Looking at MQ on z/OS performance from a high level

Dirk Marski, IBM
MQ for z/OS Support

Why are you here?

- An “MQ performance problem” is usually a manifestation of a problem occurring elsewhere
- What evidence is there?
- Audience
  - MQ Administrators, MQ sysprogs, z/OS Sysprogs familiar with
    - MQ
    - Buffer pools
    - Page sets
    - CF
    - Structures
    - SMDS
    - Logs

© 2015 IBM Corporation

Notes

- Story about PMR
  - Response time from 200 ms to 1000 ms
  - Not CICS, Network so must be MQ
  - Please fix the problem

© 2015 IBM Corporation

Why are you here?

- Do I need to worry about MQ?
- Not meeting Service Level Agreements (SLA)
  - Response time criteria
  - Batch job duration or time to empty queue
- You've been told to reduce the z/OS resources
  - You are the sysprog - but it is the applications using CPU!
- Will the work scale in the future?
  - Consolidation of LPARs and queue managers
- You need evidence for your management that the problem is not MQ

© 2015 IBM Corporation

Agenda

- Typical MQ environments
- Who are the people
- What do I need to do now
- What can cause high CPU
- What can cause long response time?
- Use new technology
- Should I tune it
- Will it scale
- Useful commands

© 2015 IBM Corporation

Typical MQ environments

- Online processing
  - Messages processed within seconds
  - End to end response time under 1 second
  - Queues typically have 10's of messages
  - Performance critical → worry about tuning
  - Work can be queued to balance workload (10s of ms)
- Batch processing
  - Message last for hours
  - Deep queues (millions of messages)
  - Performance is not critical
- Most customers have both! Need to isolate work

© 2015 IBM Corporation

What factors are there?

- Resources
  - CPU
  - Disk I/O (Response time / Throughput)
  - Storage → never allow paging to occur!
  - CF
  - Network
- Usage
  - Application design
- Metrics
  - Response time
  - Cost

© 2015 IBM Corporation

Agenda

- Typical MQ environments
- Who are the people
- What do I need to do now
- What can cause high CPU
- What can cause long response time?
- Use new technology
- Should I tune it
- Will it scale
- Useful commands

© 2015 IBM Corporation

Who are the people

Business/Strategist

Dev

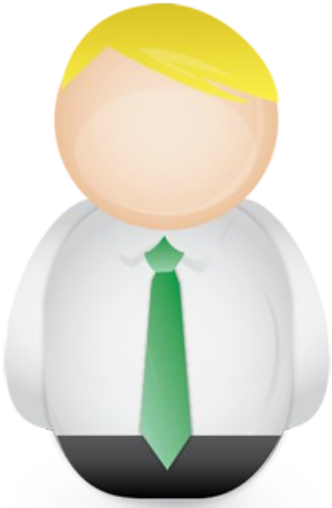

Infrastructure

© 2015 IBM Corporation

Notes

- 10 years ago infrastructure in charge – would not let application into production unless had capacity planning one
- Now with agile – want application written this week – in production next week
  - No idea of message size, or throughtput
  - Hard to plan for this
- Application developers in charge
- Any application bugs hard to fix
  - All Cobol programmers have left
  - New applications written in Java for mobile phone!

© 2015 IBM Corporation

What do I need to do now

- Need to be able to compare a good day with a bad day
- Capture data about the system.
  - Typical queue depths at peak time
  - How long messages were on xmit queue
  - Buffer pool usage
  - z/OS view of MQ
  - CF usage
    - Which LPARs were using the structures and response times
- Has the problem just started or happened over time?
- Has anything changed in the environment?
- Practice collecting the data
  - Everything will be in place and tested in case of a problem

© 2015 IBM Corporation

Agenda

- Typical MQ environments
- Who are the people
- What do I need to do now
- What can cause high CPU
- What can cause long response time?
- Use new technology
- Should I tune it
- Will it scale
- Useful commands

© 2015 IBM Corporation

Where is the CPU used?

