MP1B: IBM MQ for z/OS

Displaying IBM MQ statistics and accounting

Colin Paice

June 2017

Document Number MP1BV90 Version: 5.0.0

Take Note!
Before using this User's Guide and the product it supports, be sure to read the general information under "Notices".

V900 Edition, June 2017.

This edition applies to Version 9.0, 8.0, 7.1 and V7.0.1 of "MQ for z/OS" - Interpreting accounting and statistics data" and to all subsequent releases and modifications until otherwise indicated in new editions.

Sending your comments to IBM

You can send your comments electronically to idrcf@uk.ibm.com.

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you. You may continue to use the information that you supply.

Document generated June 16, 2017 at 13:28

Notices

This report is intended to give guidance on the use and interpretation of the statistics and accounting in IBM MQ for z/OS Version 9.0, 8.0, WebSphere MQ Version 7.1 and 7.0.1. The information in this report is not intended as the specification of any programming interfaces that are provided by z/OS or IBM MQ.

References in this report to IBM products or programs do not imply that IBM intends to make these available in all countries in which IBM operates.

Information contained in this report has not been submitted to any formal IBM test and is distributed "as is". The use of this information, and the implementation of any of the techniques, is the responsibility of the customer, and depends on the customer's ability to evaluate and integrate them into their operational environment.

The following terms, used in this document, are trademarks of the IBM Corporation in the United States or other countries or both:

CICS

IBM MQ

z/OS

Table of Contents

Changes History

This SupportPac supports IBM MQ for z/OS V9, IBM MQ for z/OS V8 IBM WebSphere MQ for z/OS V7.1 and earlier releases.

V900 April 2017

- Page set statistics showing similar information to DIS USAGE, and includes I/O statistics
- Changed grand totals in the task section to be seconds rather than float micro seconds, For example 1.268008 S instead of 1.26801e+06 uS. This is to facilitate sorting this field.
- Output of Log data to CSV - and allow columns to be specified

March 2016

- SMDS data in CSV format. Files //SMDSSCSV //SMDSBCSV, //SMDSACSV.
- Change QMST output to record rate per second of gets and put + put1s, rather than gets and puts
- Added MaxChl into the CHINITCSV header.
- Documented some missing messages
- Job information added to right hand side of some log output to facility sorting and displaying key values
- Put qmgr name and other information into some buffer statistics records to aid data extraction and sorting
- Fixed wrong field displayed in BUFFCSV
- Put qmgr name and other information into some log statistics I/O records to aid data extraction and sorting
- Experimental capability to output to CSV files and let the user specify the fields to be output.

August 2015

- CSV files are produced for gets and puts by queue per transaction, and summarised by queue. See PUTSCSV, PUTSCSVS, GETSCSV, GETSCSVS
- Output files are opened using fopen and now use the " w " attribute to (re)write the files. Before it used " a " which caused output to be appeneded to uss files
- Fields present in record are now displayed, such as Open Suspend Time.
- New CSV files for SMDS
- Summary of all queue activity is displayed in //QALL
- Queues doing page set I/O are displayed in //PSIDQIO
- More threshold keywords. SMDSWriteTime, SMDSReadTime, SMDSWaitFree, SMDSWaitBusy.
- Buffer pool CSV includes pages used when record was created. % full is now % of the highest used pages. It was % used of current pages.

December 2014

- Chinit SMF. Channels now record and display
  - DNS resolution time
  - SSL Certificate serial number
  - CN from SSLCERT
- From the accounting class 3 data, queues which directly cause I/O to the page set is reported in //BUFFIO
- SMDS calculates the MB/Second
- Numbers were sometimes inconsistent between CF report and CSV file
- Always report the summary of MQ SMF records and subtypes found

- Fixed problems with processing the data, where invalid results were displayed.
- Fixed problems where summary information in //TASKSUM was not displayed properly
- Display buffer pool pages read and written per second
- Log maximum times now report 1:1 and 2:1, where the first number is the log copy number, and the second is one page written.
- In V701 the amount of data logged was the total data logged - you could not tell if this was due to single logging or dual logging. In V710 and later, the data logging per individual log is displayed. This may appear as a reduction of data logged per second.
- The label of the Log statistics, Checkpoints, has been changed to LogLoad Checkpoints(LLCheckpoints) to show this is due to the LogLoad value, not the logs filling up.

August 2014 to support V8

1. Support of the SMF data for the chinit statistics and channel accounting
   (a) Adapter, Dispatcher, SSL, DNS TCB information
   (b) Information similar to display channel status for channels.
   (c) Report if the average SSL duration is larger than a user specified value
   (d) Report if the nettime is larger than a user specified value
2. Update buffer pool statistics to include 64 bit buffer pools and 'fixed' buffer pools
3. Update to log statistics showing average times for I/O requests. See APAR PM61284.
   (a) Using the information allows the *log busy%* to be calculated.
4. SMF statistics have the start time and duration of the records, so rates can be calculated, such as MB logged per second per log. See APAR PM61284.
5. Improved selection criteria
   (a) Display queue records only if the curdepth is greater than user specified value.
   (b) Specify Jobname, such as MQ02CHIN to display records only from that job.

V7.1 APAR PM61284: Addition to log manager statistics. This adds information about log I/O times and the longest I/O time in the SMF interval. It also adds the duration of the SMF record.

What you can do with the program

You can use the MQSMF program in different ways depending on your requirements.

- What work is using my queue manager?
  - List a summary of the transactions, jobs and channels on page 6
  - List a summary of the queues being used by which application on page 6
  - List the high use queues on page 6.
- List a summary of the queues being used by which application on page 6
- List the high use queues on page 6.
- How do I monitor my queue manager? on page 7.
- Investigate potential problems, on page 7.
- Are the application well behaved? on page 8.
- Is the queue manager set up OK? on page 8.

Some of the output files are in Comma Separated Value (CSV) format which can be imported into a spread sheet. The spread sheet can then be used to draw graph and display trends over time.

Start your analysis by looking in the //TASKSUM section which summarises any problems identified. For example if a messages was written multiple times in the output file, the worst case is displayed in the //TASKSUM output. See Investigate potential problems on page 7.

Summary of the transactions, jobs and channels

To get a summary of the applications and channels using MQ use the TASKCSV file, see page 89. This summarises the Accounting class 3 records by Date, Hour and work type. This report has information on the CPU used, MB of data logged, how many MB were put and got. Data for multiple instances of a transaction or a job are summarised into one record.

Summary of the queues being used by which application

To get a summary of the queues used, use the QSUML file on page 45 for information about local queues , and QSUMS file on page 45 for information about shared queues. These reports have information on the number of MB put and got, the number of valid gets and puts, the maximum queue depth and, for local queue the number of messages read from the page set

What are the high use queues?

You can use the QALL section to identify the queues which use a lot of CPU, or are highly active.

The QALL report prints out the data summarised by the queue name, so if there were 100 CICS transactions processsing the queue, there would be one set of data for the queue, showing 100 opens and closes etc.

- Use ISPF view (or edit) to edit the file. You can use the following commands to display the high usage queues
  - X ALL
  - f 'Total Queue CPU used' all
  - del all x
  - sort 40 50 d

exit from the file without saving.

To display where time has been spent

- view the file again
- X ALL
- f 'Total Queue elapsed time' all
- del all x
- sort 40 50 d

You can use the QSUML report on page 45 and the QSUMS report on page 45, and specify parameters QueuePutMB and/or QueueGetMB to the MQSMF program. Queues which put or get less than these values are not displayed. This eliminates low usage queues, such as using a temporary dynamic queue to use the command server.

How do I monitor my queue manager

You can monitor the amount of data logged by using the log data statistics in CSV format, see page 51.

You can monitor buffer pool usage using the buffer pool statistics in CSV format, see page 57.

You should collect data for a good day, so you have a base line to compare other days with.

