# *MQ for z – An Intro to Tuning and QMGR Management*

**“Oh my lord, where do I start?”**

**Lyn Elkins – elkinsc@us.ibm.com**

# Agenda

- **Background**
- **Basic Queue Manager Management tasks**
  - What do I do when I don’t know what I am doing?
  - Starting and stopping
  - How do I make changes?
  - Prior Day/Shift Review
- Why does MQ need to be tuned?
  - And what are the easy switches to hit

# Things to remember - Lynisms

- The best infrastructure on earth will not defeat bad application design and coding.
- The best design and coding on earth will not overcome the deficiencies in a poorly designed infrastructure.
- A good stress test plan won’t make you rich or famous, but it can make you better looking.  Just because you will be able to get some beauty sleep.
- If you don’t know what normal looks likes, how will you spot abnormal?

# Background

- **The intent of this session is to outline some of the administrative and tuning tasks that someone not familiar with MQ on z/OS might need in an emergency.**
  - Many times administrator's from other backgrounds will be asked to fill in on z/OS and vice versa
    - Do not be afraid!
  - This session is NOT covering setting up a queue manager on z/OS
- **MQ on z/OS is a different code base from the distributed platforms.**
  - This means that some of the admin and tuning are a bit different.
    - That might just be an understatement
  - Many of the basic administration tasks are the same, and can be accomplished using the same tools.
  - The impact of some of these tasks may be very different on z/OS.

# Administrative Tasks - Outline

- How do I start and stop a z/OS queue manager
- What tools are available to do these tasks?
- The Daily Grind:
  - **Starting and stopping queue managers**
  - Modification
    - Adding or changing objects
  - Monitoring
  - Maintenance
    - Backing-up is NOT hard to do
  - Troubleshooting

# Starting and stopping queue managers

- **On z/OS the queue managers are typically started and stopped by systems automation.**
- **They are usually only stopped for:**
  - Preventative maintenance (system or MQ)
  - System IPLs – regular cycles
- **There can be emergencies when an admin has to stop and start a queue manager manually.**
  - Or when a queue manager is first being built.
- **Commands to start the queue managers are quite different from the distributed commands.**

# Starting and stopping queue managers - commands

- **The format of the queue manager start command is:**
  **+cpf START QMGR PARM(QML1ZPRM)**
  - Where +cpf is the command prefix string assigned to this queue manager
  - START QMGR is the command
  - PARM identifies the load module that contains the basic definitions for this queue manager
    - This include information on the type of logs
    - Whether the queue manager is part of a QSG, etc.
  - The format of the stop command is:
    +cpf STOP QMGR MODE(QUIESCE)

# Starting and stopping queue managers – commands - Notes

- STARTing and STOPing a queue manager are documented here: http://www-01.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.adm.doc/q022170_.htm
- The system parameter module is documented here: http://www-01.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.explorer.doc/e_properties_zosconfiguring.htm

# Administrative tools

- **MQ provides several methods to perform most day to day administrative tasks:**
  - MQ Explorer
    - Defining/altering most resources
    - Starting and stopping channels
    - Reviewing event messages (see SupportPac MS0P)
  - Runmqsc
    - As of V8, can client attach to queue managers
  - z/OS only
    - ISPF panels
      - Stabilized
    - CSQUTIL
      - Batch z/OS equivalent of runmqsc

# Sample CSQUTIL

```
//DEFQCPYQ EXEC PGM=CSQUTIL,REGION=4096K,
//           PARM='QML1'
//STEPLIB  DD DSN=MQ800.SCSQANLE,DISP=SHR
//         DD DSN=MQ800.SCSQAUTH,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
 COMMAND
/*
//CSQUCMD  DD DISP=SHR,DSN=QML0.COBOL.MAKEDEF(QCPYQUES)
/*
```

Target Queue Manager

Output Location

Commands to the QMGR

Input file of commands

# MAKE Commands – Special uses of CSQUTIL

- The most commonly used MAKE commands are:
- MAKEDEF
  - Display commands in the input will generate DEFINE NOREPLACE commands in the output
- MAKEREP
  - Display commands in the input will generate DEFINE REPLACE commands in the output
- MAKEDEL
  - Display commands in the input for local queues will generate DELETE NOPURGE commands in the output

# MAKEDEF Sample

```
//DEFQS    EXEC PGM=CSQUTIL,PARM='QML3'
//STEPLIB  DD DISP=SHR,DSN=SYS1.MQM.SCSQANLE
//         DD DISP=SHR,DSN=SYS1.MQM.SCSQAUTH
//         DD DISP=SHR,DSN=SYS1.MQM.SCSQLOAD
//MDOUT    DD DISP=SHR,DSN=QML0.COBOL.MAKEDEF(QCPYQUES)
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD  *
  COMMAND DDNAME(CMDINP) MAKEDEF(MDOUT)
//CMDINP   DD  *
DISPLAY QUEUE(QCPY.*) ALL
/*
```

