

Module 3: Wait Statistics

Jonathan Kehayias
Jonathan@SQLskills.com



Overview

- Introduction
 - Thread lifecycle
 - Waits, latches, spinlocks
 - DMVs
 - Common wait types
 - Demos
-
- Doctor, Doctor, my knee hurts...

Don't Assume Symptom = Root Cause

- **Performance troubleshooting is not an exact science**
 - The same symptoms can result from many root causes
- **For example, how many different things could cause I/O latencies?**
 - Overloaded/incorrectly-configured I/O subsystem
 - Synchronous I/O-subsystem mirroring
 - Buffer pool memory pressure
 - From plan cache bloat
 - From external Windows pressure
 - From an ad hoc query
 - From an inefficient query plan
 - Network latency
 - And more...

Interpreting the Data

- **Don't do 'knee-jerk' performance troubleshooting**
 - Work through the data to see what may be the root cause
 - You'll end up spending less time overall
- **Proficiency in using wait statistics data comes from:**
 - Retrieving the data correctly
 - Understanding what common wait types mean
 - Recognizing patterns
 - Avoiding inappropriate Internet advice
 - Practice!
- **Even better is to have a series of snapshots of wait statistics over time**
 - Allows identification of changes and the time of the change
 - Allows trending
- **Remember: not as valuable when SQL Server is running well**

What are Waits?

- The term 'wait' means that a thread running on a processor cannot proceed because a resource it requires is unavailable
 - It has to wait until the resource is available
- The resource being waited for is tracked by SQL Server
 - Each resource maps to a wait type
- Example resources that may be unavailable:
 - A lock (LCK_M_XX wait type)
 - A data file page in the buffer pool (PAGEIOLATCH_XX wait type)
 - Results from part of a parallel query (CXPACKET wait type)
 - A latch (LATCH_XX wait type)

Why are Resources Unavailable?

- Some other thread is holding the resource, or the 'resource' needs some process to occur (e.g. page read from disk)
- Examples:
 - For a LCK_M_XX wait, another thread holding an incompatible lock
 - For a PAGEIOLATCH_XX wait, the I/O subsystem needs to complete the I/O
 - For a CXPACKET wait, another thread needs to complete its portion of work
 - For a LATCH_XX wait, another thread holding an incompatible latch
- Resource waits are investigated using DMVs, performance counters, and other tools

Wait Statistics Analysis


- **Very powerful method to get initial direction on a problem**
 - Avoid flailing and investigating the wrong problem
 - Can also show problems that are not obvious
- **Original whitepaper from 2005:**
 - Performance Tuning Using Waits and Queues
 - Somewhat outdated advice, missing all waits added after 2005
- **New whitepaper from 2014:**
 - SQL Server Performance Tuning Using Wait Statistics: A Beginners Guide
- **Waits and latches library**
 - <http://www.SQLskills.com/help/waits>
- **Most commercial performance monitoring tools capture and show wait statistics**
 - Many free tools also do this, such as the popular Who Is Active tool
- **Various releases of SQL Server have provided wait statistics views**

Thread Scheduling

- **SQL Server performs its own thread scheduling**
 - Called non-preemptive scheduling
 - More efficient for SQL Server than relying on Windows scheduling
 - Performed by the SQLOS layer of the Storage Engine
- **Each processor core (whether logical or physical) has a scheduler**
 - A scheduler is responsible for managing the execution of work by threads
 - Schedulers exist for user threads and for internal operations
 - Use the sys.dm_os_schedulers DMV to view schedulers
- **When SQL Server has to call out to the OS, it must switch the calling thread to preemptive mode so the OS can interrupt it if necessary**

Components of a Scheduler


- All schedulers are composed of three 'parts'



- Threads transition around these parts until their work is complete

Schedulers in SQL Server

- One scheduler per logical or physical processor core
 - Plus some extra ones for internal tasks and the Dedicated Admin Connection
- For example, for a server with four physical processor cores, with hyper-threading enabled, there will be eight user schedulers




Thread States

- A thread can be in one of three states when being actively used as part of processing a query
- **RUNNING**
 - The thread is currently executing on the processor
- **SUSPENDED**
 - The thread is currently on the Waiter List waiting for a resource
- **RUNNABLE**
 - The thread is currently on the Runnable Queue waiting to execute on the processor
- Threads transition between these states until their work is complete

Transition: RUNNING to SUSPENDED

- A thread continues executing on the processor until it must wait for a resource to become available
 - The thread's state changes from RUNNING to SUSPENDED
 - The thread moves to the Waiter List
 - This process is called being 'suspended'




The Waiter List

- The Waiter List is an unordered list of threads that are suspended
- Any thread can be notified at any time that the resource it is waiting for is now available
 - Again, absolutely no ordering
- No limit to how long a thread remains on the waiter list
 - Although execution timeouts or lock timeouts may take effect
- There is no practical limit to how many threads can be on a scheduler's waiter list at any time
- The `sys.dm_os_waiting_tasks` DMV shows which threads are currently waiting and what they are waiting for