Investigate potential problems

In ddname //TASKSUM is a summary of messages produced when looking at the task and queue records.

This has data like

```
Record#    Count    Value Message
   2202       25    98908 MQTASK13E long commit time C,'CP15','IYFFC000',
     38        1   106347 MQTASK13E long commit time B,'PAICEP7A',' ',
```

This has the following meaning.

- There was a message `MQTASK13E long commit time C,'CP15','IYFFC000',`
  - It was produced 25 times
  - The longest value (of the commit time) was 98908 microseconds. This was at record 2202 in the input file
- There was a message `MQTASK13E long commit time B,'PAICEP7A',' ',`
  - It was produced 1 time,
  - the largest (only) value was 106347 in record 38 of the input file.

To investigate these in more detail you can use specify parameters to the MQSMF program: FirstRecord 2202, LastRecord 2202 and Detail 20. This will give all maximum level of detail for the one record.

Messages are written to the ddname //MESSAGE report to indicate possible problems, such as messages read from a page set, long CF response time. By increasing the value specified in the Detail, you can get more potential messages produced.

Messages from the CHINIT statistics are written to the ddname //CMESSAGE.

For example *MQQPST04E MVCA MQQ2 2013/02/01 12:54:27 VRM:701 BP 10 Many (11317) pages read from disk. This is typical of long lived messages. Buffer pool may be too small*

MQQPST02S MVCA MQQ2 2013/02/01 12:50:46 VRM:701 BP 10 Filled many(165) times. This is typical of long lived messages. Buffer pool may be too small

Once you have identified a potential problem you can use the parameter FirstRecord and

LastRecord to display a subset of the SMF data, and use Detail to display more information for this subset of records.

If you think you have problems with puts taking a long time, you can specify the Long_Put parameter, and for puts taking over this time it will display a message like

```
MQTASK08E Long Put time due to logging MYQUEUE
```

Are the application well behaved

Messages are written to the MESSAGE report when possible unusual application programming behaviour is detected. For example

- Many gets for a specific message, but the queue is not indexed.
- An application repeatedly failed to get a message from a queue. This might be caused by the common programming error, where the message-id and correlation-id are not cleared before doing a get.

Is the queue manager set up OK?

If a buffers pool is too small and frequently fills up, messages are reported in the MESSAGE report suggesting area you may need to investigate.

JCL to run the MQSMF program

When you have run the program check the output in //SYSOUT for messages like *error opening //LOG : EDC5037I The specified ddname was not found.*

This indicates the //LOG statement is missing in the JCL.

```
//S2 EXEC PGM=MQSMF,REGION=0M
//STEPLIB DD DISP=SHR,DSN=Your.load.library.name
//SMFIN    DD DISP=SHR,DSN=*.SMFDUMP.DUMPOUT
//SYSIN   DD *

Parameters - see below

//CHINIT  DD SYSOUT=*   new
//CHINCSV DD SYSOUT=*   new
//ADAPCSV DD SYSOUT=*   new
//SSLCSV  DD SYSOUT=*   new
//DISPCSV DD SYSOUT=*   new
//DNSCSV  DD SYSOUT=*   new
//ADAP   DD SYSOUT=*    new
//SSL     DD SYSOUT=*   new

//PSIDQIO DD SYSOUT=*  new
//SMDSSCSV DD SYSOUT=* new
//SMDSBCSV DD SYSOUT=* new
//SMDSACSV DD SYSOUT=* new
//PUTSCSV  DD SYSOUT=*,DCB=(LRECL=200) new
//PUTSCSVS DD SYSOUT=*,DCB=(LRECL=200) new
//GETSCSV  DD SYSOUT=*,DCB=(LRECL=200) new
//GETSCSVS DD SYSOUT=*,DCB=(LRECL=200) new
//QALL     DD SYSOUT=* new


//PSET DD SYSOUT=* new
//BUFFIO DD SYSOUT=* new

//DISP    DD SYSOUT=*   new
//DNS     DD SYSOUT=*   new
//DCHS    DD SYSOUT=*   new
//DCHSCSV DD SYSOUT=*   new
//DCHSSUM DD SYSOUT=*   new
//CMESSAGE DD SYSOUT=*  new
//MESSAGE DD SYSOUT=*
//BUFF    DD SYSOUT=*
//BUFFCSV DD SYSOUT=*,DCB=(LRECL=200)
//DATA    DD SYSOUT=*
//QCPU    DD SYSOUT=*
//CF      DD SYSOUT=*
//CFCSV   DD SYSOUT=*
//DB2     DD SYSOUT=*
//EOJ     DD SYSOUT=*
//LOCK    DD SYSOUT=*
//LOG     DD SYSOUT=*
//LOGCSV  DD SYSOUT=*
//LOGBUSY DD SYSOUT=*
//MSGM    DD SYSOUT=*
//MSGMCSV DD SYSOUT=*
//QSUML   DD SYSOUT=*,DCB=(LRECL=200)
//QSUMS   DD SYSOUT=*,DCB=(LRECL=200)
//STGCSV  DD SYSOUT=*
```

```
//SMDS    DD SYSOUT=*
//TASKSUM DD SYSOUT=*
//TASK    DD SYSOUT=*
//TASKCSV DD SYSOUT=*
//TASKET  DD SYSOUT=*,DCB=(LRECL=200)  new
//TOPIC   DD SYSOUT=*
//STG     DD SYSOUT=*
//STGSUM  DD SYSOUT=*,DCB=(LRECL=200)
//
```

Where the DD statements are

CHINIT	page 18, Chinit statistics
CHINCSV	page 18, CHINCSV
ADAP	page 18, Adapter
ADAPCSV	page 20, Adapter CSV
DISP	page 21, Dispatcher tasks
DISPCSV	page 22, Dispatcher CSV
DNSCSV	page 23, Domain Name Server(DNS) CSV
SSL	page 24, SSL
SSLCSV	page 24, SSLCSV
DNS	page 23, Domain Name Server (DNS) task
DCHS	page 25, Channel accounting
DCHSCSV	page 27, Channel accounting CSV
DCHSSUM	page 27, Channel summary CSV
CMESSAGE	page 27, Chinit Messages(CMESSAGE)
SMDSSCSV	page 41, SMDS space
SMDSBCSV	page 41, SMDS buffers
SMDSACSV	page 42, SMDS activity
PSIDQIO	page 33, Which queue caused I/O to a page set
PSET	page 29, Page set queue I/O VNEXT
BUFFIO	Page 31, Buffer pool IO statistics
PUTSCSV	Page 91, CSV with information about puts
PUTSCSVS	Page 91, CSV with information about puts, summarised by queue
GETSCSV	Page 91, CSV with information about get
GETSCSVS	Page 91, CSV with information about get summarised by queue
QALL	Page 86, detailed information about queues summarised by the queue name
BUFFCSV	Page 57, Buffer pool statistics in comma separated value format
CF	Page 34, Coupling facility statistics
CFCSV	Page 35, Coupling facility in comma separated value format
DATA	Page 62, Data manager statistics
LOG	Page 46, Log manager statistics
LOGCSV	Page 51, Log manager statistics in comma separated values
LOGBUSY	Page 51, Logbusy
LOCK	Page 66, Lock manager statistics
MSGM	Page 62, Message manager statistics (mq verbs)
MSGMCSV	Page 62, Message manager statistics (mq verbs) in comma separated value format
STG	Page 66, Storage manager statistics
TOPIC	Page 63, Topic manager statistics
QSUML	Page 45, Summary of queue usage - Local queues
QSUMS	Page 45, Summary of queue usage - Shared queues
EOJ	Page 61, Subsystem information
DB2	Page 70, DB2 server information
SMDS	Page 36, Shared Message Data Set Statistics (SMDS)
STGSUM	Page 92, Address space level storage usage
SYSOUT	Output of warnings and problems identified in the data
SYSPRINT	Reports the parameters used, and task comments
TASK	Page 73, Detailed task data
TASKCSV	Page 89, Summarized task information

You will get data in QSUML, QSUMS, TASK,TASKCSV, and TASKSUM if TRACE(A)CLASS(3) is specified. You will get information on queues if the queue attribute ACCTQ is ON, either by setting it explicitly on a queue, or having ACCTQ(QMGR) and the QMGR attribute ACCTQ has the value ON.