Target Queue Manager

Output Location

Create DEFINE commands

Input file of commands

# Sample runmqsc as a client

- **SET MQSERVER=SYSTEM.ADMIN.SVRCONN/TCP/mpx1(1417)**
  - This command sets the MQSERVER environment variable to connect to the z/OS queue manager on host mpx1, listening on port 1417
- **runmqsc -c -uELKINSC QML1**
  - The ‘-c’ tells runmqsc to use a client connection
  - The –u supplies the user id
    - Please note that the uppercase is necessary in this case because the RACF system does not automatically fold lower to upper case in this environment
  - QML1 is the z/OS queue manager
- **pause**

# Adding and modifying MQ Objects

- **Most common administrative task on MQ!**
  - Of these defining and modifying queues is probably the most frequent.
    - When adding new queues:
      - Think about queue placement
        - » On z/OS this is the storage class attribute, which maps back to a pageset and bufferpool for private queues
        - » For shared queues the definitions needs the CFSTRUCT attribute
      - Think about cluster membership
      - Think about monitoring requirements

# Queue Definition Notes

- **One area of specific difference between distributed and z/OS is in the underlying message storage**
  - Private queues
    - Bufferpools (in memory message storage)
      - If you do not know which ones are heavily in use, defining a queue in a heavily used area could cause issues
    - Pagesets - “back up storage” – used for long term message storage and when the bufferpool becomes over full
      - VSAM File, incurs real I/O when used
  - Shared queues
    - Coupling Facility – always holds at least the message control information
    - Offloaded Messages
      - Offloading always used when messages are >63K
      - Shared Message Data Sets – VSAM files
      - DB2 Blobs

# Adding and modifying MQ Objects

- **Most common administrative task on MQ!**
  - Next most common is defining and modifying channels
  - Defining channels is virtually the same on z/OS as distributed
    - Except where noted
  - When adding new channels, think about:
    - Transport type – typically TCP, but z/OS still defaults to LU6.2
    - CONNAME – limited to 48 characters on z/OS, 264 on distributed
    - Channel Authorization Required?
      - » If so, what are the rules?
    - IS SSL/TLS required?
      - » If so, need to interact with the security (RACF, etc.) on both sender and receiver side for set-up

# Adding and modifying MQ Objects

- When defining send/receiver channels some additional things to think about:
  - Who is defining the counterpart?
    - Make sure the connection is tested before first production use
  - Both channel authorization and SSL/TLS considerations may involve multiple groups:
    - Network/distributed platform security administrator
    - RACF or other ESM administrator

# Adding and modifying MQ Objects

- Next most common is defining and modifying channels (continued)
  - Sender channels
    - » What is the target queue manager?
    - » Has my transmission queue been defined?
    - » Should it be triggered?
  - Receiver channels
    - » What is the sending queue manager?
    - » Do I use the DLQ for undeliverable messages?
  - Defining SVRCONN channels into MQ for z/OS
    - – Remember that using clients connecting directly to z/OS queue managers can drive up MIPS!

# IBM MQ – Thinking about tuning?

- **Where do I begin?**
  - **MQ JES logs**
  - **IBM MQ Statistics**

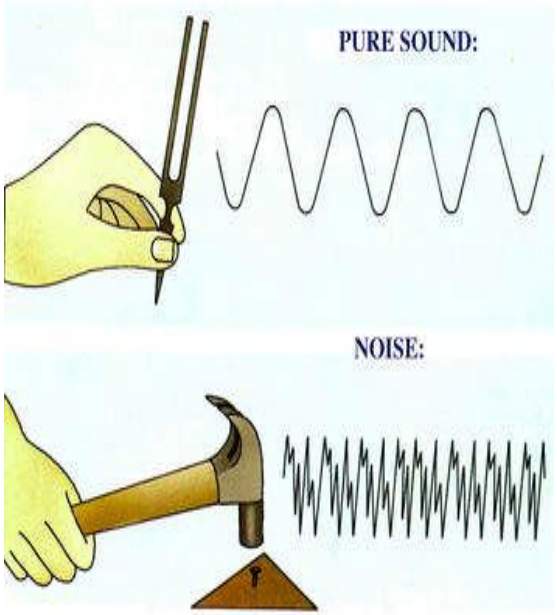

# What data needs to be collected/reviewed?

- **Typically to really start tuning a queue manager:**
  - One week of MQ Statistics data – SMF 115
    - For MQ V8 queue managers this should include classes 1 and 4
    - For pre V8 queue managers, if gathering data about channels using the MQCMD program (or similar) if using is also helpful
  - One high volume day's JES logs for the queue manager and channel initiator

- **For those not familiar with z/OS SMF is System Management Facility**
  - This is a tools used by just about every subsystem to emit data abut how that subsystem is processing
  - MQ can emit both Statistics and Accounting information
    - Statistics are all we shall talk about in here
    - Started and stopped via the:
      - System parameter module
      - START TRACE(STAT) command
- **One week of MQ SMF 115 usually provides a good idea of what is going on in both online and batch process cycles.**
  - There can be a real danger in tuning for one workload, impacting the other workload!