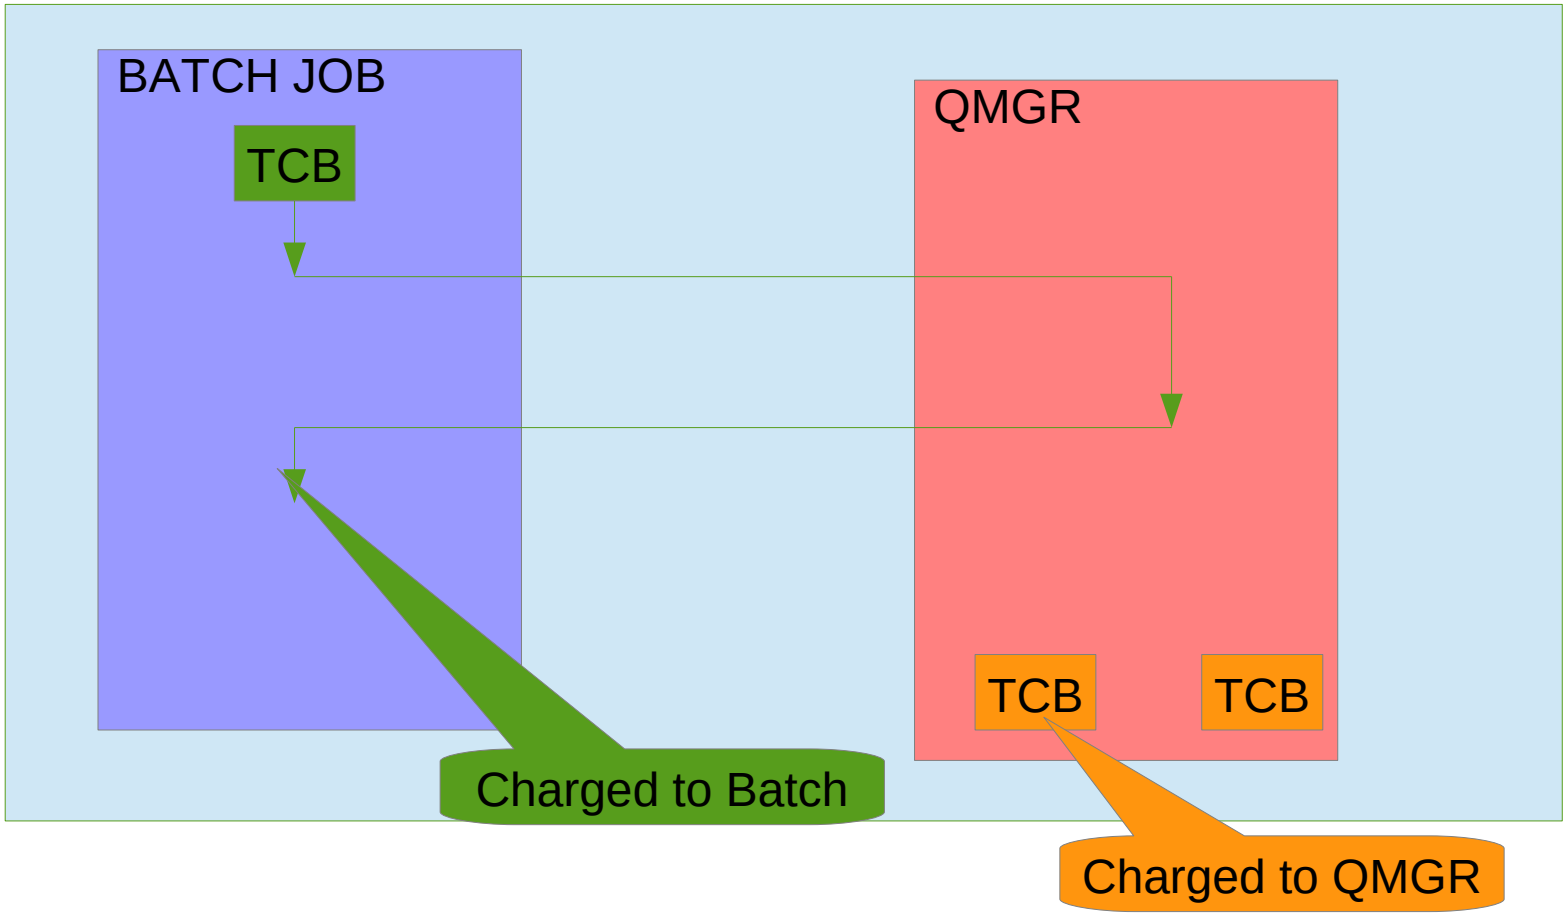

© 2015 IBM Corporation

How do I know if I have high CPU?

- MVS Reports CPU used by Jobs
- CICS/IMS statistics
  - Reports CPU used by transaction
- MQ accounting class
- Programs like Strobe(R) and Application Performance Analyser
  - Both of these sample the address space
    - Report which areas of code are hot
    - Reports DB2 and MQ usage

© 2015 IBM Corporation

Notes

- Strobe / Application Performance Analyser are near real time.
- Application Performance Analyser has little value for the QMGR (or subsystems) as it works best for applications (jobs)

© 2015 IBM Corporation

What can cause high CPU

- Queues not indexed
  - –Many messages scanned
  - –Msgid/Correlid not cleared – so getting the wrong message
  - –CSQI004I MQ7A CSQIMGE3 Consider indexing CP0000 by MSGID for BATCH connection PAICEP7A, 36345 messages skipped
- Too many servers
- Still using DB2 blobs
- Dont understand getwait – MQGET returns when message arrives or times out

© 2015 IBM Corporation

Application design problems

- Trigger every - with medium to high throughput
  - Additional messages produced, cost of starting and ending transaction
  - Long running transaction is better in most cases
  - May need logic to monitor and start more instances
- Multiple hops eg broker
  - Many stages putting to MQ and getting from MQ
  - Convert 'tree' to XML. MQPUT, MQCOMMIT, MQGET, XML → Tree
- Use of dynamic shared queue
  - Often better to have a common reply-to queue

© 2015 IBM Corporation

High CPU – Application coding