Record selection parameters

The following section gives the parameters which can be passed to the program, through the //SYSIN ddname.

Value	Range	Default	meaning
CFStruct			CF Structure name
Channel			Only display channels beginning with the specified value.
Curdepth	0 to 999999999	0	If curdepth is greater than 0 then information about the queues are only displayed when the the queue's *Curdepth maximum* is larger than the value specified. You can use this to display queue with unusually deep queues, such as transmission queues.
CICSTRAN			CICS transaction name. Selects transactions beginning with the specified value.
FirstRecord	0 to 999999999	0	specify starting record - 0 means not set
LastRecord	0 to 999999999	0	specify ending record - 0 means not set
Jobname			The prefix of jobs to be selected for example MQ02CH
QM			Select records by this queue manager
Queue			Select accounting records by this queue
QueueGetMB	0 to 9999	0	When this value is greater than 0, queues getting more than this value of MB of data, are reported in the QSUM* files
QueuePutMB	0 to 9999	0	When this value is greater than 0, queues putting more than this value of MB of data, are reported in the QSUM* files
StartTime			Use records if the start time is after hh:mm:ss
EndTime			Use records if the End time is before hh:mm:ss
zOS			Select record from this z/OS image

Interval for calculating data rates.

Value	Range	Default	meaning
SMF_Interval_time	1 to 999999999	60	The interval between SMF records in minutes. This is used when converting values to rates/second. For example the amount of data logged per second. This should be used when running MQ V7 without APAR PM61284, as this APAR adds the interval duration.

Thresholds for reporting out-of-line conditions in statistics

Value	Range	Default	meaning
SMDSWriteTime	1 to 999999999	1000	Report if the SMDS average Write time theshold in microseconds is larger than this value
SMDSReadTime	1 to 999999999	1000	Report if the SMDS average Read time theshold in microseconds is larger than this value
SMDSWaitBusy	1 to 999999999	1000	Report if the applications had to wait for a busy buffer is larger than this value
SMDSWaitFree	1 to 999999999	1000	Report if the applications had to wait for a free buffer is larger than this value

ReadAheadMsgCount	0 to 999999999	0	Report if the number of messages read ahead was larger than this value
MsgsReadFromDisk	0 to 999999999	0	Report if the number of messages read from disk was larger than this value
BPIMW	1 to 999999999	1000	Report if buffer pool had more than this number of immediate writes
BPReadIOS	0 to 999999999	100	Report if buffer pool had more than this number of disk reads:Serious
BPReadIOW	0 to 999999999	0	Report if buffer pool had more than this number of disk reads:Warning
CFSTime	1 to 999999999	100	CF single response time in micro seconds
DeferredWriteTaskS	0 to 999999999	50	Number of DWT started-Serious
DeferredWriteTaskW	0 to 999999999	0	Number of DWT started-Warning
HighLogRateMB	1 to 9999	100	Specify high log rate in MB/Sec. A value larger than this will be reported as a high log rate. Many systems can achieve over 100MB/second. The value you get will depend on your hardware configuration.
OKLogRateMB	1 to 9999	50	OK log rate in MB/Sec
ReadLogBuffers	0 to 999999999	0	Report if the number of log records read exceeds this value

Thresholds for reporting out-of-line conditions in accounting data

Value	**Range**	**Default**	**meaning**
ApplicationLogMB	1 to 9999	1	Report in TASKCSV if application logged more than this value. Use this to select high use applications.
ApplicationCPUSec	1 to 9999	1	Report in TASKCSV if application used more than this CPU seconds. Use this to select high use applications. See BatchCPU etc below.
BPGetPGetN	1 to 999999999	20	Report if ratio of get Old pages:get New pages exceeds this. A large ratio can indicate queues are not indexed, or a lot of browse activity.
CFSTime	1 to 999999999	100	CF single response time in micro seconds
CFSyncRT	1 to 999999999	10000	Report in MESSAGE if synchronous CF request greater than this value, in microseconds
CommitET	1 to 999999999	10000	If the commit time is longer than this value, produce message MQTASK12S, MQTASK13E or MQTASK14W
DB2Time	1 to 999999999	100	If the average time for DB2 request in micro seconds is greater than this value, produce message MQQ5ST11W, MQQ5ST12W, MQQ5ST13W or MQQ5ST14W
DNSLookupTime	1 to 999999999	1000	Report if average time doing DNS lookup is longer than this value in microseconds.
ExitTime	1 to 999999999	10000	Report if the average ExitTime for a channel is longer than this value in microseconds.
HighLogRateMB	1 to 9999	100	Specify high log rate in MB/Sec. A value larger than this will be reported MQQJST10E in MESSAGE

Long_Open	50 to 999999999	1000	For IBM use only. Value to be considered a long open in micro seconds.
Long_Put	50 to 999999999	10000	If the average put times is greater than this value, produce message MQTASK08E or MQTASK09E .
LongLatchWait	50 to 999999999	1000	IBM use only. If the longest latch wait value is greater than this value report messages MQTASK15S or MQTASK15W
MaxDepth	0 to 999999999	0 (not used)	If queue depth was greater than this value produce message MQTASK10W.
NetTime	1 to 999999999	10000	Report if the average Nettime for a channel is longer than this value in microseconds.
OKLogRateMB	1 to 9999	50	OK log rate in MB/Sec
SSLDuration	1 to 999999999	10000	Report if the average duration for an SSL request a channel is longer than this value in microseconds.
BatchCPU	0 to 999999999	0 (not used)	Print message if CPU used by a Batch job is greater than this value in microseconds
ChannelCPU	0 to 999999999	0 (not used)	Print message if CPU used by a Channel is greater than this value in microseconds
CICSCPU	0 to 999999999	0 (not used)	Print message if CPU used by a CICS transaction is greater than this value in microseconds

Control what is displayed

Value	**Range**	**Default**	**meaning**
Debug			For IBM use only
Detail	0 to 20	5	Which level of detail 0 to 20. 20 provides the most detail
printHex			Debug value to print raw record in hex

IPLookup	0 to 1	0	displays the ip address for channels in format
			0 – with ip address in numbers
			1 – lookup the ip address to get the name
			for example
			`MQPV.MQPH 9.20.4.5`
			or
			`MQPV.MQPH winmvsca.ibm.com`

MQ_LC_ALL		En_GB.IBM-285	Locale to be use in printing numbers. For example using MQ_LC_ALL 'En_GB.IBM-285' gives channel accounting data as 50,292,900 instead of the default 50292900. With De_AT.IBM-924 this value would be displayed as 50.250.324 See XL `C/C++ XL C/C++ Programming Guide SC09-4765-12` for more information.

Processing Comma Separated Values

This version of the supportPac has experimental code to see if it is of value to the end user.

You can have the log data put out to files in the USS file system, with one file per queue manager. You can select which fields you want to display.