# IBM MQ for z/OS Tuning – Where do I start looking?

- **There are two primary locations for basic tuning data**
  - IBM MQ JES logs
    - Both the MSTR and CHIN address spaces have messages that can show tuning opportunities
      - Storage usage  - Example
        - » CSQY220I BWF0 Queue manager is using 607 MB of local storage, 1128 MB are free
      - Also contains the results of commands (DISPLAY QMGR, etc.)
    - Review regularly
      - Often left to automation tools

# IBM MQ – Tuning

- As a warning, leaving the review of the JES logs up to automation completely can be detrimental to your health.

# Intro to Tuning - Review of SMF115

- **The SMF 115 data is the statistical information produced by a MQ for z/OS queue manager.**
  - Primarily used to track major trends and resolve performance problems with the queue manager
  - Very lightweight
    - Two or three records per queue manager per SMF interval
  - Broken down into the major ‘managers’ within MQ
  - SupportPac MP1B, no matter which version, provides at least 2 views into the MQ statistics
    - Report form
    - Column or CSV form
      - This is used to import into spreadsheets for analysis
    - Both are necessary because the CSV form may be missing fields you need to review when looking for a problem.

# What is this stats data good for?

- Detecting patterns of use (and overuse!)

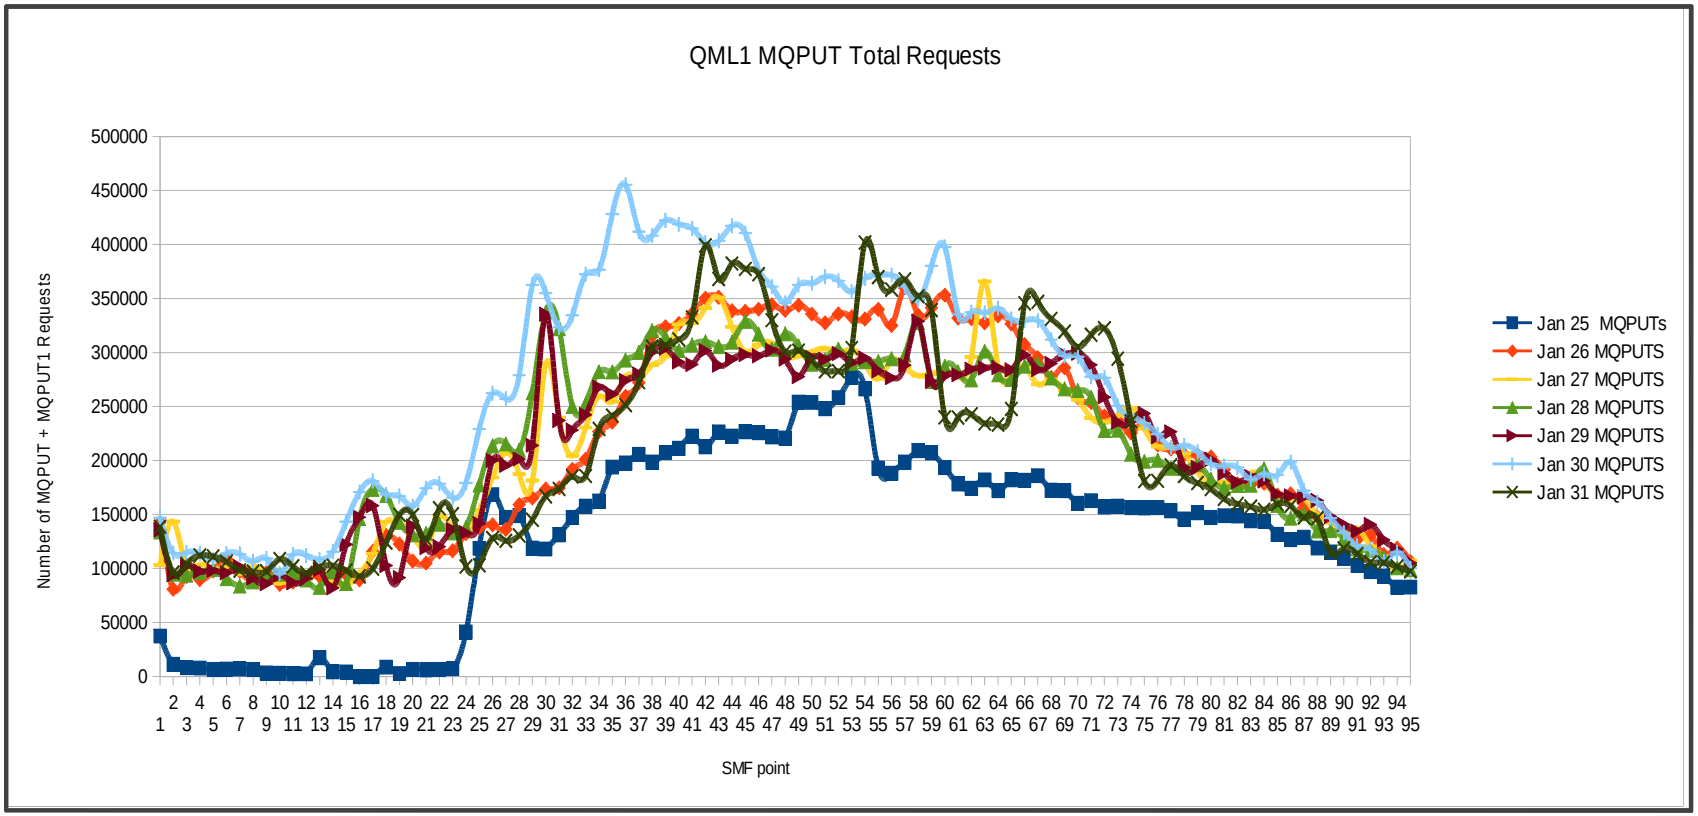

- **One very common question is ‘how busy is my queue manager?’**
  - The Message manager statistics tell you how many MQ API requests have been made during the SMF Interval
    - This is not ‘How many messages have been processed’
- **The statistics give you a high level look into the queue manager**