Special Case: Quantum Exhaustion

- If a thread does not need to wait for a resource, it will continue executing until its quantum is exhausted
 - Thread quantum is fixed at 4 milliseconds and cannot be changed
- If this occurs, thread moves to bottom of the Runnable Queue
 - The thread's state changes from RUNNING to RUNNABLE



Transition: SUSPENDED to RUNNABLE

- A thread continues to wait until it is told that the resource is available
 - The thread's state changes from SUSPENDED to RUNNABLE
 - The thread moves to the bottom of the Runnable Queue
 - This process is called being 'signaled'




The Runnable Queue


- The Runnable Queue is mostly a First-In-First-Out (FIFO) queue
 - Special case to avoid unfair scheduling in 2016
 - E.g. two threads where T1 can use entire 4ms, but T1 only 0.5ms
 - Pre-2016, T1 will get 8x the CPU time of T2
 - 2016: T1 and T2 will get roughly equal CPU time
 - See blog post at <http://bit.ly/2c6CElb> (capital-i)
 - Special case with Resource Governor High/Medium/Low priority workload groups
 - Never used, doesn't work properly
- Threads enter the queue at the bottom and progress to the top
- The thread at the top of the queue is the one that will execute on the processor when the processor becomes free
 - When the currently executing thread is suspended or exhausts its quantum
- The size of the Runnable Queue can be seen from the `runnable_tasks_count` column in `sys.dm_os_schedulers`

Transition: RUNNABLE to RUNNING

- The thread waits on the Runnable Queue until it reaches the top and the processor becomes available
 - The thread's state changes from RUNNABLE to RUNNING




Pulling It All Together



Wait Times Definition (1)

- **Total time spent waiting:**
 - Known as 'wait time'
 - Time spent transitioning from RUNNING, through SUSPENDED, to RUNNABLE, and back to RUNNING
- **Time spent waiting for the resource to be available:**
 - Known as 'resource wait time'
 - Time spent on the Waiter List with state SUSPENDED
- **Time spent waiting to get the processor after resource is available:**
 - Known as 'signal wait time'
 - Time spent on the Runnable Queue with state RUNNABLE
- **Wait time = resource wait time + signal wait time**

Wait Times Definition (2)



sys.dm_os_waiting_tasks DMV

- This DMV shows all threads that are currently suspended
- Think of it as the 'what is happening right now?' view of a server
- Most useful information this DMV provides:
 - Session ID and execution context ID of each thread
 - Wait type for each suspended thread
 - Description of the resource for some wait types
 - E.g. for locking wait types, the lock level and resource is described
 - Wait time for each suspended thread
 - If the thread is blocked by another thread, the ID of the blocking thread
 - Useful to find what's at the head of a blocking chain
 - Can show non-intuitive patterns
- Usually very first thing to run when approaching a 'slow' server
 - The data is more useful when joined with other DMV results

sys.dm_os_wait_stats DMV

- This DMV shows aggregated wait statistics for all wait types
 - Aggregated since the server started or the wait statistics were cleared
- Think of this as the 'what has happened in the past?' view of a server
- This DMV provides:
 - The name of each wait type
 - The number of times a wait has been for this wait type
 - The aggregate overall wait time for all waits for this wait type
 - The maximum wait time of any wait for this wait type
 - The aggregate signal wait time for all waits for this wait type
- Some math is required to make the results useful
 - Calculating the resource wait time and averages
- Also sys.dm_exec_session_wait_stats coming in SQL Server 2016
 - Cumulative (not per batch) wait stats for a single SPID

Using sys.dm_os_wait_stats on Azure

- There are a few differences using the sys.dm_os_wait_stats DMV on Azure SQL Database
 - sys.dm_os_wait_stats gives stats for the container the database is in
 - Use the new sys.dm_db_wait_stats for waits for just the database
 - Some system table references in the scripts I'll show you have to change
- Same thing for using sys.dm_io_virtual_file_stats
- Tim Radney explains how to get it all working here:
 - <http://bit.ly/1pHB7dn>

Filtering Benign Waits

- An extremely important point to bear in mind is that waits ALWAYS occur inside SQL Server
 - I.e. just because waits exist does not mean there is a perf problem
- Rather than looking at all waits, most useful is to focus on highly prevalent wait types
 - More processing of the sys.dm_os_wait_stats results is required
 - Common method is to show the top 95% of all waits by wait time
- Some wait types are almost always benign and can be safely ignored
 - Some have pathological, very rare cases where they can be problematic
- For example, the WAITFOR wait type
 - Only occurs when a WAITFOR DELAY statement is executed
 - When filtering the top 95% of waits by total wait time, not filtering out this wait can badly skew the results

Demo

Simple example: page latches and locks

Storing Wait Statistics

- **Capturing wait statistics information over time allows:**
 - Trending
 - Point-in-time analysis to see when a problem started to occur
- **Simple method:**
 - Use sys.dm_os_wait_stats demo script and add a GETDATE () call
 - Store the results in a table
 - Create SQL Agent job to capture the wait statistics every hour or so
 - Create another SQL Agent job to purge wait statistics older than a month

Methodology (1)

- Gather information about exactly when the performance problem arose and the user-visible characteristics of the problem
- Gather information about what changed before the problem arose
- Is the problem happening now or was it in the past?
- Examine any historical data sets from before the change and correlate through the time the problem arose
 - Look to see how the pattern of waits changes over time
- Examine the output from `sys.dm_os_waiting_tasks/dm_os_wait_stats`
 - What is happening on the server right now?
 - What has happened in the past?