In the parameter input you can specify USS path where you want the files to be stored. For example

path /u/paice/

In the output file //OFIELDS is a list of fields you can have displayed in the CSV files. For example

```
Type.value                      (CSV heading) description
LOGCSV.DATE                     (Date) Date from SMF record
LOGCSV.TIME                     (Time) Time fom SMF record
LOGCSV.QMGR                     (QMGR) QMGR name from SMF record
LOGCSV.STIME                    (StartTime) Interval start time
LOGCSV.ETIME                    (EndTime) Interval end time
LOGCSV.DURATION                 (Duration) Interval duration
LOGCSV.PagesPerIO               (PagesPerIO) Overall pages written per IO
...
LOGIOCSV.DATE                   (Date) Date from SMF record
...
LOGIOCSV.WRITERATE              (logWriteRate) MB per second per log
LOGIOCSV.LOGBUSYPC              (IOPercentBusy) How busy the log was in %% of duration
LOGIOCSV.PagesPerIO             (PagesPerIO) Overall pages written per IO
LOGIOCSV.11Count                (11Count) Count of Writing 1 page to Log 1
LOGIOCSV.11CIs                  (11CIs) Count of single CIs written to Log 1
LOGIOCSV.11PPERIO               (11Pages/IO) Pages per I/O log 1
...
```

Specify the field names in

```
//FIELDS DD *
OGCSV.DATE                     (Date) Date from SMF record
LOGCSV.TIME                     (Time) Time fom SMF record
LOGCSV.QMGR                     (QMGR) QMGR name from SMF record

* LOGCSV.DURATION                 (Duration) Interval duration
LOGCSV.PagesPerIO               (PagesPerIO) Overall pages written per IO
/*
```

Where lines beginning with * are ignored on input

Records beginning withy LOGCSV are written to the logcsv file. This JCL and input specify the field names DATE, TIME, QMGR, STIME, ETIME, ~~Duration~~ and PagesPerIO, to be be written to the output file.

The ouput file name is by default logcsv.qmgr.csv in the specified path. There will be one file per queue manger, and the high level qualifier is specified in the Logcsvdd keyword, and defaults to logcsv.

On my system this produced file /u/paice/logcsv.MQ27.csv with content

```
Date,Time,QMGR,PagesPerIO,Wait_buffer,
2016/02/10,00:00:00,MQ27,  2,0,
2016/02/10,00:15:00,MQ27,  3,0,
2016/02/10,00:30:00,MQ27,  2,0,
```

and a file /u/paice/logcsv.MQPA.csv with similar content

Processing Comma Separated Values

Some of the output files have data in Comma Separated Values(CSV) format.

You can have the output go to MVS data sets, or files in OMVS.

For example to create a file in USS.

//DNSCSV DD PATH='/u/paice/DNS.csv',

// PATHOPTS=(OWRONLY,OCREAT),PATHMODE=(SIRWXU,SIRGRP)

This file can then be imported into a spread sheet.

You can use cut and paste to copy the data into a spread sheet. For example with Open Office you can use paste, and use the options, *Separated by Comma, Space* and *Merge Delimiters.*

Output from the MQSMF program

Chinit statistics

In //CHINIT the channel initiator control information block contains basic information for the CHINIT, including:

- CHINIT job name (*qcctjobn*)
- QSG name if it is in a QSG (*qcctqsgn*)
- Peak number of current channels (*qcctnocc*)
- Peak number of active channels (*qcctnoac*)
- MAXCHL - maximum permitted current channels (*qcctmxcc*)
- ACTCHL - maximum permitted active channels (*qcctmxac*)
- TCPCHL - maximum permitted TCP/IP channels (*qcctmxtp*)
- LU62CHL - maximum permitted LU62 channels (*qcctmxlu*)
- Storage used by CHINIT (*qcctstus*)

You can use this information to see if the number of active channels is approaching the configured maximum value. Note that the number of current and active channels are the values when the record was created. So, between the two intervals there might have been more than this number of channels active.

Channel information from SMF data

Here is an example of channel information from SMF data:

```
MVCA,MQPV,2014/03/18,13:00:00,VRM:800,
From 2014/03/18,12:45:00.015222 to 2014/03/18,13:00:00.083630 duration 900.068408 seconds
Peak number used of current channels...........   4
Peak number used of active channels ...........   0
MAXCHL. Max allowed current channels...........9999
ACTCHL. Max allowed active channels............9999
TCPCHL. Max allowed TCP/IP channels............9999
LU62CHL. Max allowed LU62 channels.............. 200
Storage used by Chinit......................... 436MB
```

You can monitor the storage usage and see whether the value is trending upwards. If the total used is approaching the total storage available, you might be running out of storage, and so might not be able to support many more channels.

If the numbers of active current channels are tending towards the maximum number of channels, you might need to increase the maximum number of channels.

CHINCSV

This provides the same information as above but in a CSV format.

The column headings are

```
mvs,qm,qsg,date,time,QSG,HWMCCHL,HWMACHL,MaxCHL,MaxChl,MaxAct,MaxTCP,MaxLU,StgMB
```

Adapter tasks

Example data

```
Task,Type,Requests,Busy %,     CPU used, CPU %,"avg CPU","avg ET"
```

```
   ,    ,        ,      ,      Seconds,      , uSeconds,uSeconds
   0,ADAP,  470297,  10.2,    41.290670,  4.6         88,     194
   1,ADAP,   13907,   0.6,     1.589428,  0.2        114,     365
   2,ADAP,    2517,   0.2,     0.185325,  0.0         74,     746
   3,ADAP,    1095,   0.1,     0.085774,  0.0         78,     907
   4,ADAP,     535,   0.1,     0.040743,  0.0         76,     947
   5,ADAP,     220,   0.0,     0.016228,  0.0         74,    1175
   6,ADAP,      82,   0.0,     0.005521,  0.0         67,    1786
   7,ADAP,      80,   0.0,     0.004248,  0.0         53,    1160
Summ,ADAP,  488733,   1.4,    43.217938,  0.6,        88,     205
```

The fields are calculated from:

- Duration: qwhs.qwhsdurn
- Requests: qctreqn
- Busy %: qcteltm and duration
- CPU used: qctcptm
- CPU %: qctcptm and duration
- Average CPU: qctcptm and qctreqn average
- ET: qcteltm and qctreqn

This example shows that there were eight adapter tasks.

Adapter number 0

- Processed the majority of the requests (470297 out of 488733)
- Was busy 10.2% of the interval
- Used 41.3 seconds of CPU

Overall

The average CPU per request was 88 microseconds of CPU and took 205 microseconds

The times are in STCK values. To convert the data in the records to microseconds divide the value by 4096 (which removes the bottom 12 bits).

The adapters process MQ requests. Some of these requests might wait, for example, for log I/O during a commit, so the average Elapsed Time per request has little meaning.

When a MQ request is made the first free adapter task is used.

- If there is at least one adapter that has been little used (less than 1% busy), you have enough adapters.
- If at least one adapter was not used, you have enough adapters defined.
- If all the adapters were used, you might need to allocate more adapters.
- If all of the adapters were used, and they were all busy for most of the interval, you need to allocate more adapters.

You can use the ALTER QMGR CHIADAPS() command to change the number of adapters used. Any changes come into effect the next time the CHINIT is started.

Attention: If there are too many adapters acting on a small set of queues, you might get contention within the queue manager.

Understanding the data

Usually the first adapter processes most of the requests, and the other adapters process fewer and fewer requests. The first few adapters may be busy for most of the period (close to 100 %busy).

In a system constrained for CPU the data looks different for example

```
Task,Type,Requests,Busy %,     CPU used, CPU %,"avg CPU","avg ET"
    ,    ,        ,      ,      Seconds,      , uSeconds,uSeconds
   0,ADAP,   30005,  67.1,     2.802067,  6.0,        93,    1051
   1,ADAP,   27121,  64.1,     2.561604,  5.5,        94,    1111
   2,ADAP,   25668,  61.7,     2.418879,  5.1,        94,    1131
   3,ADAP,   24154,  58.5,     2.269495,  4.8,        94,    1138
   4,ADAP,   22698,  55.2,     2.127903,  4.5,        94,    1142
   5,ADAP,   20906,  51.4,     1.956664,  4.2,        94,    1156
   6,ADAP,   19094,  47.1,     1.781381,  3.8,        93,    1160
   7,ADAP,   16827,  42.3,     1.568889,  3.3,        93,    1181
Summ,ADAP,  186473,  55.9,    17.486883,  4.7,        94,     100
```

Here, none of the task are close to 100% busy.