# SMF 115 – using the data

- If using the MP1B SupportPac for format and print the data, the CSV file or the column format can be downloaded from z/OS and then imported into a spreadsheet for evaluation.
- For example, the log data loaded into a spreadsheet looks as follows:

| Date | Time | QMGR | wr_wait | wr_nwait | Aug08 Force Writes | Aug08 Log Buffer Waits | read_buf | read_act | read_arc | r_delay | N_CheckP | Aug08 Num I/O | Aug08 Control Intervals Written | paging |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2011220 | 00:02:01.06 | QML1 | 0 | 69242 | 447 | 0 | 0 | 0 | 0 | 0 | 0 | 6726 | 20658 | 0 |
| 2011220 | 00:03:01.05 | QML1 | 0 | 74069 | 403 | 0 | 0 | 0 | 0 | 0 | 0 | 6866 | 22446 | 0 |
| 2011220 | 00:04:01.04 | QML1 | 0 | 72830 | 453 | 0 | 0 | 0 | 0 | 0 | 0 | 6726 | 22550 | 0 |
| 2011220 | 00:05:01.08 | QML1 | 0 | 70513 | 478 | 0 | 0 | 0 | 0 | 0 | 0 | 6672 | 20870 | 0 |
| 2011220 | 00:06:01.07 | QML1 | 0 | 77959 | 489 | 0 | 0 | 0 | 0 | 0 | 0 | 7166 | 23458 | 0 |
| 2011220 | 00:07:01.06 | QML1 | 0 | 60920 | 483 | 0 | 0 | 0 | 0 | 0 | 0 | 6342 | 17732 | 0 |
| 2011220 | 00:08:01.06 | QML1 | 0 | 65818 | 444 | 0 | 0 | 0 | 0 | 0 | 0 | 6662 | 18960 | 0 |
| 2011220 | 00:09:01.03 | QML1 | 0 | 68325 | 476 | 0 | 0 | 0 | 0 | 0 | 0 | 6740 | 19982 | 0 |
| 2011220 | 00:10:01.08 | QML1 | 0 | 50712 | 478 | 0 | 0 | 0 | 0 | 0 | 0 | 6000 | 12088 | 0 |

# SMF115 – Using the data

- **By itself the SMF 115 data for many of the internal component managers are useful, but some manipulation of the raw data can make it much more consumable and useful.**
- **Examples include:**
  - The message manager MQPUT chart as shown earlier
  - Calculating the log data written per second
  - Charting, making it easier to spot patterns, problems, and trends
- **Some of the calculated data is done for you in the 'new' MP1B, but beware – it uses hard coded values for intervals, etc. which can skew results.**