- Not clearing msgid/correlid
  - Search of queue
  - Gets worse with increased depth
- Customer with 12 broker msg flows. Reduced them to 1 flow doing all steps. CPU reduced by 50%, transaction speed improved x10
- Polling the queue
  - Use get wait with long interval
  - Some people wrongly code 0.1 second and loop around
  - Opportunity to use get signal to post an ECB

© 2015 IBM Corporation

What can cause for high CPU in the chinit.

- Clients doing connect, one get/put, disconnect
  - Connect/Disconnect is expensive
  - Consider front end concentrator
  - Can you use connection pooling?
  - Data conversion? → Move to remote queue manager?
- Channel batch size too small
  - Use batch size > 50 and > default batchlim
- Some TLS ciphers are more expensive
- V8 Suppress message – channel start/stop
- Client should use listener with INDISP QMGR instead of GROUP
  - Eliminate DB2 requests
- What are Channel exits doing?
  - Use DIS CHS EXITTIME

© 2015 IBM Corporation

Notes

- Customers think batch size 1 speeds things up
- TLS cipher – see MP16
- Cheaper not to produce the channel start/stop (V8) rather than produce and throw away (V7)
- Note that for client connections, in order to use extended transactional client, the server needs to be V701 or higher.

© 2015 IBM Corporation

Agenda

- Typical MQ environments
- Who are the people
- What do I need to do now
- What can cause high CPU
- What can cause long response time?
- Use new technology
- Should I tune it
- Will it scale
- Useful commands

© 2015 IBM Corporation

Where is time spent?