The adapter task records how busy it is by calculating the start and end times of the request. If there is a delay in z/OS dispatching the TCB then the adapter task will appear busy. So it looks like the first adapter was not dispatched 33 % of the time (100-67.1%).

If you see this pattern of data then check the chinit address space is not constrained for CPU. Also check the dispatcher data. See below.

Adapter CSV

This is the same information as the adapter section, but in CSV format, which can imported into a spread sheet

Dispatcher tasks

Example data

```
Task,Type,Requests,Busy %,     CPU used, CPU %,"avg CPU","avg ET"
    ,    ,        ,      ,      Seconds,      , uSeconds,uSeconds
   0,DISP,   26587,   0.4,     0.592463,  0.1         22,     127
   1,DISP,   26963,   0.3,     0.588092,  0.1         22,     112
   2,DISP,  864329,   2.7,     2.545668,  0.3          3,      28
   3,DISP,   26875,   0.4,     0.590825,  0.1         22,     120
   4,DISP,   26874,   0.4,     0.603285,  0.1         22,     123

Summ,DISP,  971628,   0.8,     4.920332,  0.1,         5,      38
```

The example data shows that there were five dispatchers. A channel is associated with a dispatcher, and channels are distributed across all the dispatchers. This example shows that one dispatcher is processing more requests than other dispatchers. This is normal, as some channels might stop, so the dispatcher is processing fewer channels, and some channels can be busier than others.

- 4.9 seconds of CPU were used by the dispatchers.
- The average request used 5 microseconds of CPU and took 38 microseconds elapsed time.
- The average CPU used per request depends on the message traffic, for example, bigger messages use more CPU than smaller messages.

The fields are calculated from:

- Duration: qwhs.qwhsdurn
- Requests : qctreqn
- Busy %: qcteltm and duration
- CPU used: qctcptm
- CPU %: qctcptm and duration
- Average CPU: qctcptm and qctreqn
- Average ET: qcteltm and qctreqn

Usually, the number of dispatchers should be less than, or equal to, the number of processors in the LPAR. If you have more dispatchers than processors in the LPAR they might compete for CPU resources.

Channels have an affinity to a dispatcher, so you might find that some dispatchers process many more requests than another dispatcher.

You can use the ALTER QMGR CHIDISPS() command to change the number of dispatchers used. Any change comes into effect the next time the CHINIT is started.

Understanding the data.

A dispatcher is used to send and receive data over a communications network, and this is not usually dependent on external events. The average elapsed time should, therefore, be close to the average CPU time used. If the CHINIT is delayed due to lack of CPU, then the ratio of average Elapsed Time to average CPU time is much larger, compared to when the CHINIT is not delayed for CPU.

In a test systems which was constrained for CPU the elapsed time was more than 10 times the CPU used.

Dispatcher CSV

This is the same information as the dispatcher section, but in CSV format, which can imported into a spread sheet

Domain Name Server (DNS) task

The data is similar to the adapter and dispatcher tasks, but there is additional information on the longest request at the end of the record.On the TCB for the DNS requests there is

```
longest ,date      ,time
uSeconds,          ,
     463,2014/03/18,12:56:33.987671
```

The longest DNS request took 463 microseconds elapsed time, and this occurred at 12:56:33 local time.

The fields are calculated from:

- Duration: qwhs.qwhsdurn
- Requests : qctreqn
- Busy %: qcteltm and duration
- CPU used: qctcptm
- CPU %: qctcptm and duration
- Average CPU: qctcptm and qctreqn
- Average ET: qcteltm and qctreqn
- Longest: qctlgdu
- Longest at: qctlgtm

The DNS is used when a channel starts, to convert an ouput connection name to an IP address, and from an input IP address to a connection name.

The DNS task can go out of your enterprise to look up a value. If the average elapsed time is significantly more than the average CPU time used, you might have some long requests.

If the value of the longest request time is unacceptable you should work with your network team to investigate why you are having long requests. It might be that you have an invalid name in your conname.

If the DNS task is busy for 25% of the duration, consider investigating the cause further.

Note: There are requests to the DNS task that are not DNS lookups, so you might have the number of requests being greater than zero – but no longest request information.

Domain Name Server (DNS) CSV

This is the same information as the DNS section above, but in CSV format, which can imported into a spread sheet.

SSL

This is the same information as the Dispatcher section above. It has additional information for each task.

```
Task,... longest ,date      ,time
    ,... uSeconds,          ,
   0,... 8864,2014/03/18,12:46:40.237697
   1,... 4714,2014/03/18,12:46:18.938022
   2,... 7273,2014/03/18,12:46:35.358145
   3,... 13164,2014/03/18,12:46:44.514045
   4,...22438,2014/03/18,12:46:22.134123
Summ,...22438,2014/03/18,12:46:22.134123
```

The longest request was for SSL task 4, took 22,438 microseconds, and occurred at 12:46:22.134123 local time. The overall longest value is displayed in the Summ record.

The fields are calculated from:

- Duration: qwhs.qwhsdurn
- Requests : qctreqn
- Busy %: qcteltm and duration
- CPU used: qctcptm
- CPU %: qctcptm and duration
- Average CPU: qctcptm and qctreqn
- Average ET: qcteltm and qctreqn
- Longest: qctlsdu longest at: qctlstm

A running channel is associated with an SSL task, in a similar way that a channel is associated with a dispatcher. The SSL tasks can use the cryptographic coprocessors available to the LPAR. So, the elapsed time can include time spent on a coprocessor. You should monitor the average elapsed time throughout the day. If this time increases significantly during peak periods you should work with your MVS systems programmers, as your coprocessors might be over utilized.

If the SSL tasks are busy for a significant proportion of the interval, increasing the number of SSL tasks might help. If the SSL tasks are waiting for external resources such as a coprocesor, increasing the number of SSL tasks may have little effect.

You can use the ALTER QMGR SSLTASKS() command to change the number of SSL tasks used. Any changes come into effect the next time the CHINIT is started.

SSLCSV

This record has the information from the Summ section of the SSL records above

Channel accounting

This information is based on the output from Display Channel Status.

```
Connection name               winmvsca.hursley.ibm.com
Connection name               9.20.4.159
Channel disp                  PRIVATE
Channel type                  SENDER
Channel status                RUNNING
Channel STATCHL               HIGH
Remote qmgr/app               MQPV
Channel started date & time  2014/06/04,08:05:50
Channel status collect time  2014/06/04,08:08:16
Last msg time                2014/06/04,08:07:50
Active for                       145 seconds
Batch size                                50
Messages/batch                          48.8
Number of messages                    33,896
Number of persistent messages         33,896
Number of batches                        695
Number of full batches                   601
Number of partial batches                 94
Buffers sent                          33,992
Buffers received                         697
Xmitq empty count                          0
Message data                        120,263,008   114 MB
Persistent message data             120,263,008   114 MB
Non persistent message data                 0     0  B
Total bytes sent                    120,266,176   114 MB
Total bytes received                     19,996    19 KB
Bytes received/Batch                        28    28  B
Bytes sent/Batch                       173,044   168 KB
Batches/Second                               4
Bytes received/message                       0     0  B
Bytes sent/message                       3,548     3 KB
Bytes received/second                      137   137  B/sec
Bytes sent/second                      829,421   809 KB/sec
Compression rate                             0
Exit time average                            0 uSec
Net time average                           116 uSec
Net time min                                71 uSec
Net time max                            14,481 uSec
Net time max date&time                 2014/06/04,08:07:41
```