# Logging Rates - Charted

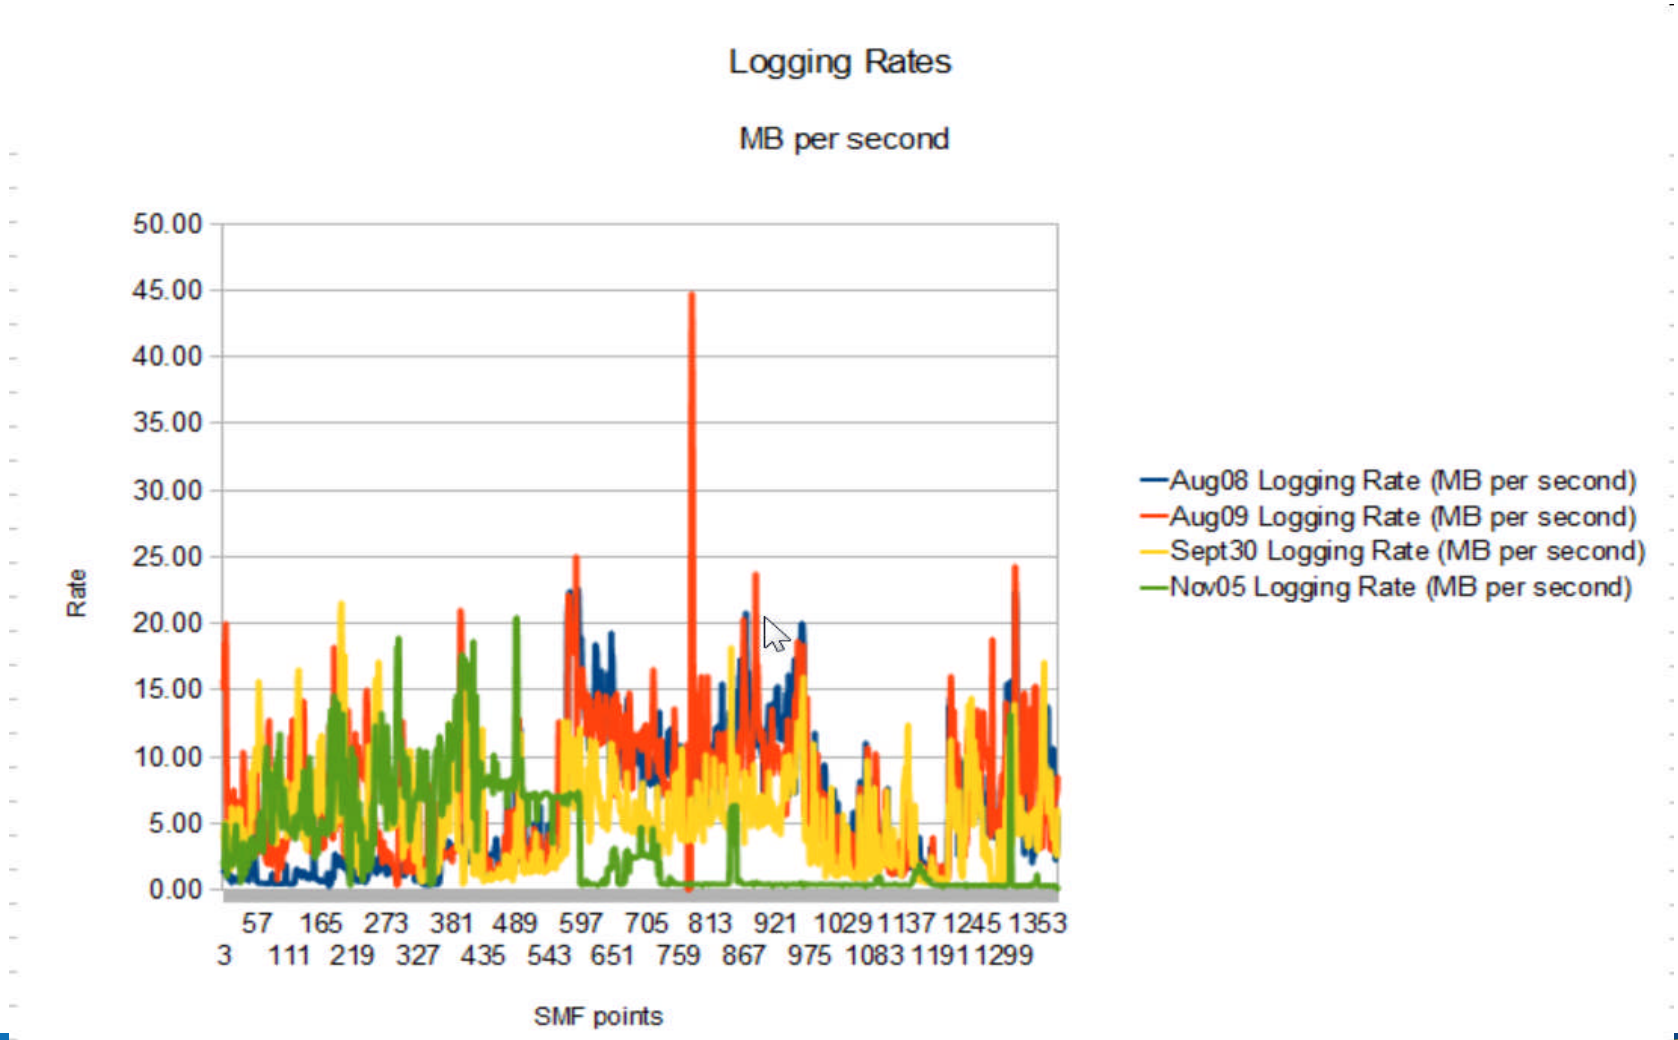

- **Writing to the MQ log is often the critical factor in a queue manager being able to absorb more workload**
  - It can also be key to performance on a heavily persistent messaging workload
- **This logging chart was taken from a problem determination situation – the queue manager "could not keep up". There are a couple of important things to note:**
  - The big spike was first thought to be an extremely high logging rate due to workload volume OR that the log during that period was on different hardware.
    - It was not, it was because the SMF interval had been altered dynamically on that day.
  - Other than the big spike, the highest MB/second was about 25. In reviewing the detailed data, there were a large number of waits on logging. Further investigation uncovered the real problem – the I/O subsystem had a maximum rate of 25MB/sec and MQ and DB2 were heavily competing!
    - In looking at the data, the logging rate is frequently at the 20/25 MB per second rate. For the environment, this was quite high. It was discussed with the capacity planning team. At the time they were looking into the purchase of a new hardware, which was rated at up to 100 MB/second. This information provided additional emphasis on the need to upgrade, which has since been done to good effect.

# SMF115 – Using the data – continued

- **Example of using a spreadsheet to calculate the log writes:**

| Num I/O | Num_CI _W | paging | | Seconds | Interval Duration | | Log writes in MB per second |
|---|---|---|---|---|---|---|---|
| 142446 | 142460 | 0 | | 980 | | | |
| 13352 | 13352 | 0 | | 1877 | 897 | | 0.06 |