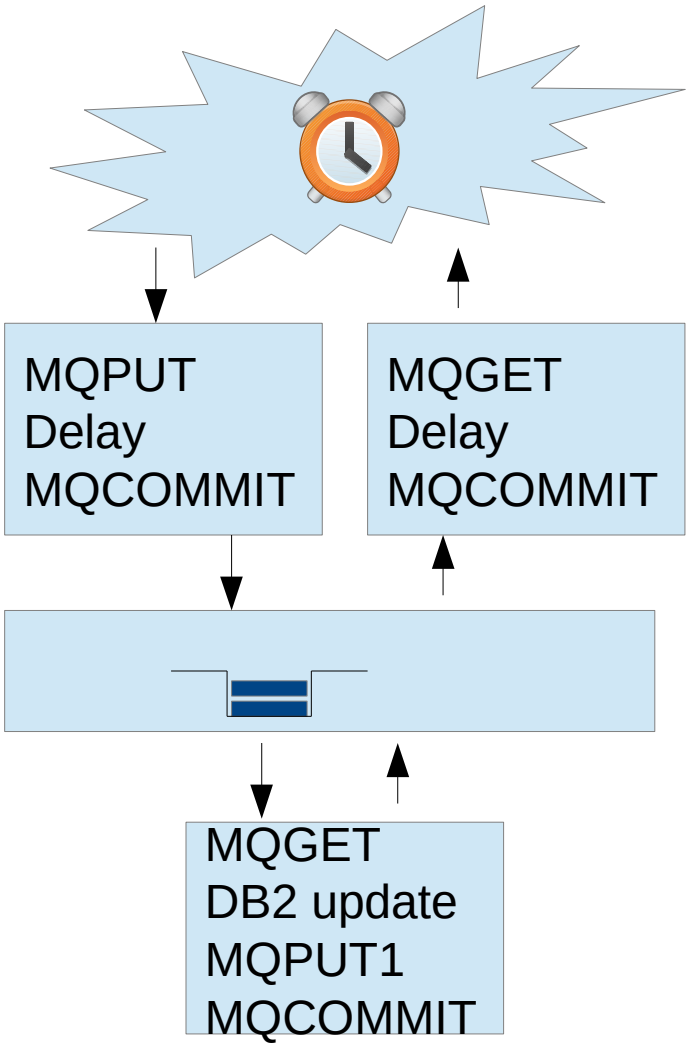

- Delay before/after MQ
- Messages queued up
  - Server applications not processing them
    - Is application dispatched?
    - Delays due to database?
- MQPUT+MQGET longer
- MQCommit taking longer
  - 2 Phase
  - Disk I/O
- Channel is an application
  - And has network time

© 2015 IBM Corporation

Where do you start?

- If applications or MQ short of CPU (big problem)
  - Fix this before doing anything else – Talk to z/OS sysprog
- What can increase time of MQPUTs and MQGETS
  - Use MQ accounting class(3)
  - Check Buffer Pools and CF
- What can cause increase in MQCOMMIT
  - Use MQ accounting class(3)
  - Use RMF to look at log response time.
- Network problems (big problem)

© 2015 IBM Corporation

Workload isolation

- Check 'batch' workload is isolated from 'online' workload
  - Different page set
  - Different structures
    - Different CF?
  - Different buffer pools
  - Different channels
- Consider a dedicated QMGR for large messages
- You need to know the MQ Objects used by each application type

© 2015 IBM Corporation

Deep server queue: What can delay the (server) application

- Is z/OS ok? → Ask sysprog; short on CPU, paging , WLM
- Are there enough application instances?
  - Is this just a peak workload?
- Is the problem in CICS/WAS/IMS? - Check statistics
  - Transactions slow to run
  - Is the application busy?
    - Doing database updates?
  - CICS/IMS limiting number of instances
    - So only 2 transaction instances of “PAYROLL” can run at a time
- Trigger every – slow to start

© 2015 IBM Corporation

What can make MQPUT and MQGET take longer?

