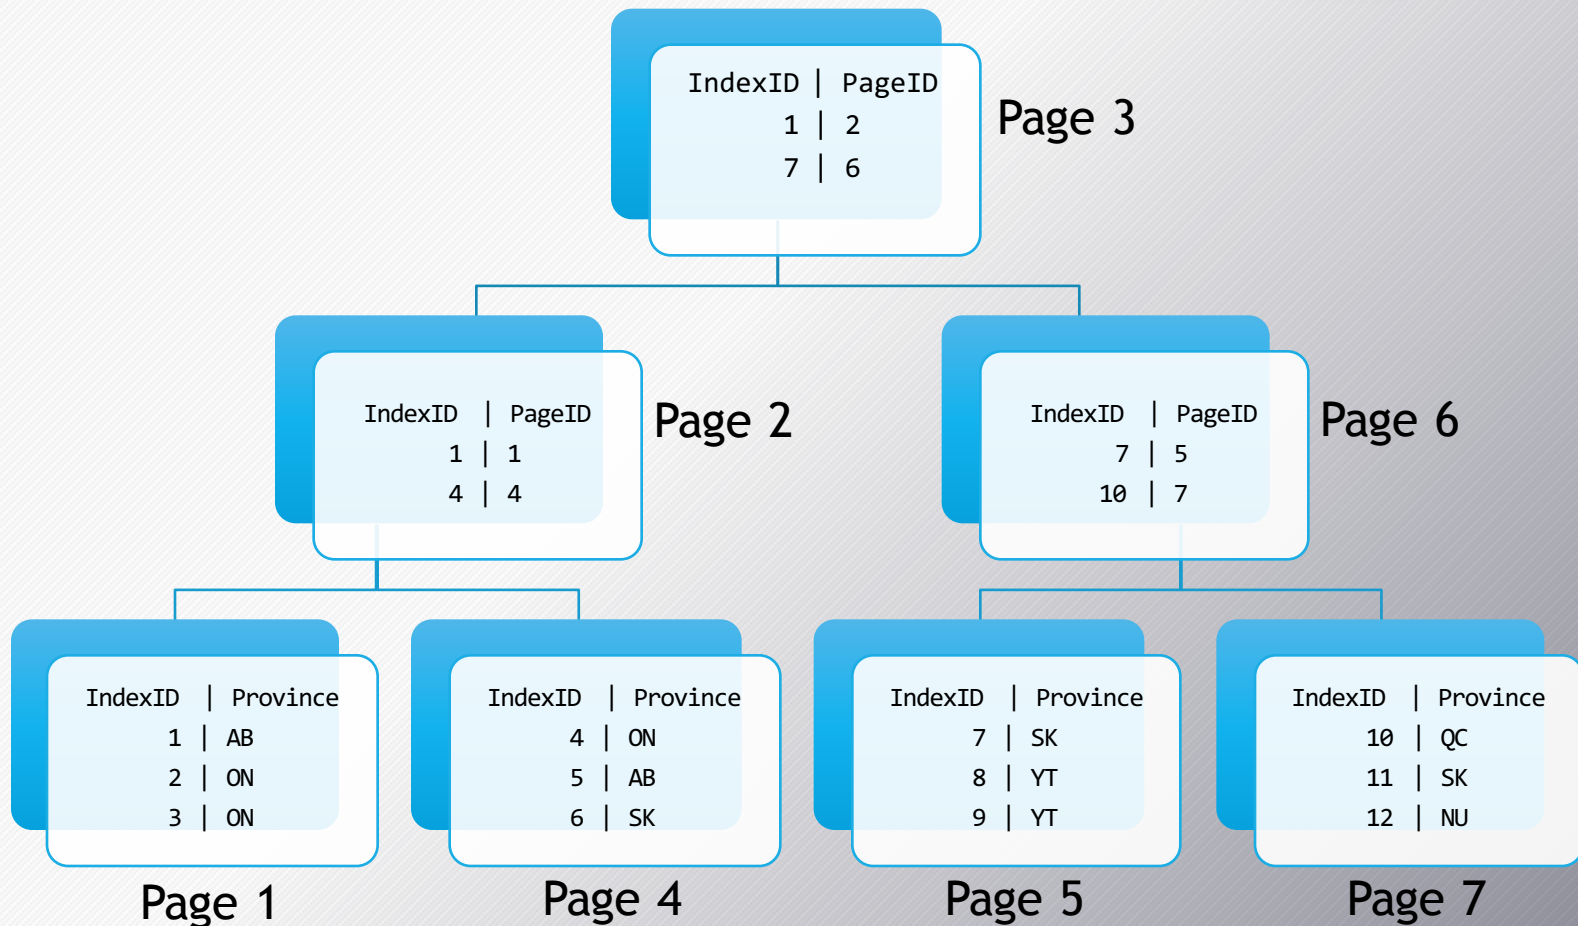


Paul Randal (SQLskills.com):