| 12420 | 12420 | 0 | | 2775 | 898 | | 0.05 |
| 10956 | 10956 | 0 | | 3670 | 895 | | 0.05 |
| 10484 | 10484 | 0 | | 4567 | 897 | | 0.05 |
| 9274 | 9274 | 0 | | 5463 | 896 | | 0.04 |
| 8946 | 8946 | 0 | | 6359 | 896 | | 0.04 |
| 8864 | 8864 | 0 | | 7256 | 897 | | 0.04 |
| 1476 | 1476 | 0 | | 8154 | 898 | | 0.01 |

- **The formulas used to calculate seconds and duration:**
  - Seconds =SUM(HOUR(B2)*3600,(MINUTE(B2)*60),SECOND(B2))
  - Interval duration =(Q3-Q2)
  - Interval crossing a date =SUM((86399-Q96),Q97)
- **The formula for log writes in MB/Second**
  - =ROUND(((N3*4096)/(1024*1024))/(R3),2)
  - Where N3 = the number of CI Writes column and R3 is the Interval duration

# Reviewing the SMF 115 data – Are there problems?

- **While not always the culprit, tuning the buffer pools has been a never ending cycle of opportunities**
  - Lazy queue definitions
    - I copy what works, might not be what is best for the queue manager
  - Workload pattern changes
    - What flows today, might become a log jam without any warning
      - Well except the business knew they were opening 2600 new branches on Tuesday

# SMF 115 data – Urgent symptoms

- **In the next few slides, an analysis of a bufferpool under stress is shown**
- **First the raw SMF data for two weeks was processed thru the "old MP1B" MQCSMF and the MQ1150 format and print programs**
- **The Buffer Manager statistics were downloaded into a spreadsheet**
- **The spread sheet was sorted to find:**
  - Non-Zero Short on storage counts
  - Non-Zero DMC counts
  - Percent of free pages
- **This showed the areas that needed to be looked at in greater detail, and it became apparent that there were some processing patterns that need evaluation**