Methodology (2)

- Look at the top 3-4 relevant waits
 - If LATCH_XX is present, examine the output from `sys.dm_os_latch_stats`
- Avoid the temptation to knee-jerk and equate symptoms with the root-cause
- Gather further info from relevant sources to pin-point problem
 - DMVs, query plans, performance counters, code analysis
 - Try one solution to see if that solves the problem
 - Repeat analysis, etc.

Using Extended Events

- **Extended Events were added in SQL Server 2008 to all Editions**
 - Very lightweight mechanism for tracing activity in SQL Server
 - However, mis-configuration can make Extended Events very expensive
- **When a wait starts and ends, the `sqlos.wait_info` event fires**
 - Captures similar information to `sys.dm_os_wait_stats`
 - Also the `sqlos.wait_completed` event added in SQL Server 2014
- **When a preemptive wait starts and ends, the `sqlos.wait_info_external` event fires**
 - Used when a thread is waiting for a call out to the OS and has to switch from non-preemptive to preemptive scheduling
- **Using the Extended Events system allows:**
 - Capturing of all wait types for a single operation
 - Monitoring for specific wait types occurring
 - Advanced analysis of SQL Server internals


What are Latches?

- **A latch is a synchronization mechanism between threads**
 - Many people equate latches with locks, but they are quite different
- **A latch protects access to an in-memory data structure**
 - Whereas a lock protects transactional consistency
- **Latches are lightweight and are held only for a short time**
 - Whereas a lock may be held until the end of a transaction
- **Latches cannot be controlled by SQL Server users**
 - Whereas locks can be controlled with hints and configuration options
- **Latches have a variety of modes, equating to the level of access to the in-memory data structure that is required**
 - E.g. an EX latch is required to change a data structure, and a SH latch is required to read most data structures
 - This is similar to the modes that a lock can have
- **SQL Server tracks latch wait times just like other waits**

Types of Latches

- **There are three types of latches:**
 - Latches waiting for data file pages to be read from disk into memory
 - Manifest as PAGEIOLATCH_XX waits
 - Latches for access to in-memory data file pages
 - Manifest as PAGELATCH_XX waits
 - Latches for access to all other data structures
 - Manifest as LATCH_XX waits
- **Examples of non-page latches:**
 - FGCB_ADD_REMOVE
 - ACCESS_METHODS_HOBT_VIRTUAL_ROOT

B-tree Page Split Example



From slide deck by
Thomas Kejser
(with permission)

Latch Contention

- **Just like with locks, latches can be a source of contention**
 - This means that what appears to be traditional blocking involving locks may actually be blocking involving latches
- **If one thread has a latch held exclusively then other threads must wait until that thread releases the exclusive latch**
- **This does not become a performance problem until there are many concurrent threads competing for access to the same latch**
 - As latches are only held for a short duration, a single thread waiting a very short time for another thread does not cause a problem
 - However, if hundreds of threads are waiting for a single thread, then that aggregates into a noticeable performance problem
- **Whitepaper on investigating latch contention:** <http://bit.ly/1uGha25>

Superlatches

- **When the Engine detects lots of SH latch requests on a buffer in the buffer pool, it will promote the latch to a super latch**
- **The latch is partitioned so there is one latch per scheduler, rather than one latch overall**
 - Reduces contention, as even SH latch access requires coordination
- **The superlatch will be demoted again if a series of EX latch requests are detected**
 - As an EX request for a superlatch requires EX latching each superlatch
- **May see this error:**
 - *Message Warning: Failure to calculate super-latch promotion threshold.*
 - Benign message as thresholds are recalculated every 60s by the lazy writer

sys.dm_os_latch_stats DMV

- This DMV shows aggregated wait statistics for all non-page latch classes
 - Aggregated since the server started or the latch statistics were cleared
- This DMV provides:
 - The name of each latch class
 - The number of times a wait has been for this latch class
 - The aggregate overall wait time for all waits for this latch class
 - The maximum wait time of any wait for this latch class
 - It does NOT list the latch modes being acquired
- Some math is required to make the results useful
 - Calculating the average times rather than the total times

Clearing Wait and Latch Statistics

- Clearing the aggregated wait statistics can be done at any time using the code below:
 - DBCC SQLPERF ('sys.dm_os_wait_stats', CLEAR);
- And for latch statistics:
 - DBCC SQLPERF ('sys.dm_os_latch_stats', CLEAR);
- Clearing the wait statistics allows the effect of a workload change to be measured against previous wait statistics
- Be careful if you are taking periodic snapshots of wait statistics as this will invalidate your series of snapshots
- When were they last cleared?
 - <http://www.sqlskills.com/blogs/erin/figuring-out-when-wait-stats-were-last-cleared/> (<http://bit.ly/OkXQVC>)

Using Extended Events

- **For very advanced troubleshooting there are events that allow tracking of latches**
 - `sqlserver.latch_suspend_begin`
 - `sqlserver.latch_suspend_end`
 - These are similar to the `sqlos.wait_info` and `sqlos.wait_info_external` events but have a lot more information about the latch itself

What are Spinlocks?