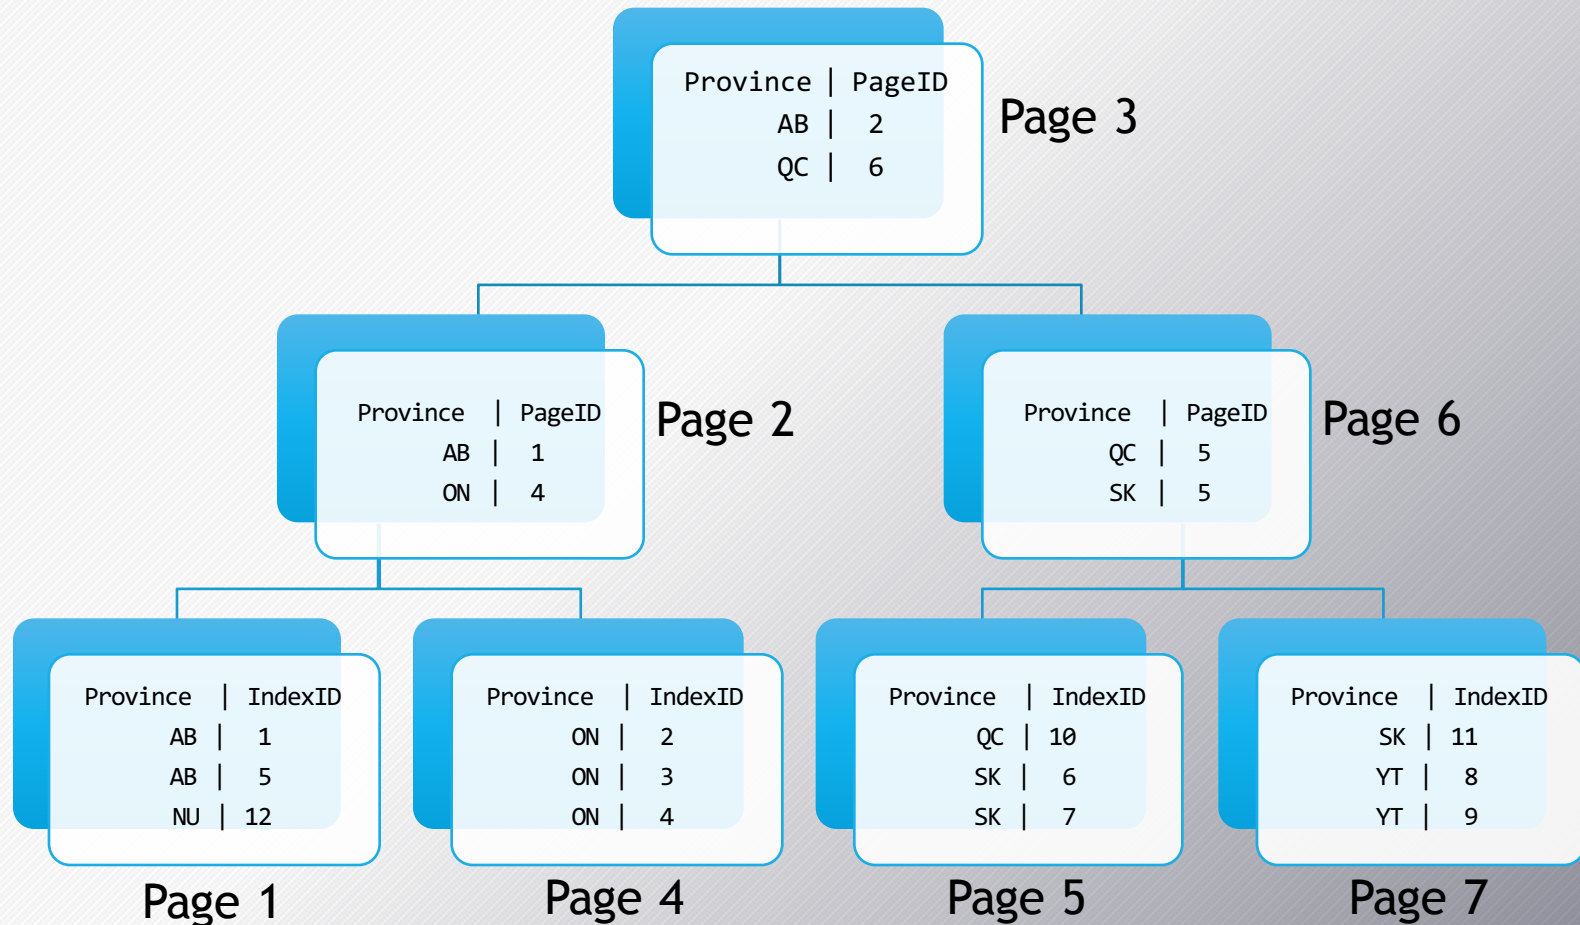
Anatomy of a Page

<http://goo.gl/gWNjT0>

Basic Internals: Tables



Basic Internals: Indexes



Basic Internals: Indexes (Contd.)

- Heap
 - No B-tree, no links - all info stored in Index Allocation Map
- Clustered Index
- Non-Clustered Index
- Unique Indexes and Constraints
- Filtered Index
- Indexed View
- Fill Factor

Basic Internals: Statistics

- Distribution of Data (up to 200 steps)
- Density of Data (rows per step)
- Tipping Point (Selectivity of Data)
- Auto updated after 20% plus 500 rows changed
- All query plan costs are based on these values
- DBCC SHOW_STATISTICS
- Read more by Gail Shaw and Kimberly L Tripp

Where Do I Begin?

- Hardware Configuration
- Operating System
- SQL Server Configuration
- Monitoring Tools

Hardware

- CPU (SMP or NUMA?)
- RAM (Memory pressure?)
- Storage (Lots of things)
- Network (SAN? TCP offloading?)

Operating System

- Physical or Virtual?
- Version and Edition
- Drive allocation
 - Also cluster size, partition alignment
 - Instant File Initialization
- Power Saving
- What Else Is Running On This Server?

SQL Server Configuration

- Version, Edition, Build, and Licence
- Part of a Cluster, Mirror or Availability Group?
- Trace Flags (1118, 1117, 3226, and possibly 4199)
- sp_configure
 - Cost Threshold for Parallelism? Appropriate to workload
 - Max Degree of Parallelism (MAXDOP)? Cores in NUMA node
 - Optimize for Ad-Hoc Queries? Turn it on
 - Max Server Memory? Appropriate to environment
 - Backup Compression? Turn it on
 - Lightweight Pooling and Priority Boost? Turn it off
- Agent Jobs
 - Maintenance tasks
- System Databases

User Databases

- Number of Databases
- Data and Transaction Log file allocation