- These are normally small impacts
- High cost MQGETs
- Bigger messages? - Has the application changed?
- Using DB2 blobs? (Use DB2 monitors)
- Using SMDS instead of CF. MQ Display command
- Now doing page set I/O - MQ Display usage command
  - Buffer pool > 85% full starts moving pages to disk
  - Buffer pool > 95% full - every page does disk I/O
  - Buffer pool with mixed message types
    - Isolate key business transactions
- Do you need to use persistent messages?

© 2015 IBM Corporation

What can make MQPUT and MQGET take longer – CF?

- Impact usually small
- Use RMF
- Response time depends on location and duplex (10 to 1000 uSecs)
- Message arrives on CF1. LPAR1 likely to get message
  - Overload LPAR1? OTMA
- Response time depends on load on CF
  - CF CPU < 50% busy
  - CF Channels
- Other structures can impact MQ eg CICS and DB2

© 2015 IBM Corporation

What can make commits take longer

- This is normally small impact
- MQ is busier every week
- Longer I/O request (1ms to 5 ms) – use RMF or MQ stats
  - Due to increased load
  - Due to slower DASD
  - Mirrored DASD
- 2 Phase commit
  - delays in DB2 commit processing
- CF requests may take longer

© 2015 IBM Corporation

Notes

- Typical log I/O under 1-2 ms
- Mirrored DASD takes longer has to go to remote site
  - The further – the longer it takes!
- 2 phase is 3 log forces
  - So adding an MQPUT into the transaction makes longer transaction
  - Eg was 5ms to DB2 before (with mirrored DASD) now 10 = 5ms from DB2 + 2 from MQ

© 2015 IBM Corporation

What can delay a channel?

- Network can have major impact
  - Use TCPIP Ping and netstat (see later)
  - TCP packet trace
  - MQ Display commands
- Put retry
  - Queue full at the remote end – retry after a period – messages queued
- Interference
  - Same channel used for different applications
  - Application 1 has more traffic , impacts application 2
  - Application 1 get queue full – message for application 2 delayed.

© 2015 IBM Corporation

Check the channel

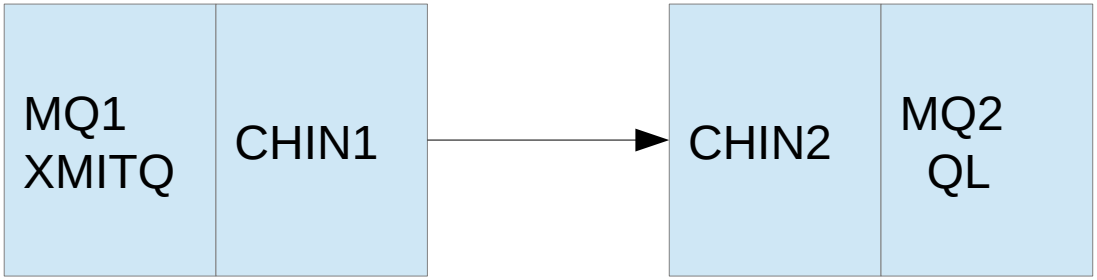

- MQPA dis chs(...)xqtime nettime msgs xbatchsz
  - MSGS(2000)
    - So you can calculate message rate
    - Issue DIS CHS wait.. DIS CHS
  - XQTIME(65,53) uSeconds → high investigate?
    - How long was the message on the XMITQ for ? < 100 is good
  - XBATCHSZ(1,1)
    - low means channel was waiting for messages to send
    - If XBATCHSZ = BATCHSIZE always messages to send. Problem ?
  - NETTIME(422,399) uSeconds → if high then investigate?
    - Might be a network, channel or remote end issue

© 2015 IBM Corporation

What is Nettime?

