

SQL Server 2014

Additional Performance Features

Bob Beauchemin
bobb@SQLskills.com



Outline

- Performance Features
- Reliability and Security Features
- Hybrid Cloud Features



2

© SQLskills. All rights reserved.
<http://www.SQLskills.com>



Delayed Durability

- **Transaction outcome is not returned to caller until after transaction log records are written**
- **This can slow down any transaction, especially**
 - In-memory OLTP (it's the biggest bottleneck)
 - When SQL Server transaction log on Windows Azure storage
- **Delayed durability is about loosening transaction guarantee**
 - Control return to user before transaction log committed
 - Can lose work on SQL Server or OS crash
- **Can be specified**
 - Per-database
 - On COMMIT TRANSACTION statement
 - On OLTP compiled stored procedure atomic block

Partitioned Statistics

- **Statistics on partitioned tables are table-level**
 - Recorded at a table-level
 - Updated at a table-level
- **Often, only one partition is the active partition**
 - Other partitions contain historical data
- **In SQL Server 2014**
 - Statistics can be created per-partition
 - INCREMENTAL = YES
 - Auto-created partitioned statistics can be created per-database
 - Individual partition's statistics can be updated
 - UPDATE STATISTICS RESAMPLE (ON PARTITIONS)
 - Update can recreate statistics as partitioned
 - INCREMENTAL = YES

Partitioned Stats Statements Affected

- **Partitioned statistics affect DDL**
 - CREATE STATISTICS
 - UPDATE STATISTICS
 - sp_createstats
 - CREATE INDEX
 - ALTER INDEX
 - ALTER DATABASE SET options
 - DATABASEPROPERTYEX
- **Metadata**
 - sys.databases
 - sys.stats

Partitioned Online Index Rebuild

- **Pre-2014**
 - Online index rebuild works only on all partitions
 - Long-running job to rebuild historical partitions
 - Offline index rebuild can rebuild one partition
 - Shorter job
 - "Current" partition is offline during entire job
- **In SQL Server 2014**
 - Online index rebuild works per-partition

Parallel SELECT INTO

- **No code changes needed**
 - Although insert volume must benefit from parallelism
- **Requires Compatibility Level 110 or greater**
- **No parallelism on INSERT... SELECT or INSERT... EXECUTE**

Performance Improvement in TEMPDB

- **SQL Server has always used "eager writes"**
 - To prevent flooding buffer pool with newly written pages
- **Eager writes no longer written to disk as quickly**
 - Allow pages to be stored in RAM, queried, table dropped
 - Never written to disk unless needed in this case
- **Allows speedup of some TEMPDB operations**
 - SELECT INTO #temp table
 - Table-valued parameters
 - Index operations with SORT IN TEMPDB

Low Priority Waiting

- **Used in index rebuild operations**
 - Short Sch-M must be taken during online index rebuild
 - Partition switching must take Sch-M on both partitions
 - Sch-M lock is incompatible with shared locks (all readers)
- **If rebuild uses WAIT_AT_LOW_PRIORITY**
 - Waiters for shared locks satisfied before low-priority Sch-M
- **Can specify duration to wait in minutes**
- **Can specify what happens after duration timer expires**
 - ABORT_AFTER_WAIT = **[NONE | SELF | BLOCKERS]**
 - SELF - Low priority tasks exits with error
 - BLOCKERS - Blocks killed (kill command)
 - NONE - Continue waiting
 - BLOCKERS option requires ALTER ANY CONNECTION privilege

Buffer Pool SSD Extension

- **Performance improved by allocating more memory to buffer pool**
 - Buy and install more memory
 - Use SSD storage to increase buffer pool
- **Since memory usually cheaper than SSD**
 - Use SSD extension on older machines where more memory not an option
 - When you run out of memory slots
- **SSD used as "paged overflow" to store clean pages only**
- **Configured with ALTER SERVER CONFIGURATION statement**
- **Some limitations**
 - Extension must be less than max server memory
 - To increase "SSD memory" – must drop and re-alter
 - To reduce "SSD memory" – must drop, start/restart SQL Server, then alter

