



B75

Looking Beyond IMS for those Pesky Performance Problems

Dave Viguers

dviguers@us.ibm.com

IMS
Technical Conference

Sept. 27-30, 2004

Orlando, FL

Session Description

- Although IMS tends to take the blame for a lot of the performance problems there are many times when the problem lies outside of the control of IMS. Many times there may be some operating system service, security product, or other service or subsystem which is contributing to the performance difficulties being experienced. This session will look at some of those issues, how to identify them, and what might be done to keep your IMS running its best.

Topics

- Cyclical delays
- Syncpoint delays
- SMQ thruput
- Delays in Appl
- WLM
- GRS
- RACF
- CF

Occasional or cyclical delays

- Every few minutes during busy times there is a delay in processing
- You already looked at checkpoint
 - ▶ Ruled out because times don't match
- Checked for system messages
 - ▶ Nothing obvious - no OLDS switches
- CPU not pegged
 - ▶ At least not consistently at times of delays
- DASD seems ok
 - ▶ Logger data sets seem fine
- No GRS delays

Where else can you look?

- Go to SYS1.PARMLIB (D PARMLIB)
 - ▶ Member SMFPRMxx
 - Where xx is whatever is in IEASYSxx (SMF=)
 - Check the DDCONS keyword
 - ◆ Default is YES
 - ◆ If set to NO need to continue looking
 - ◆ Also check INTVAL value
 - ▶ Normally this is only a problem if using DB2 with IMS
 - IMS does not do much SMF recording
 - ▶ Can impact even without DB2 but less likely

Delays in syncpoint

- Longer tran elapsed times
- Higher region utilization
- Rule out the obvious
 - ▶ Logging
 - ▶ I/O bottlenecks
- Is DB2 data sharing in the picture?
 - ▶ Check the Immediate Write option
 - PH1 setting means updates written then
 - Yes means at actual update
 - This would show up in call detail in IMS mon or online monitor

Delays in syncpoint

- Real problem may be Group Buffer Pool too small
 - ▶ Switching to Y from PH1 may only move the problem
 - Look for occasional really long times
- One other point while we're talking DB2
 - ▶ Don't be concerned with long Class 1 times
 - This is 'normal' with IMS since
 - Thread reuse for better performance
 - Time may include the time between trans

Shared Queues thruput

- There are the IMS (CQS) things to consider
 - ▶ Structure checkpoints
 - ▶ Queue structure response times
 - ▶ Overflow processing
- CQS uses the MVS logger
 - ▶ Log stream is duplexed
 - Dataspace - best performance
 - ICF use may dictate not using for integrity
 - DASD - can severely limit thruput
 - But allows ICF use and PPRC of log data

Shared Queues thruput

- Duplexed CF structures
 - ▶ System managed
 - ▶ MVS Logger structures
 - ▶ CQS structures
 - ▶ Good for availability
 - ▶ Not so good for performance
 - Requires CF - CF communications
 - Can significantly impact CF times
 - Also reflected in higher CPU usage

Long delays in application

- Monitors show IMS call time ok
- DB2 calls also seem normal
- May show up as trans waiting for locks
 - ▶ Held by delayed / hung trans
- Need to find the trans holding locks
- Use IMSPA (or other monitor) call trace
 - ▶ Check individual call times
 - ▶ Check times between calls
 - Just shows up as time in application
 - Use the call trace to talk with appl folks
 - What is happening between calls?
 - ◆ APPC / MQ / ?

Long delays in application

■ Sample call trace

- ▶ Call elapsed on right - looks ok
- ▶ Look at Pgm Time between calls
 - ▶ Also not ISRT before long delay
 - ▶ Holding locks

Relative Time Elap	Pgm Time	Call	ST	PCB	Feedbck	Mod	IWT	DDname	IWT Elapsed	Breakdown of Call Time				Call		
Secs.Mil.Mic Sc.Mil.Mic	Sc.Mil.Mic	No.	PCBname	Func	Cd	Segname	Lvl	ULE	(# IWAITs)	Sc.Mil.Mic	Pct	CPY	Pct	DLA	Pct	IWT
0.007	3.238	0.037	10	HISTDB	GN	GE MYROOT	01				.00%	100.00%		.00%		
1.288	4.470	1.225	11	HISTDB	ISRT	HIST	02		1	1.112	.00%	100.00%		86.34%		
1.173	4.603								FPD DEDB HISTDD	1.112						
1.173	24.009.143	24.003.385	12	HISTDB	GU	DEP02	02		1	0.958	.00%	100.00%		81.67%		
0.009	24.009.313								FPD DEDB HISTDD	0.958						
0.009	24.010.332	0.016	13	HISTDB	GN	GE MYROOT	01				.00%	100.00%		.00%		
				24.010.374		0.033	14	HISTDB	GHU	DEP02	02					
	.00%	100.00%	.00%													0.005