- **A spinlock is an even lighter-weight thread synchronization mechanism than a latch**
 - Used like a latch for data structure access control
- **Spinlocks are used when the data structure access will be for an extremely short time so the overhead of acquiring a latch is too much**
- **Examples of spinlocks:**
 - `FGCB_PRP_FILL`
 - `BUF_FREE_LIST`
- **Troubleshooting spinlocks usually requires very deep knowledge of SQL Server internals**
 - However, it is interesting and useful to know what spinlocks are

Spinlock Internals

- There is no waiting mechanism for spinlocks like there is for latches
- A thread tests the spinlock to see if it can be acquired
- If not, the thread sits in a loop checking whether it has the spinlock
 - When the thread cannot acquire the spinlock, this is called a 'collision'
 - When the thread loops, this is called spinning
 - Spins required after a collision do not count as more collisions
 - The number of collisions, and the number of spins are tracked
- After a certain number of spins, the thread yields
 - This is called a 'backoff'
 - Does not incur a wait, simply calls the Windows sleep() function
 - Can cause other threads to have high signal wait times
- SQL Server tracks all of this

Spinlock Contention

- When a large number of threads are contending for access to a single spinlock, this can lead to performance problems
- All these symptoms must be present for high CPU usage to potentially be from spinlock contention:
 - High and increasing spins (billion or more) and backoffs for a spinlock
 - High CPU usage with many connections to the server, OLTP workload
 - CPU usage, spins, and backoffs increasing much faster than the workload is increasing (possibly an exponential divergence)
- However, it is likely to NOT be spinlock contention so investigate other waits and latches first
 - Common for some spinlocks to have very high spins with an OLTP workload
- Troubleshooting spinlock contention is very advanced
- Whitepaper on investigating spinlock contention
<http://bit.ly/1rbZCuK>

Some Common Spinlocks

- **OPT_IDX_STATS**
 - This is the Query Optimizer updating the structures used for sys.dm_db_index_usage_stats and sys.dm_db_missing_index_stats
 - This could be from many concurrent updates to table with lots of indexes
 - See KB article 2003031 (<http://bit.ly/OgG7BE>)
- **LOCK_HASH**
 - This is the Lock Manager looking in the list of hash buckets for lock hash collisions
 - Consider smaller transactions, using NOLOCK, turning off page locks
- **LOGFLUSH_ACCESS and LOGFLUSHQ**
 - Involved with writing completed transaction log buffers to disk
 - Contention could be from very heavy load of very small transactions
 - See the WRITELOG wait for more details on how to reduce
- **Great post on spinlocks from Chris Adkins at <http://bit.ly/1HWr4XD>**

Transaction Log Example


- **Taking the transaction log and the logging system as an example, there are waits, latches, and spinlocks associated with it**
- **Waits:**
 - WRITELOG, LOGBUFFER, LOGGENERATION, LOGMGR, LOGMGR_FLUSH, LOGMGR_QUEUE, LOGMGR_RESERVE_APPEND
- **Latches:**
 - LOG_MANAGER, LOGBLOCK_GENERATIONS
- **Spinlocks:**
 - BUF_WRITE_LOG, LOGCACHE_ACCESS, LOGFLUSHQ, LOGLC, LOGLFM
- **From waits to latches to spinlocks, understanding the uses and troubleshooting becomes progressively harder and less likely to be required**
 - This is common across the SQL Server Engine

What's Relevant?

- **Just because there are waits, does not mean they are the problem**
 - Look for actionable items and filter out things like background tasks
 - Look at the demo code to see what I mean
- **Need to identify the top, relevant waits and then drill in**
- **Example:**
 - 1000 waits for LCK_M_S
 - Is it a problem?
 - No, if that was over 8 hours, there were 10 million locks acquired, and total wait time for the LCK_M_S locks was only 50s altogether
 - Yes, if each wait was for 50s

Top Wait Types

- **Survey results from 1700+ SQL Server instances across Internet**



- **Source:** <http://www.sqlskills.com/blogs/paul/common-wait-stats-24-hours/>
(<http://bit.ly/1n1m1lF> - capital i before the F)