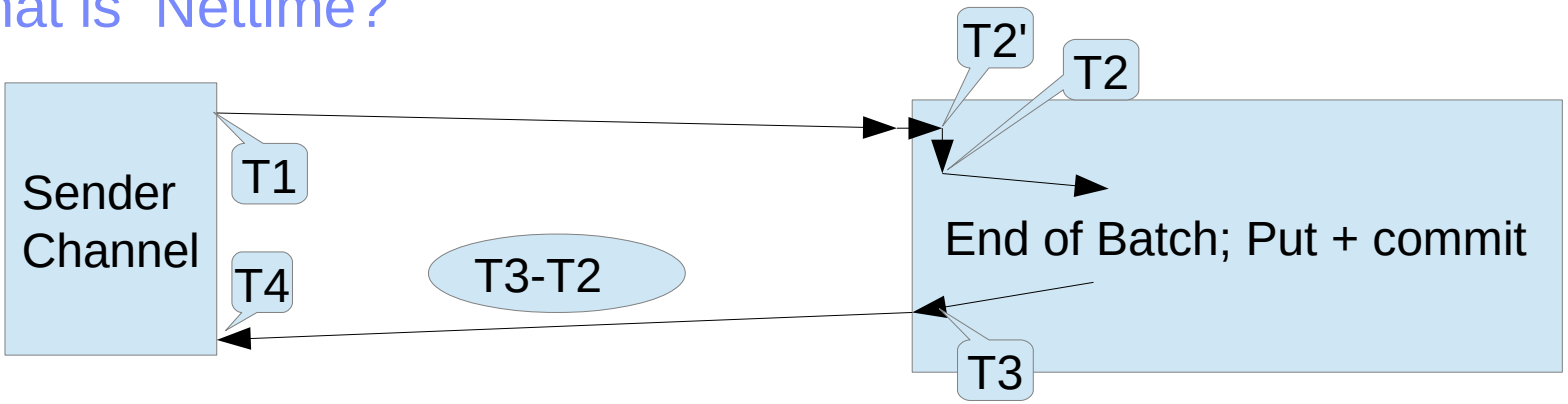

- Time T1 and T3 are taken immediately before send()
- Time T2 and T4 are taken immediately after receive()
- Time T2' is when the TCP/IP buffer holds the data
- T2' – T2 is the time between data arriving and the channel being ready to receive it
  - busy/blocked channel means long delay
  - Network may be ok for high NETTIME
- NETTIME short. Network is OK and no delays (T2'-T2 small)

© 2015 IBM Corporation

What is Nettime? - Queue retry

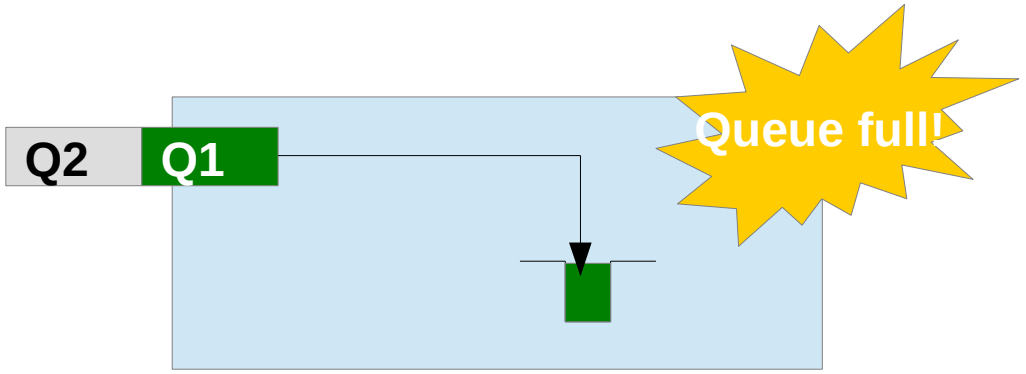

- The green queue has filled up
  - Put to the queue waits for period
    - MRRTY(10) this many times
    - MRTMR(1000) this many milliseconds between retries
- Q2 delayed!