Field name	**Example value**		`Field name`
Connection name	winmvsca.hursley.ibm.com		Output of the C run time function inet_pton with qcstcnnm
Connection name	9.20.4.159		qcstcnnm
Channel disp	PRIVATE		qcstchdp values MQCHLD_*
Channel type	SENDER		qcstchty values MQCHT_*
Channel status	RUNNING		qcstchst values MQCHS_

Channel STATCHL	HIGH		qcststcl values MQMON_*
Remote qmgr/app	MQPV		qcstrqmn
Channel started date & time	2014/06/04,08:05:50		qcststrt
Channel stopped time			qcstludt
Channel status collect time	2014/06/04,08:08:16		qcstcltm
Last msg time	2014/06/04,08:07:50		qcstlmst
Active for	145 seconds		Start time = later of channel start time qcststrt, and interval start time End time = earlier of channel ended time qcstludt, and end of interval qwhsstck
Batch size	50		qcstcbsz
Messages/batch	48.8		qcstnmsg/qcstbatc
Number of MQI requests (for svrconn)			qcstnmsg
Number of messages	33896		qcstnmsg
Number of persistent messages	33896		qcstnpmg
Number of batches	695		qcstbatc
Number of full batches	601		qcstfuba
Number of partial batches	94		qcstbatc-qcstfuba
Buffers sent	33,992		qcstbfst
Buffers received	697		qcstbfrc
Xmitq empty count	0		qcstqetc
Current shared connections			qcstcscv
Message data	120263008	114 MB	qcstnbyt
Persistent message data	120263008	114 MB	qcstnpby
Non persistent message data	0	0 B	qcstnbyt - qcstnpby
Total bytes sent	120266176	114 MB	qcstbyst
Total bytes received	19996	19 KB	qcstbyrc
Bytes received/Batch	28	28 B	qcstbyrc/ qcstbatc
Bytes sent/Batch	173044	168 KB	qcstbyst/ qcstbatc
Batches/Second	4		qcstbatc/active time
Bytes received/message	0	0 B	qcstbyrc/ qcstnmsg
Bytes sent/message	3548	3 KB	qcstbyst/ qcstnmsg
Bytes received/second	137	137 B/sec	qcstbyrc/active time
Bytes sent/second	829421	809 KB/sec	qcstbyst/active time
Compression rate	0		qcstcpra
Exit time average	0 uSec		qcstetav
Exit time min			qcstetmn

Exit time max			qcstetmx
Exit time max date&time			qcstetmx
Net time average	116 uSec		qcstntav
Net time min	71 uSec		qcstntmn
Net time max	14,481 uSec		qcstntmx
Net time max date&time	2014/06/04,08:07:41		qcstntdt
Put retry count			qcstptrc, see below

Note: qcstbyst and qcstbyrc are the values from the Display Channel Status command. The values in these fields wrap at 999,999,999. You can get the true number of bytes sent by adding 999,999,999 to this value until the value exceeds the qcstnbyt value.

The Put retry count value is the number of times a receiver channel tried to put a message to a queue and had problems.

For example

```
+CSQX038E MQPV CSQXRESP Unable to put message to MPQAZ, MQCC=2 MQRC=2192 (MQRC_PAGESET_FULL)
+CSQX565E MQPV CSQXRESP No dead-letter queue for MQPV, channel MQPH.MQPV
```

The put request is retried. See channel attributes MRTMR and MRRTY for how many times and how often it retries. During this time the channel is paused and is not processing any messages.

Channel accounting CSV

This is the same information as the Channel section, but in CSV format, which can imported into a spread sheet. The columns are

MVS	z/OS system name
MQ	Queue Manager name
date	Date the record was created
time	time the record was created
channelType	Sender etc
channelName	
BSZ	Negotiated batch size
ABSZ	Achieved batch size
Bytes/sec	

Channel summary CSV

This provides a summary of the activity for different channel types.

For example

```
MVS,MQ,date,time,VRM,channelType,count,Persistent,NonPersistent,'P/Sec','NP/Sec'
MVCA,MQPV,2014/06/30,11:30:00,VRM:800,RECEIVER,2,75720,0,3786,0
MVCA,MQPV,2014/06/30,11:30:00,VRM:800,total,2,75720,0,3786,0
MVCA,MQPH,2014/06/30,11:30:00,VRM:800,SENDER,2,75720,0,2611,0
MVCA,MQPH,2014/06/30,11:30:00,VRM:800,total,2,75720,0,2611,0
MVCA,MQPH,2014/06/30,11:34:04,VRM:800,SENDER,23,86237508,0,559983,0
MVCA,MQPH,2014/06/30,11:34:04,VRM:800,total,23,86237508,0,559983,0
```

Chinit Messages(CMESSAGE)

This file contains information about the Chint SMF Data. Checks are made with user specified values, such as SSL Duration. See the value that can be specified in Record selection parameters on page 12.

The number in parenthesis () before the message is when the message is displayed. The message is displayed if the value of Detail in the input parameters is >= the value in ().

The messages text is a paraphrase of the message, to make it easier to read.

(0) MQCHIN000S	The high water mark of the number of active channels >90 % of max channels.
(10) MQCHIN001W	The high water mark of the number of active channels >50 % of max channels.
(10) MQCHIN001I	The high water mark as a percentage of the max channels.
(10) MQCHIN006S	Longest SSL request >> specified SSL duration (SSLDuration).
(0)MQCHIN011S	High water mark of active channels nn% > 90% of max active channels.
(10)MQCHIN003W	High water mark of active channels nn% > 50% of max active channels.
(10)MQCHIN013I	High water mark of active channels nn% max active channels.
(0) MQCHIN007I	Dispatcher task is nn% busy on average .
(0) MQCHIN008I	Adapter task is nn% busy on average.
(0) MQCHIN009I	SSL task is nn% busy on average.
(0) MQCHIN006S	Longest SSL request (n uSeconds) >> specified SSL specified duration(nn uSeconds). The longest SSL request was greater than 10 time the value of SSLDuration .
(0) MQCHIN005W	Longest SSL request (n uSeconds) > specified SSL specified duration(nn uSeconds). The longest SSL request was greater than 10 value of SSLDuration .
(0) MQCHIN007I	DNS task is nn% busy on average.
(0) MQCHIN004S	Longest DNS request (n uSeconds) >> specified DNS duration(nn uSeconds). The longest DNS request was greater than 10 time the value of DNSLookupTime. Investigate why the DNS request took so long. An invalid address may have been specified, and the request may have been routed out of your domain.
(0) MQCHIN005W	Longest DNS request (n uSeconds) specified SSL specified duration(nn uSeconds).

Page set statistics

The page set statistics gives information similar to that available with the `DIS USAGE` command, as well as some performance indicators.

For the details about the fields see the layout for qis1 in member SCSQC370(CSQDSMFC).

The records are written to ddname //PSET.

The information about size and number of pages in use is printed in CSV format suitable for importing into a spread sheet in ddname //PSETCSV. See see Page set page CSV data on page 31.

The data is also printed in a slightly different format, by buffer pool and page set, see Page set statistics by buffer pool on page 31.

If the I/O has been done to the page set data is printed in CSV format in //PSIOCSV, see see Page set I/O CSV data on page 31.