PAGEIOLATCH_XX Wait

- **What does it mean:**
 - Waiting for a data file page to be read from disk into memory
 - Common modes to see are SH and EX
 - SH mode means the page will be read
 - EX mode means the page will be changed
- **Avoid knee-jerk response:**
 - Do not assume the I/O subsystem or I/O path is the problem
- **Further analysis:**
 - Determine which tables/indexes are being read
 - Take the page ID and follow steps in this post: <http://bit.ly/1F0ftog>
 - Analyze I/O subsystem latencies with sys.dm_io_virtual_file_stats and Avg Disk secs/Read performance counters
 - Correlate with CXPACKET waits, suggesting parallel scans
 - Examine query plans for parallel scans and implicit conversions
 - Investigate buffer pool memory pressure and Page Life Expectancy

PAGEIOLATCH_XX Wait Solutions

- Create appropriate nonclustered indexes to reduce scans
- Update statistics to allow efficient query plans
- Move the affected data files to faster I/O subsystem
- If data volume has simply increased, consider increasing memory
- Possibly In-Memory OLTP in SQL Server 2014+
- A quick band-aid could be to add more memory regardless to increase the buffer pool size
 - Cheap to do, provides temporary relief, maybe less risky than immediate code change

PAGELATCH_XX Wait

- **What does it mean:**
 - Waiting for access to an in-memory data file page
 - Common modes to see are SH and EX
 - SH mode means the page will be read
 - EX mode means the page will be changed
- **Avoid knee-jerk response:**
 - Do not confuse these with PAGEIOLATCH_XX waits
 - Does not mean add more memory or I/O capacity
- **Further analysis:**
 - Determine the page(s) that the thread is waiting for access to
 - Analyze the queries encountering this wait
 - Analyze the table and index structures involved

PAGELATCH_XX Wait Solutions

- **Classic tempdb contention**
 - Add more tempdb data files
 - Enable trace flag 1118
 - Reduce temp table usage
- **Excessive page splits occurring in indexes**
 - Change to a non-random index key
 - Avoid updating index records to be longer
 - Provision an index FILLFACTOR to alleviate page splits
- **Insertion point hotspot in a clustered index with an ever-increasing key**
 - Spread the insertion points in the index using a random or composite key, plus provision a FILLFACTOR to prevent page splits
 - Shard into multiple partitions/tables/databases/servers

LCK_M_XX Wait

- **What does it mean:**
 - A thread is waiting for a lock that cannot be granted because another thread is holding an incompatible lock
- **Avoid knee-jerk response:**
 - Do not assume that locking is the root cause
- **Further analysis:**
 - Follow the blocking chain using sys.dm_os_waiting_tasks to see what the lead blocking thread is waiting for
 - Use the blocked process report to capture information on queries waiting too long for locks
 - See Michael Swart's blog post for details about the various methods and further links (<http://bit.ly/ki3bYI>)
 - Are there any LCK_M_RS_XX locks? If so serializable isolation level was used


LCK_M_XX Wait Solutions

- **Lock escalation from a large update or table scan**
 - Possibly configure partition-level lock escalation, if applicable
 - Consider a different indexing strategy to use nonclustered index seeks
 - Consider breaking large updates into smaller transactions
 - Consider using snapshot isolation, a different isolation level, or locking hints
 - All the general strategies for alleviating blocking problems
- **Unnecessary locks for the data being accessed**
 - Consider using snapshot isolation, a different isolation level, or locking hints
- **Something preventing a transaction from releasing its locks quickly**
 - Determine what the bottleneck is and solve it appropriately
- **Serializable isolation level being used erroneously**
 - Distributed transactions, .Net TransactionScope default

WRITELOG Wait

- **What does it mean:**
 - Waiting for a transaction log block buffer to flush to disk
- **Avoid knee-jerk response:**
 - Do not assume that the transaction log file I/O system has a problem (although this is often the case)
 - Do not create additional transaction log files
- **Further analysis:**
 - Correlate WRITELOG wait time with I/O subsystem latency using `sys.dm_io_virtual_file_stats`
 - Look for LOGBUFFER waits, showing internal contention for log buffers
 - Look at average disk write queue length for log drive
 - If constantly 31/32 (111/112 on SQL 2012+) then the internal limit has been reached for outstanding transaction log writes for a single database
 - Look at average size and volume of transactions
 - Are all the log files for all databases on the same volume?

Transaction Log Flushes



WRITELOG Wait Solutions

- Move the log to a faster I/O subsystem
- Spread log files over multiple volumes
- Upgrade if hitting the 32 outstanding I/Os limit
- Increase size of transactions to reduce small log block flushes to disk
- Implement delayed durability in SQL Server 2014
 - Beware trade-off of potential data loss in the event of a crash
- Remove unused nonclustered indexes to reduce logging overhead from maintaining unused indexes during DML operations
- Change index keys or introduce FILLFACTORs to reduce page splits
- Investigate whether synchronous database mirroring/AGs/SAN replication is introducing delays
 - AG log send wait is tracked by HADR_SYNC_COMMIT
- Potentially split the workload over multiple databases or servers
- **Potentially isolate log writer thread** – see <http://bit.ly/1UaC4sl>