# SMF 115 data – Urgent symptoms – Call 911

- **Unless there has been an unexpected spike in volume, this should never been seen!**
  - SOS
  - Freepages at 5% or less

| QMGR | BP | NumBuff | %now | %low | dwt | dmc | stl | stla | sos |
|---|---|---|---|---|---|---|---|---|---|
| QML2 | 3 | 70000 | 18 | 0 | 109 | 198906 | 922354 | 1 | 50 |
| QML2 | 3 | 70000 | 19 | 0 | 68 | 143872 | 387873 | 1 | 13 |

| Date | Time | QMGR | BP | NumBuff | %now | %low | dwt | dmc | stl | stla | sos |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 201133 | 408:15:21. | QML1 | 3 | 70000 | 98 | 5 | 9 | 27 | 32557 | 0 | 0 |
| 201133 | 420:41:19. | QML1 | 3 | 70000 | 95 | 5 | 2 | 364 | 61145 | 0 | 0 |

# SMF 115 data – Urgent symptoms

- **Red Flags for Bufferpools - Continued**
  - DMC – synchronous write process is requested
    - Note that it did not run this often, but this is the number of times that conditions were right!

| QMGR | BP | NumBuff | %now | %low | dwt | dmc | stl | stla | scs |
|---|---|---|---|---|---|---|---|---|---|
| QML3 | 3 | 70000 | 16 | 0 | 58 | 210092 | 853991 | 1 | 0 |
| QML3 | 3 | 70000 | 22 | 3 | 132 | 36526 | 1232774 | 2 | 0 |

# SMF115 – Bufferpool Trends and Analysis - Notes

- In the chart shown two high volume days were compared to see if there was a pattern to the BP use.
- BP 0, 1 an 2 showed almost no utilization.
- BP 3 was in very heavy use, some of the time.
- BP 3 is under some stress.
- Having multiple days worth of data is vital, had there just been one heavy day it may have been an anomaly. Data from longer periods of time, when compared like this can be very useful in tracking usage, etc.
- In this case there was a clear pattern of overuse of bufferpool 3, in further evaluation the SMF116 data showed that all the queues that were being used for this queue manager were defined on the same pageset/bufferpool. By moving some of the queues to another resource pool, the stress was reduced, work flowed faster and the CPU usage was reduced.
- In attempting to replicate the issues, the information on the previous slides was used to create the charts, but also to show that charting the pattern might be helpful in the evaluation.

# SMF 115 data – Clairvoyance – Ask Doctor

- **May be approaching a problem**
  - Consistently Approaching/Achieving  20 % Free pages

| QMGR | BP | NumBuff | %now | %low | qwt | dmc | stl | stla | sos |
|---|---|---|---|---|---|---|---|---|---|
| QML4 | 2 | 70000 | 53 | 19 | 0 | 0 | 46571 | 0 | 0 |
| QML4 | 3 | 70000 | 98 | 20 | 0 | 0 | 46028 | 0 | 0 |
| QML4 | 3 | 70000 | 75 | 20 | 0 | 0 | 0 | 0 | 0 |

# Bufferpool Use - Warnings

- In the data shown, two bufferpools were approaching the 20% freepage threshold.
- At 20% the async write task is initiated, which is not catastrophic, but if it can be avoided it should be.
- In this case, when several weeks worth of data were examined the 20% threshold was being broken on a regular basis.  After evaluation fo the SMF116 class-3 data it was found this was batch oriented workload, and messages were expected to queue up for long periods of time this was not a problem.  It is something to watch.

# SMF 115 data – Reviewing the Message Manager information for gross use numbers

- **Message Manager Information**
  - Good indication of queue manager usage
    - This is only a count of API calls, not one of successful calls
    - Volume trends can be approximated from the MQPUT and MQPUT1 calls, as these are generally successful
    - MQGETs may or may not have data returned

| QMGR | Open | Close | Get | Put | Put1 | Inq | Inq | Set | Total API calls | Total Puts |
|---|---|---|---|---|---|---|---|---|---|---|
| QML1 | 160 | 151 | 2,925,084 | 3,417,313 | 0 | 1 | 0 | 0 | 6,342,709 | 3,417,313 |
| QML1 | 248 | 228 | 2,256,084 | 3,150,666 | 0 | 5 | 0 | 0 | 5,407,231 | 3,150,666 |
| QML1 | 897 | 895 | 3,468,114 | 3,093,355 | 0 | 50 | 0 | 0 | 6,563,311 | 3,093,355 |

# Message Manager Statistics

- This data was taken from the message manager output from the old MQCSMF format and print program.
- Two columns were added to calculate the Total API calls and Total Puts (sum of MQPUT and MQPUT1 calls)
- When charted over a few weeks an upward curve was noticed.