Resource Governor Extensions

- **Resource Governor introduced in SQL Server 2008**
 - Meant to guard against runaway tasks that use all resources
 - Goal is more predictable throughput
 - Uses workload groups to define level of service
 - Workload groups are assigned to resource pools
 - Logins assign to workload group by classifier function
- **Pre-SQL Server 2014, you could configure (resource pools)**
 - CPU percent – MAX/MIN
 - Memory percent – MAX/MIN
 - Scheduler Affinity – NUMA or AUTO
 - Addition parameters on workload groups
- **In SQL Server 2014**
 - MAX_OUTSTANDING_IO_PER_VOLUME – global option for IO operations
 - IOPS_PER_VOLUME – MAX/MIN per resource group

New Diagnostics

- **Performance information available while query executes**
 - sys.dm_exec_query_profiles
 - Counts reported are per-plan operator, per-thread
 - Can be joined with performance monitor metadata queries
 - IO counts
 - Can compare actual vs. estimated row counts if "show actual plan" on
- **New features all are instrumented through**
 - Extended events – all
 - Performance Counters – new counters and extension to existing
 - Messages to errorlog – Low priority waiting
 - New DMVs and new columns in existing DMVs

Always On Enhancements

- **Availability Groups**
 - Maximum number of secondary replicas is increased from 4 to 8.
 - Readable secondary replicas now remain available when disconnected from primary
 - Easier setup and replica adding
 - More diagnostics through extended events
- **Failover cluster instances (FCI)**
 - Can use shared cluster volumes (Windows Server 2012+) as shared disks

Security

- **New granular permissions**
 - Makes principal of least privilege easier to implement
 - CONNECT ANY DATABASE Permission
 - IMPERSONATE ANY LOGIN Permission
 - SELECT ALL USER SECURABLES Permission
- **Encrypted backups without TDE**

SQL Server Configurations

- **There are four main ways SQL Server runs**
 - SQL Server on hardware
 - SQL Server virtualized on premises
 - SQL Server on cloud VM (IaaS)
 - Azure SQL Database (PaaS)
- **Four main differences**
 - Resource Capacity
 - Latency
 - Hardware
 - ASD and Virtual - problems cannot be solved with hardware
 - Instead solve problems with programming, scale out, queuing
 - Throttling

Hybrid Solutions

- **Hybrids introduced in SQL Server 2014**
 - SQL Server in Azure VM as Availability Group Readable Secondary
 - For on-premises Availability Group
 - SQL Server databases on Windows Azure Storage
 - SQL Server backup and restore to Windows Azure Storage
 - Includes encryption
 - Automated backup (a.k.a. Smart Backup)
 - Not available for Azure SQL Database (PaaS)
 - Backup to Azure storage tool
 - Backup to Azure or local storage
 - Backup encryption and compression supported
 - Backup available for pre-SQL Server 2014 versions
- **Windows Azure storage can be geo-replicated**
 - SQL Server databases on Windows Azure Storage does not support this

Management for SQL Server in Azure VM

- **Cloud adapter service**
 - Allows command to be sent from on-premises to SQL on Azure VM
 - Installed by default with SQL Server on Azure VM
 - Not installed with on-premises
- **Used with SSMS "Deploy SQL Database to Azure VM" wizard**
- **Additional SQL Server PowerShell cmdlets use this service**
 - Get-SqlInstance
 - Get/Set-SqlDatabase
 - Stop/Start-SqlInstance
 - Add/Remove-FirewallRule
 - Set-SqlNetworkConfiguration

Summary

- **Incremental changes**
 - Reduce contention during long jobs
 - Better granularity for partitioned tables
 - Allow resource governor to manage IO limits
 - Allow loosening transaction durability for speed
- **Enhancements for HADR**
- **Incremental security changes**
- **Enable hybrid on-premises and cloud solutions**