53


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

Demo

Slow transaction log


Parallel Threads Example

- As part of a query plan, you may see the  operator, for example
- This is a Repartition Streams operation
 - Uses producer and consumer threads, plus a control thread
- For a degree-of-parallelism = 4 operation, the threads would look like:



Threads in a Parallel Query

- A parallel query can have branches/zones that execute at same time




- Each branch can use up to degree-of-parallelism threads
- Always a single-task (thread) that runs the query, thread ID 0
- So total possible threads a query can use (reserved at start) is:
 - DOP x number of concurrent branches, plus the control thread
- Total possible threads a single operator can use is:
 - DOP x 2, plus the control thread
- Post on this from Paul White at <http://bit.ly/1y8wwAg>
- Post describing parallel thread placement on CPUs at <http://bit.ly/24KV2ej>


CXPACKET Wait Explanation

- **What does it mean:**
 - Parallel operations are taking place
 - Accumulating very fast implies skewed work distribution amongst threads or one of the workers is being blocked by something
- **Avoid knee-jerk response:**
 - Do not set server-wide MAXDOP to 1, disabling parallelism
- **Further analysis:**
 - Correlation with PAGEIOLATCH_SH waits? Implies large scans
 - Also may see with ACCESS_METHODS_DATASET_PARENT latch or ACCESS_METHODS_SCAN_RANGE_GENERATOR latch
 - Examine query plans of requests that are accruing CXPACKET waits to see if the query plans make sense for the query being performed
 - What is the wait type of the parallel thread that is taking too long? (i.e. the thread that does not have CXPACKET as its wait type)

CXPACKET Wait Example (1)



CXPACKET Wait Example (2)



CXPACKET Wait Solutions

- **Possible root-causes:**
 - Just parallelism occurring (e.g. <http://bit.ly/1kPymkZ> from CSS)
 - Table scans because of missing nonclustered indexes or incorrect query plan
 - Out-of-date statistics causing skewed work distribution
- **If there is actually a problem:**
 - Make sure statistics are up-to-date and appropriate indexes exist
 - Consider MAXDOP for the query, or just for the database (in 2016 onwards)
 - Consider MAXDOP = physical cores per NUMA node, or 8 for non-NUMA
 - Consider MAXDOP for the instance, but beware of mixed workloads
 - Consider using Resource Governor for MAX_DOP
 - Consider setting 'cost threshold for parallelism' higher than the query cost shown in the execution plan
 - CTFP means parallel plan if serial plan cost exceeds CTFP
 - Examine plan cache to find a good CTFP value to use
 - Jon's blog post at <http://bit.ly/1rTs9UX>

Demo

Parallelism

SOS_SCHEDULER_YIELD Wait

- **What does it mean:**
 - A thread exhausted its 4 millisecond quantum and voluntarily yielded
- **Avoid knee-jerk response:**
 - Do not assume that CPU pressure is the problem
 - High signal wait times show CPU pressure
 - Do not assume that spinlock contention is the problem
- **Further analysis:**
 - Examine query plans to see whether scans are occurring
 - Check if there is a very small or non-existent number of PAGEIOLATCH_XX waits occurring, which indicates that the workload is memory-resident
 - Look for long Runnable Queues
 - Capture SQL Server code call stacks to see where the waits are occurring
- **Note: these waits have zero resource wait time so regular methods of aggregating and prioritizing waits will miss them**
 - They do not appear in sys.dm_os_waiting_tasks

SOS_SCHEDULER_YIELD Solutions

- **Possible root-causes:**
 - SQL Server is executing code that can use a lot of CPU without having to wait for a resource (e.g. a large scan with few PAGEIOLATCH_SH waits)
 - Look also for long Runnable Queues, indicating CPU pressure
- **Solutions:**
 - For the quantum exhaustion case on slower processors, potentially enable hyper-threading to give more schedulers and more potential for concurrent work, especially for OLTP workloads

Using Extended Events to Examine Call Stacks

- **The only way to see exactly why SOS_SCHEDULER_YIELD waits are occurring is to examine SQL Server code call stacks**
- **Download the correct symbols**
 - See my blog post How to download a sqlservr.pdb symbol file (<http://bit.ly/Lc1cpi>)
- **Enable trace flag 3656 to allow call stack symbol resolution**
 - Also disable error log printing about dbghelp.dll version
- **Create an Extended Event session that:**
 - Captures sqllos.wait_info events for wait_type = 120 (or 124 for 2012+)
 - Captures the package0.callstack action
 - Uses the package0.histogram target
- **Run the workload and examine the captured call stacks**
- **Very advanced!**
 - See my blog post for a walk-through example (<http://bit.ly/vJXjA6>)

Demo


Using Extended Events to investigate SOS_SCHEDULER_YIELD waits

LATCH_XX Wait

- A non-page latch is the point of contention
- Further analysis:
 - Use sys.dm_os_latch_stats to investigate which latch(s) are experiencing high wait times
 - Correlate with other prevalent wait statistics
 - For example, CXPACKET waits with LATCH_EX waits where the prevalent latch class is ACCESS_METHODS_SCAN_RANGE_GENERATOR
- Possible root-causes and solutions:
- Depend on the latch class
 - These are not documented so use a search engine to look for any information
 - My blog category on latches has information (<http://bit.ly/Lc3J35>)