# Message Manager - Trend Chart

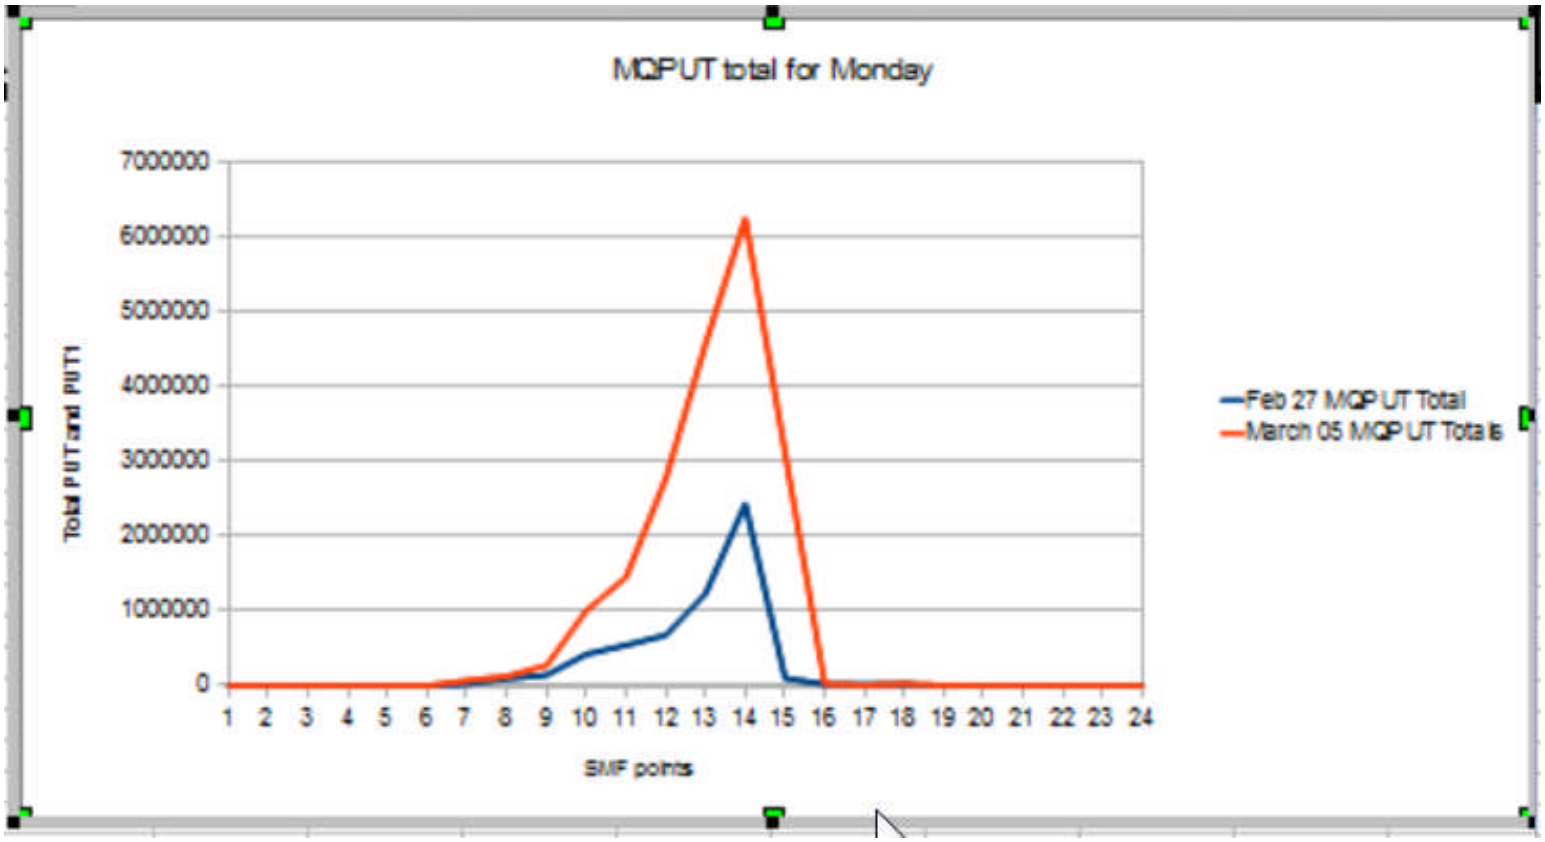

# Message Manager Trend Chart

- Two days data is not really a trend analysis, but it’s a start
- If more Mondays are charted, a real trend may emerge and show that volume is increasing allowing a good admin to plan for additional workload.
- This is an overall count for the queue manager, individual queue activity can be evaluated from the SMF116 class 3 data.

# Reference Material

- **Manuals & Info Centers:**
  - http://www-01.ibm.com/support/knowledgecenter
- **SupportPacs:**
  - http://www-01.ibm.com/support/docview.wss?rs=977&uid=swg27007205
- **Developer Works:**
  - http://www.ibm.com/developerworks/
- **My blog**
  - http://www.lynsmq4zos.com

# Q&A

- Thank you for your attention!