Workload Manager

- Goal mode is only option in 1.3 and above
- If WLM settings incorrect
 - ▶ Similar to poor compat mode settings
 - ▶ Typically shows up as:
 - Delays in IMS trans during high CPU
 - If 100% then delays might be inevitable
 - Trans waiting for locks
 - Possibly held by lower priority work
- So what is correct?
 - ▶ It depends - of course

Workload Manager

- Region priority is adjusted to meet goal
 - ▶ Or at least try to
- How does IMS schedule trans into regions?
 - ▶ By class for FF (specific regions for IFP)
 - ▶ How many different trans use the same class?
- Unless specific requirements
 - ▶ Don't try to manage below class (eg: tran)
 - ▶ WLM might get confused
 - Too much work for no benefit

Workload Manager

- Possible settings
 - ▶ IMS & DB2 address spaces
 - Importance 1 (maybe CPU critical)
 - ▶ MPP's
 - Importance 2 (maybe CPU critical)
 - ▶ BMP's
 - Importance 3 - 5
 - ▶ If DBCTL then AOR's similar to MPP's
 - ▶ IRLM
 - SYSSTC

GRS / RECONS

- RESERVE / ENQ
 - ▶ Exclude or Convert - but do at least one
 - Exclude if
 - possibility of 'other' systems outside this sysplex with access
 - Consider using SYNCRES option (GRSCNFxx)
 - ◆ eliminates DBRC OBTAIN to force I/O
 - Convert if
 - GRS star
 - Other data sets on same volume
 - ▶ Resource - DSPURI01 (SYSTEMS level)
 - DSPURI02 at SYSTEM level (not systemsg)

GRS / RECONS

- GRS uses system resources
 - ▶ CF structure(s)
 - Delays in access can affect entire sysplex
 - Monitor with RMF
- Don't forget basic RECON placement guidelines
 - ▶ Separate volumes for each RECON
 - ▶ Each one with own catalog
- DBRC priority - WLM Importance 1

RACF

- We all know that security just gets in our way
 - ▶ Unfortunately we can't usually control
- RACF (or whatever security manager)
 - ▶ Uses system resources
 - CF structures
 - DASD
 - Data spaces
 - Storage
 - Getmains/Freemains
- Use caching to avoid I/O
- RACF enhancements
 - ▶ Selective refresh

RACF

- Different IMS implementations
 - ▶ Depends on source of request
 - Terminal
 - Sign on/off infrequently
 - ◆ Hopefully at least
 - ◆ ETO ASOT / ALOT ?
 - APPC
 - Has been source of high CPU & delays
 - Recent IMS FIX to reduce overhead
 - ◆ Fewer calls to RACF
 - ◆ No reduction in function
 - ◆ Uses RACF OBJECT support
 - ▶ Use DFSBSEX0 to reduce calls

Coupling Facility

- Many types of configurations
 - ▶ CF engine speed
 - Should always be at least the same as host
 - ▶ # of engines
 - 1 usually enough, 2 better to handle spikes
 - ▶ Dedicated or Shared processors
 - Dedicated best for production
 - ▶ Dynamic CF dispatching
 - OK for test, not for production
 - ▶ Utilization
 - Keep under 50%

Coupling Facility

- Link types
 - Peer mode ICB4 on z990's
 - more buffers, increased data rates
- Duplexing
 - Software (eg: VSO)
 - 2x single structure
 - System managed
 - 3-5x single structure
 - Need to weigh value versus 'cost'
 - increased cpu / response time
 - faster recovery (compared to rebuild or log)

Summary

- Just some of the recurring issues
- Lots of other possibilities
- IMS uses lots of system services
 - ▶ Reserve
 - ▶ Enqueue
 - ▶ Getmain
 - ▶ Page fix
 - ▶ etc.
- Make sure WLM goals are appropriate
- Get documentation