Top Latch Classes

- Survey results from 581 SQL Server instances across Internet



- Source: <http://www.sqlskills.com/blogs/paul/most-common-latch-classes-and-what-they-mean/> (<http://bit.ly/QXQB44>)

FGCB_ADD_REMOVE Latch

- Access to the File Group Control Block (FGCB) when adding, removing, shrinking, or growing files in the filegroup
- Further analysis:
 - Analyze the auto-growth settings of the file in all filegroups
 - Extended Events must be used to determine which database is involved
 - In SQL Server 2012+, use the latch_suspend_begin and latch_suspend_end events with latch class 48, and correlate to the database_data_file_size_change event using causality tracking
 - In SQL Server 2008 and 2008 R2, FGCB_ADD_REMOVE latch class number was 54
- Possible root-causes:
 - Auto-growth settings for a file are very low, requiring frequent growth, coupled with heavy, concurrent use of the filegroup
 - File sizes set too low for the rate of data entry into the database, requiring frequent auto-growth

Demo

Using Extended Events to determine which database is growing

DBCC_XX Latches

- **Examples:**
 - DBCC_MULTIOBJECTSCANNER (the most common to see)
 - DBCC_CHECK_AGGREGATE
 - DBCC_OBJECT_METADATA
- **Do not stop running consistency checks**
- **Further analysis: none necessary as these are all DBCC CHECKDB**
- **Possible root-causes:**
 - DBCC_MULTIOBJECTSCANNER latch was identified as a contention point and fixed in 2012 and under a trace flag in SQL Server 2008 R2
 - See Bob Ward's post (<http://bit.ly/yAFY7W>) and KB article 2634571
 - DBCC_OBJECT_METADATA is a bottleneck with computed column indexes
 - See my post at <http://bit.ly/YAzY6h>

PREEMPTIVE_OS_XX Waits

- **What does it mean:**
 - A thread has called out to the OS
 - Threads must switch to preemptive mode when doing so
 - Note that the thread status will be RUNNING instead of SUSPENDED
- **Further analysis:**
 - 194 PREEMPTIVE_OS_XX waits in SQL Server 2012/2014
 - These waits are very minimally and poorly documented
 - To determine what SQL Server is asking the OS to do, usually can remove the PREEMPTIVE_OS_ prefix and search in MSDN for the remainder of the wait type, as it will be the name of a Windows API
- **Possible root-causes and solutions:**
 - Depends on the wait type
 - For instance, increasing PREEMPTIVE_OS_CREATEFILE waits occur when using FILESTREAM on an incorrectly prepared NTFS volume

PREEMPTIVE_OS_WRITEFILEGATHER Wait

- **What does it mean:**
 - A thread is calling out to Windows to write to a file
- **Avoid knee-jerk response that I/O subsystem has a problem**
- **Further analysis:**
 - What database operations are under way? E.g. restore or file growth
- **Possible root-causes and solutions:**
 - Zeroing a large transaction log file during a restore or log file growth
 - Zeroing a large data file during restore or data file growth
 - Enable instant file initialization and set manage growth appropriately
 - Do not delete existing database files before performing a restore
 - Described in KB article 2091024 (<http://bit.ly/Lpz88m>)

PREEMPTIVE_OS_WAITFORSINGLEOBJECT Wait

- **What does it mean:**
 - Thread calling Windows to wait on state change of synchronization object
 - Commonly seen with ASYNC_NETWORK_IO wait
- **Further analysis:**
 - Follow instructions/solutions as for ASYNC_NETWORK_IO
 - Check whether transactional replication is running
- **Possible root-causes and solutions:**
 - As for ASYNC_NETWORK_IO (see next slide)
 - Could also be transactional replication Agent jobs (such as the Log Reader and Distribution Agent jobs)
 - See Joe Sack's blog post for more details (<http://bit.ly/AAtn5i>)

ASYNC_NETWORK_IO Wait

- **What does it mean:**
 - SQL Server is waiting for a client to acknowledge receipt of sent data
- **Avoid knee-jerk response:**
 - Do not assume that the problem is network latency
 - It is a network delay as far as SQL Server is concerned though
- **Further analysis:**
 - Analyze client application code
 - Analyze network latencies
- **Possible root-causes and solutions:**
 - Nearly always a poorly-coded application processing results one record at a time (RBAR = Row-By-Argonizing-Row), or slow application server
 - Very easy to demonstrate using a large query and SQL Server Management Studio running on the same machine as SQL Server
 - Otherwise look for network hardware issues, incorrect duplex settings, or TCP chimney offload problems (see <http://bit.ly/aPzoAx>)