 - File growth, free space, location, VLFs
- I/O Performance
 - Latency < 10ms
- Recovery Model
- Compatibility Level
 - SQL Server 2014
- Page Verify Option
 - Always be CHECKSUM
- Auto-Update Statistics
 - Always be on

Monitoring Tools

- WHAT IS THE BASELINE?
- Dynamic Management Views
 - Missing Indexes:
 - `sys.dm_db_missing_index_details`
 - `sys.dm_db_missing_index_groups`
 - Wait Statistics: `sys.dm_os_wait_stats`
 - Sessions: `sys.dm_exec_sessions`
- Free Tools:
 - `sp_WholsActive` (Adam Machanic)
 - Diagnostic Scripts (Glenn Berry, SQLskills.com)
 - `sp_Blitz` (BrentOzar.com)
 - `sp_BlitzIndex` (BrentOzar.com)
 - `sp_AskBrent` (BrentOzar.com)

Tuning Process

- tempdb performance
- Indexes
 - Unique clustered indexes on all tables
 - Non-clustered covering indexes (avoid key lookups)
 - Included columns (2005+)
 - Filtered indexes (2008+) (WHERE clause)
 - Indexed views (materialised views)
- Statistics
 - 20% plus 500 rows changed
- Query Tuning (the last resort)

Query Tuning

- Examine the longest running queries
 - Memory consumers
 - Sort, hash aggregate, hash join
 - Blocking operators
 - Sort, scalar aggregates, hash join, eager spool, any operation that can spill to disk (e.g. worktable)
 - Types of join (nested loop, merge, hash)
 - Key / Bookmark Lookups (try a covering index)
 - Cursors vs CROSS APPLY
 - Temp tables vs table variables
 - Collation
 - Implicit conversions
 - Craig Freedman: <http://goo.gl/H7ywSq>

Tips and Tricks

- Stored procedures beat ad-hoc queries any day
- Statistics are updated at 20% plus 500 rows
- Avoid query hints at all costs
- Parameterised dynamic T-SQL can be good
- Tipping point (scans vs seeks)
- Parameter sniffing can be bad
- OPTIMIZE FOR UNKNOWN could be worse
- It depends ...