Page set statistics report

```
MV41,MQ03,2014/11/26,11:01:00,VRM:900,
  From 2014/11/26,11:00:00.000000 to 2014/11/26,11:01:01.000000, duration    61 seconds.
PS00 BP 0, Pages 1798, Size 7192 KB, free 86.0%, used 14.0%, P 14%, NP 0%, #Full 0,
     Pages written in checkpoint           126
     Pages written in not in checkpoint      1
     Number of stripes                       1
     Expansion type:User expansions          4
     <Page set expansion occurred>
PS00 Type :I/O requests,   Pages, Avg I/O time, pages per I/O, MB/Sec,  busy%
PS00 Write:       46469, 1487008,         5394,          32.0,     23,  4102%
PS00 IMW  :          17,      17,          887,           1.0,     23,     0%
PS00 GET  :        7477,    7477,         1109,           1.0,     23,    13%
```

Where

PSID 0	Page set number (qis1psno)
BP 0	Buffer pool number (qis1bpno)
Pages 1798	Number of pages in the page set (qis1topg)
Size 7192 KB	size of page set in KB, MB or GB (qis1topg * 4096)
free 86.0%	How much of the page set was free when the data was collected (100* qis1unpg/qis1topg)
used 14.0%	How much of the page set was in use when the data was collected (100 - free)
P 14%	The percentage number of pages with Persistent data on them when the data was collected (100 * qis1ppg / qis1topg)
NP 0%	The percentage number of pages with Non Persistent data on them when the data was collected (100 * qis1nppg / qis1topg)
#Full 0,	Cumulative number of times the page set had page set full.

Pages can be for

- Persistent data
- Non persistent data
- Internally used to manage the messages and space in the page set

The % of Persistent + % of Non Persistent will usually be less than the used % value.

`Pages written in checkpoint          126`

This is how many pages were written to the page set due to checkpoint activity (qis1ckpp).

`Pages written not in checkpoint      1`

This is how many pages were written to the pageset not due to checkpoint activity (qis1dwpp + qis1imwn -qis1ckpp)).

If this number is small (less than the number of pages written in checkpoint) this can be ignored. There are often a small number of pages written outside of checkpoing, for example, if a page is in use during checkpoint processing then when the page is released, it may be immediately written to disk.

If the number is large, then this indicates the buffer pool was too small, and pages had to be written out to free up space in the buffer pool.

```
Number of stripes                    1
```

This field tells you how many stripes the page set has (qis1strn).

If you have a lot of disk I/O then you may get performance benefits from striping the pageset[1]. When the buffer pool fills up, an internal task, the deferred write processor(DWP), writes pages out to the buffer pool. In MQ Version 8, DWP can write up to 32 pages in an I/O request. Striping the page set should reduce the I/O time, and so pages should be freed up faster.

```
Expansion type:User   expansions        4
```

The Expansion type value is USER, SYSTEM, NONE (qis1expn).

The expansions is the number of times the page set has expanded since the queue manager was started.

```
Page set expansion occurred
```

This text is displayed if page set expansion occurred at least once during this internal (qis1expf).

I/O to the page set

```
PS00 Type :I/O requests,   Pages, Avg I/O time, pages per I/O, MB/Sec,  busy%
PS00 Write:       46469, 1487008,         5394,          32.0,     23,   410%
PS00 IMW  :          17,      17,          887,           1.0,     23,     0%
PS00 GET  :        7477,    7477,         1109,           1.0,     23,    13%
```

This section has the I/O statistics for the page sets.

PS00	The page set number (00).
Write	This is the number of pages written by DWP. DWP processing is performed at checkpoint, or when the buffer pool's stealable (immediately usable) pages falls below a threshold.
IMW	The buffer pool was close to being full, and request to use buffer pool pages are synchronously written to the page set. For persistent messages, this will also cause a write and wait to the log.
Get	When a request needs a page and the page is not in the buffer pool, it has to be read from the page set.

The fields for each type are

I/O requests	Number of I/O requests (qis1dwpn, qis1imwn, qis1getn)
Pages	Total number of pages moved between buffer pool and pageset by the operation type. IMW and Get process one page per I/O, DWP writes multiple pages per I/O).
Avg I/O time	This time is (total I/O time)/(number of I/O requests). I/O requests with multiple pages will take longer than a request with only one page. Concurrent I/O to the same track on DASD can cause I/O delays. ((qis1+++t/4096) [2] / qis1+++n)
pages per I/O	This is the average number of pages per I/O request. (qis1+++p / qis1+++n)
MB/Sec	This is the number of 4KB pages * 4096 / I/O time. (qis1+++n * 4096 / (1024 * 1024) / ((qis1+++t/4096)^2/1000000))
busy	This is the I/O time/ SMF duration. This value can be > 100% if there are concurrent requests to the page set. Where it says busy 410% there are at least five requests in parallel. (100 * qis1+++t/4096^2 / qwhsdurn)

[1] Striping at the MVS level, so multiple channels can be used, rather than striping within a DASD subsystem

[2] qis1+++t/4096 to convert from STCK format to microsecond

Page set statistics by buffer pool

For each buffer pool, this reports the page set I/O statistics above, within each buffer pool. This is written to //BUFFIO.

```
MV41,MQ03,2014/11/26,11:01:00,VRM:900,
  From 2014/11/26,11:00:00.000000 to 2014/11/26,11:01:01.000000, duration    61 seconds.
BP   PSID Type :I/O requests,   Pages, Avg I/O time, pages per I/O, MB/Sec,  busy%
BP00   00 Write:       46469, 1487008,         5394,          32.0,     23,  410%
BP00   00 IMW  :          17,      17,          887,           1.0,     23,    0%
BP00   00 GET  :        7477,    7477,         1109,           1.0,     23,   13%

BP00   04 Write:           1,       2,          636,           2.0,     12,    0%
BP00      Total:       53963, 1494504,         5372           27.6,     23,  425%
```

The fields are the same as above.

You can have multiple page sets within a buffer pool, and the data is summarised for the buffer pool. You can identify the buffer pools with a lot of page set activity by editing the file, and using the following ISPF editor commands

X ALL	exclude all the lines in the file
f ' ' 9 10 all	find the blanks in the page set column
top	go to the top
f3	use the line command to display the first 3 lines of the file with the column header information
delete all x	to remove the excluded lines
sort 31 40 d	to sort on the number of pages

At the top of the file will be the buffer pool with the largest number of pages moved between buffer pool and page set.

Page set page CSV data

The information about size and number of pages in use is printed in CSV format suitable for importing into a spread sheet in ddname //PSETCSV. The column headings are *z/OS, QM, Date, Time , StatDur ,PS ,BP ,pages ,MB , %free, %used, %P, %NP, #full* . Where the column names are the same as described above.

Page set I/O CSV data

If there has been I/O to the page set then data will be written to //PSIOCSV. The columns are *z/OS, QM, Date, Time, StatDur, PS, BP, wIO, wPages, wAvgt, iIO, iPages ,iAvgt, gIO, gPages, gAvgt, stripes,! chkptp, chkptp, expands.* Where

z/OS	The z/OS system name
QM	The Queue manager name
Date	the date when the record was created
Time	The time the record was created
StatDur	The duration of the interval in seconds
PS	The page set (qis1psno)
BP	The buffer pool (qis1bpno)
wIO	The number of write IO requests for DWP (qis1dwpn)
wPages	The number of pages written for DWP (qis1dwpp)
wAvgt	The average IO time for write requests for DWP (qis1dwpt)
iIO	The number of immediate write requests (IMW) (qis1imwn)
iPages	The number of pages written for immediate write (IMW) (qis1imwp)
iAvgt	The average IO time for immediate write requests (IMW) (qis1imwt)
gIO	The number of read request (qis1getn)
gPages	The number of pages got for the read requests (qis1getp)

gAvgt	The average IO time for read requests (qis1gett)
stripes	The number of stripes used by the log (qis1strn)
chkptp	The number of pages written in checkpoint (qis1ckpp)
! chkptp	The number of pages written outside of checkpoint (qis1dwpp + qis1imwn - qis1ckpp)
expands	The number of page set expansions(qis1expn).

Additional messages produced in MESSAGE output file.

Message: MQPSET00W ... Page set *m* status *sss*
When: Page set status is one of Defined,Offline,Not defined or Suspended.

Action: Investigate why the page set is not available.

Message: MQPSET01S ... Page set *m* was full *n* times
Detail: 1

When: qis1full $>$ 0

Reason: The page set filled up.