OLEDB Wait

- **What does it mean:**
 - The OLE DB mechanism is being used
- **Avoid knee-jerk response:**
 - Do not assume that linked servers are being used
- **Further analysis:**
 - What are the queries doing that are waiting for OLEDB?
 - If linked servers are being used, what is causing the delay on the linked server?
- **Possible root-causes:**
 - DBCC CHECKDB and related commands use OLE DB internally
 - Many DMVs use OLE DB internally so it could be a third-party monitoring tool that is repeatedly calling DMVs (especially if they're very short waits)
 - Poor performance of a linked server

Demo

Other wait types

Miscellaneous Wait Types (1)

- **EXECSYNC**
 - Parallel threads exchanging information during execution
- **ASYNC_IO_COMPLETION**
 - Waiting for a non-data-file I/O to complete, for example when zeroing a transaction log file or writing to backup media (especially over a network)
- **IO_COMPLETION**
 - Similar, for example when reading from a transaction log file during transaction rollback or with transactional replication
- **WRITE_COMPLETION**
 - Waiting for a non-buffered I/O to complete, for example when creating certain allocation bitmap pages, such as PFS pages

Miscellaneous Wait Types (2)

- **THREADPOOL**
 - Waiting for a worker thread to become available, for example on a heavily-loaded system with a lot of parallel queries running
- **RESOURCE_SEMAPHORE**
 - Waiting for a query execution memory grant, for example for a sort
 - Usually indicates concurrent, memory-hungry queries
- **MSQL_XP**
 - Waiting for an extended stored procedure call to complete
- **LOGBUFFER**
 - Waiting for a log buffer to become available when flushing log contents

Real-World Example: Symptoms

- **Auto dealership hosting service**
 - Lots of auto dealers from across the US hosted on one site
 - Each auto dealer uploads inventory each day
 - One large Listing table storing all inventory for all dealers
 - One large Visitor table tracking clicks on web pages
 - No DBA
- **System had performance problem:**
 - User queries on inventory and prices regularly timed out
 - Inventory updates regularly timed out
 - Climbing CPU usage
 - Response time getting longer and longer
 - Car dealers pressuring hosting service for fixes
- **First step: analyze wait statistics...**

Real-World Example: Analysis

- **No historical data so gathered wait statistics data using the queries shown in earlier modules**
- **Both DMVs showed the same three wait types:**
 - CXPACKET wait = parallelism
 - PAGEIOLATCH_SH wait = reading data file pages from disk
 - WRITELOG wait = waiting for log writes, with average wait more than 20ms
- **Possible issues from just wait statistics**
 - Many queries doing parallel table scans of data that is not memory resident
 - I/O subsystem for the log file over-loaded and/or high number of log flushes
- **Investigated further using DMVs to analyze:**
 - Query plans
 - Index and table structures, index usage, fragmentation, and statistics
 - I/O subsystem latencies
- **Next step: determine root-causes...**

Real-World Example: Root-Causes

- **Both large tables had random GUID cluster keys**
 - High fragmentation in the clustered indexes leading to poor readahead
 - Lots of page-split transaction log activity during web page click tracking
- **Both large tables had more than 50 single-column nonclustered indexes**
 - Indexes not being used for seeks, resulting in table scans
 - Large amounts of nonclustered index maintenance from inserts, updates, deletes contributing to transaction log activity
- **Insufficient buffer pool memory for application workload data**
- **Poorly laid out I/O subsystem contributing to high latencies**
- **Poorly written code from using an ORM system**
- **Final step:** propose and implement solution

Real-World Example: Solution

- **Solution included:**
 - Increasing server memory and provisioning more appropriate I/O subsystem
 - Changing main tables to have bigint IDENTITY cluster keys
 - Removing useless nonclustered indexes
 - Analyzing ORM-generated code and query plans to determine appropriate nonclustered indexes
 - ORM system could not be removed for political reasons
 - Implementing index maintenance and periodic health checks
- **End result:** no performance problems and a happy client, with minimal investigation time

Key Takeaways

- Wait statistics are a key part of performance problem diagnosis
- Don't knee-jerk, follow the easy methodology
- Practice gathering and analyzing wait statistics using the DMVs
- Remember that waits always happen
- Don't get into latches and spinlocks unless absolutely necessary
 - Too many potential red herrings
- Know what the common waits mean and don't mean

Resources

- Waits/latches repository
 - <https://www.SQLskills.com/helps/waits>
- Whitepapers:
 - SQL Server Performance Tuning Using Wait Statistics: A Beginners Guide
 - <https://www.sqlskills.com/help/sql-server-performance-tuning-using-wait-statistics/>
 - Performance Tuning Using Waits and Queues
 - Diagnosing and Resolving Latch Contention on SQL Server
 - Diagnosing and Resolving Spinlock Contention on SQL Server
 - Gnarly links – see our whitepapers page at <http://bit.ly/19j0cOd> (zero then oh)
- Blog post categories
 - <https://www.sqlskills.com/blogs/paul/category/wait-stats/> and /latches/ and /spinlocks/
- Pluralsight: SQL Server: Performance Tuning Using Wait Statistics

Review

- Introduction
- Thread lifecycle
- Waits, latches, spinlocks
- DMVs
- Common wait types
- Demos

Questions?