© 2015 IBM Corporation

Agenda

- Typical MQ environments
- Who are the people
- What do I need to do now
- What can cause high CPU
- What can cause long response time?
- Use new technology
- Should I tune it
- Will it scale
- Useful commands

© 2015 IBM Corporation

Technology answers to old applications

- Often applications owners do not want to change applications
  - Want technology solution
  - V8 64 bit buffer pools
  - Bigger page sets
- Increased workload → running out of log space
  - add more logs or make them bigger
- Is changing the hardware better than you spending 1 week on it?
  - Application tuning
  - Encryption hardware
  - Stripe data sets
- 1990's applications filling up buffer pools?
  - Go to V8 and use very large buffer pools (and real storage)

© 2015 IBM Corporation

Notes

- 1990's applications filling up buffer pools?
  - When queues had 10 messages
  - Now queues have 1 Million messages
  - They have no COBOL programmers
  - Go to V8 and use very large buffer pools (and real storage)

© 2015 IBM Corporation

How much effort should I spend?

- High activity - business critical
- Low activity - low importance
- Understand the business applications
  - Impact on the business
  - What business applications are important to the business - high focus from management
  - Response time critical?
- Is good enough, good enough?
  - If it is good enough today, and good enough for tomorrow
- How much of transaction is MQ – 2% of cost or 90% of cost?

© 2015 IBM Corporation

Should I tune it? Manage the risks

- Tuning is not always beneficial
  - Merge two channels to save CPU
  - Causing longer response times
  - Combination of very big messages and short message
  - Availability
- You may accept higher costs for improved availability
  - Shared queue, multiple LPARs

© 2015 IBM Corporation

Your mileage may vary

- Different implementations have different characteristics
  - Same application in concept, different implementation, DB2 etc, message size
  - Resources available – storage for buffer pools
  - Some customers have spare capacity - some have none
  - Response time from DASD and CF
  - Throughput rates
- General tuning common to all customers
  - Customer specific tuning
- Some test at max volume + 25%, some 'Implement and pray'

© 2015 IBM Corporation

Agenda

- Typical MQ environments
- Who are the people
- What do I need to do now
- What can cause high CPU
- What can cause long response time?
- Use new technology
- Should I tune it
- Will it scale
- Useful commands

© 2015 IBM Corporation

Will it scale?

- Combining queue managers
- Scenario 1
  - High throughput, messages short lived
- Scenario 2
  - Up to 1 Million messages on a queue

© 2015 IBM Corporation

Will it scale? High throughput – short lived messages (1)

- Two queue managers
  - Logging 60 MB/Second
  - Max depth on queue 10
  - Buffer pool max 2% full
  - 5 channels XBATCHSZ 35

- Merged
  - Check log stats
    - 120 MB/Second
    - Logs fill more often
    - More checkpoints/hour
  - Assume same servers
  - Max depth stay low
  - Buffer pool 2% full
    - More pageset set I/O due to checkpoint frequencey
  - Increase channel batch size?
  - Use 10 channels?
  - Increased storage usage
    - More connections

© 2015 IBM Corporation

Will it scale? High throughput – short lived messages(2)

- Two queue managers
  - CF
  - One path to CF from each LPAR
- Merged
  - Same sized CF
  - Same CPU used
  - One path to CF
    - Higher Utilisation
  - May offload to SMDS
  -
  -
  - CPU impact – more work

© 2015 IBM Corporation

Will it scale? Deep queues

- Two queue managers
  - Logging 60 MB/Second
  - Max depth on queue 1 Million
  - Buffer pool max 95% full
- Merged
  - Check log stats
    - 120 MB/Second
    - Logs fill more often
    - More checkpoints/hour
  - Expect 2 million messages
    - Bigger page sets?
    - More data, more checkpoints, more I/O
    - Extended format?
    - Stripe page sets
    - Need to measure the time to drain queues
    - Use multiple page sets and 2 Buffer pools
- Need to test

© 2015 IBM Corporation

Agenda

- Typical MQ environments
- Who are the people
- What do I need to do now
- What can cause high CPU
- What can cause long response time?
- Use new technology
- Should I tune it
- Will it scale
- Useful commands

© 2015 IBM Corporation

Where do I start to look?

- Where do I start to look?
  - SMF data – see deep dive
  - DB2 monitors for BLOB activity
- Use online z/OS tools like RMFWDM
  - Reports every minute
  - Shows you delays due to CPU, DASD, CF
- Offline tools like RMF
  - Period 30 minutes – may not show spikes
- MQ monitoring eg Tivoli
- SupportPac MP1B – free!
  - MQCMD issues display command and stores output in csv file

© 2015 IBM Corporation

What's happening to my queue?

- DIS Q(xx) curdepth
  - Tells you depth now
- RESET QSTATS(SERVER)
  - RESETINT(14)
  - HIQDEPTH(6)
  - MSGSIN(3482)
  - MSGSOUT(3482)
  - Processing rate 3482/14 = 249 messages a second
- Performance events
- Use monitoring product or MQCMD in SupportPac MP1B
  - Issues command periodically and stores output in CSV file

© 2015 IBM Corporation

What's happening to my Channel?

- Use Display Chstatus
  - MSGS(2000)
    - So you can calculate message rate
  - XQTIME(65,53) → high investigate?
    - How long was the message on the XMITQ for ?
  - XBATCHSZ(1,1) → low means channel was waiting for messages to send
  - NETTIME(422,399) → high investigate?
- Use SMF (V8)

© 2015 IBM Corporation

What's happening to my Channel?

- Ping
  - //RUN1GB EXEC PGM=IKJEFT01,REGION=0M
  - //SYSTSPRT DD SYSOUT=*
  - //SYSTSIN   DD *
  - ping winmvsac (count 1000 verbose length 32768
  - /*

  – Ping statistics for winlnxn7.hursley.ibm.com (9.20.4.136)
    - Packets: Sent=1000, Received=1000, Lost=0 (0% loss)
    - Approximate round trip times in milliseconds:
    - Minimum=1.03 ms, Maximum=10.19 ms, Average=1.13 ms, StdDev=0.50 ms

  – Seen customers with over 80 milliseconds across pacific.

  – Where is the remote box?

  – Measure your round trip time

-

© 2015 IBM Corporation

What's happening to my Buffer pool

- DIS USAGE
  - Tells you now
  - CSQI065I MQPC Buffer pool attributes ... 359

Buffer pool	Available buffers	Stealable buffers	Stealable percentage	Page class	Location
0	10000	9981	99	4KB	BELOW
1	10000	9979	99	4KB	BELOW
2	10000	9999	99	4KB	BELOW
3	10000	9994	99	4KB	BELOW
4	50000	49962	99	4KB	BELOW

  - 99% free !

© 2015 IBM Corporation

What is my logging rate?

- Log message
  - CSQJ002I MQPD END OF ACTIVE LOG DATA SET DSNAME=...
    STARTRBA=0000000E72A1E000 ENDRBA=0000000F32B9DFFF
  - 0F32B9DFFF−0E72A1E000 = 3074MB
  - Every 2 minutes
  - Logging rate = 25 MB/Second
- Monitor how often your logs fill up!
- Edit joblog x all;f CSQJ002I all; del all x;look at timestamps

© 2015 IBM Corporation

Summary of what to collect

- Look at queue depths for application queues
  - Large depths → problem else where
- Check Buffer pool < 85%
- Check CF response time
  - And other CF information
- Check Response time of the MQ Logs
- Check DIS CHS
  - XBATCH < BATCHSZ
  - XQTIME
  - NETTIME
- TCP PING

© 2015 IBM Corporation

Summary

An “MQ performance problem” is usually a manifestation of a problem occurring elsewhere

© 2015 IBM Corporation

Thank you

© 2015 IBM Corporation