Action: Investigate why the page set is filling up.

Depending on the setting the Page set *m*ay expand.

You can prevent any future expansions using the `ALTER PSID(m) EXPAND(NONE)` command.

Message: MQPSET02W ... Page set *m* expansion occurred. Current expansion count *n* times
When: Bit qis1expf is on.

Action: Investigate why the page set filled up.

Message: MQPSET03E ... Page set *m* I/O for writing to page set was *n%* busy
When: The total I/O time for page set writes was greater that the value in the *PSETbusyPC* parameter.

Message: MQPSET04E ... Page set *m* I/O for Immediate Writing to page set was *n%* busy
When: The total I/O time for page set Immediate writes was greater that the value in the *PSETbusyPC* parameter.

Message: MQPSET05E ... Page set *m*, I/O for getting from page set was *n%* busy
When: The total I/O time for reading from the page set was greater that the value in the *PSETbusyPC* parameter.

Message: MQPSET05E ... Page set *m*, Total I/O time to page set was *n%* busy
When: The total I/O time for the page set was greater that the value in the *PSETbusyPC* parameter.

Page set Queue I/O

The Page set Queue I/O reports on the queues which had reads or writes to the page set. It processes the Class(3) accounting records The information is accumulated and reported when the buffer pool stats for that queue manager are reported. So if there were multiple application instances getting from queue, and the gets had page set I/O, there will be one records with the accumulated data, for the Queue Manager, buffer Pool, page set and queue combination.

You can use this section

1. To identify the queues which a large number of pages read from the page set. You can use this information to help you keep the queue depth low by starting your getting applications earlier
2. To review queue and page set mapping. If you have multiple queues on the same page set with a large number of gets from the page set you should consider moving a queue to a different page set to split the disk I/O.
3. To identify buffer pool problems. If a queue has a large number of pages written, then the buffer pool has been more than 95% busy and the requests for pages have been written to disk asynchronously.

The information is written to ddname //PSIDQIO. When a buffer pool statistics record is processed, it calls Page Set Queue I/O routine to report on the I/O records for that queue manager. When the end of the SMFIN file has been detected, any remaining records are printed out but with ... instead of the time.

```
                    QMGR, BP, PS,          Count,  Total_time,Avg_time,   Rate,Queue
HUR1,MQ01,2014/11/03,16:00:00,VRM:710,
2014/11/03,16:00:00,MQ01,006,057,Read,     24462,    12864391,     525,     27,PAYROLL
2014/11/03,16:00:00,MQ01,006,023,Read,     45838,    17367589,     378,     51,ACCOUNTS
2014/11/03,16:00:00,MQ01,006,019,Write,        2,        1065,     532,      0,ACOUNTS
Last entries
.................,MQ01,001,000,Write,        1,        1147,    1147,      0,MYLOG
.................,MQ01,004,006,Read,       371,      140479,     378,      6,ACCOUNTS
```

Where

```
HUR1,MQ01,2014/11/03,16:00:00,VRM:710,
```

is the time from the buffer pool record. The time portion of this used prefixed to each record to facilitate sorting of the data.

```
MQ01,006,057,Read,     24462,    12864391,     525, 27,PAYROLL
```

MQ01	Is the name of the queue manager
006	Is the buffer pool number
057	Is the page set
24462	Is the number of requests to read a page from the page set
2864391	Is the total I/O time
525	Is the average time to read a page
27	Is the rate of pages per second. This is calculated from the number of request/duration of the statistics interval
PAYROLL	Is the queue name

```
Last entries
```

This is the start of the remaining records which have not been printed out by a buffer pool record.

Ways of using the data

If you sort by the total time column you will see the queues which have the biggest delay. You can then use the date and time to see why the queue has so much I/O.

With MP1B you can select which queue to use for example *Queue PAYROLL* in //SYSIN to display the detailed task records for that queue.

If you find you have several queues with high I/O activity, you may want to consider moving a queue to a different page set. You may want to consider putting queues with a very large number of gets from the page set in their own buffer pool.

Coupling Facility statistics

The CF statistics give information about the Coupling Facility usage.

When a message is put or got from a queue the request to the CF has a single update. During a commit, the request to the CF may change several messages, so this is counted as a multiple request.

The data is displayed in the ddname //CF.

See the record layout in SCSQMACS(CSQDQEST) for interpretation of these fields.

```
MVCA MQ7B 2013/01/08 11:25:44 VRM:710
  APP1        , Structure-fulls      0
    Single     1000, avg et in uS    14, Retries      0
    Multiple     24, avg et in uS  1185, Retries      0
    Max entries      1033,  Max elements      2048
```

Where

APP1	The name of the structure
Structure-fulls	The number of times the structure filled
Single	The number of requests where there was a single request in the CF request
avg et in uS	The average elapsed time where there was a single request in the CF request, in microseconds
Retries	The number of times a request was retried
Multiple	The number of requests where there were multiple request in the CF request
avg et in uS	The average elapsed time where there were multiple request in the CF request, in microseconds
Retries	The number of times a request was retried, where there were multiple request in the CF request
Max entries	The maximum number of entries used in the structure
Max elements	The maximum number of elements used in the structure.

Additional messages produced in MESSAGE output file.

Message: MQQEST00E QEST ... structure sss full n times When: qestsful>0

Reason: The CF Structure has reached its capacity.

Action: Investigate to see if this is a short term problem, or a longer term problem. If this is a long term problem, you will need to increase the size of the CF structure.

Message: MQQEST01S QEST ... structure sss extremely long average CF response time n uS When: The average CF Single response time is > 100 * value specified in CFStime parameter.

Reason: The average response time of the single requests is taking a long time.

Action: Review the performance of the Coupling Facility. For example there may be a remote Coupling Facility. You may have specified a value of CFStime which is unrealistic.

Message: MQQEST02E QEST ... structure sss very long average response time n uS When: The average CF Single response time is > 10 * value specified in CFStime parameter.

Reason: The average response time of the single requests is taking a long time.

Action: Review the performance of the Coupling Facility. For example there may be a remote Coupling Facility. Determine if this occurred a time of peak workload.

You may have specified a value of CFStime which is unrealistic.

Message: MQQEST03W QEST ... structure sss long average response time n uS When: The average CF Single response time is greater than the value specified in CFStime parameter.

Reason: The average response time of the single requests is taking a long time.

Action: Review the performance of the Coupling Facility. For example there may be a remote Coupling Facility. Determine if this occurred a time of peak workload.

You may have specified a value of CFStime which is unrealistic.

It is acceptable to get message MQQEST03W, but you can use it as a warning if the structure response time increases.

Coupling Facility CSV

The CF statistics give information about the Coupling Facility usage.

The data reported in the ddname //CFCSV is Coupling Facility in a single line, n comma separated values.

```
MVS,QM,DATE,TIME,Structure,Full,'Max entries','Max elements','avg S','avg M','Num S','Num M'
MVCA,MQ7A,2013/02/09,14:18:39,CSQ_ADMIN   ,     0,   345,   390,925,1287,1999,5
MVCA,MQ7A,2013/02/09,14:18:39,APP1        ,     0,    35,    68,1027,1112,3000,999
```

Where the fields are

MVS	The MVS system ID
QM	The queue manager name
DATE	The date in YYYY/MM/DD format
TIME	The time in hh:mm:ss
Structure	The name of the structure
Full	The number of times the structure was full
Max entries	The maximum number of entries in time period
Max elements	The maximum number of elements in the time period
avg S	The average CF response time for single requests
avg M	The average CF response time when there are multiple data requests in a single CF request
Num S	The number of single requests
Num M	The number of requests when there are multiple data requests in a single CF request.

Shared Message Data Set Statistics (SMDS)

Shared Queue messages can be offloaded from the coupling facility to Shared Message Data Set(SMDS).

The output is displayed in the ddname //SMDS. For the details about the fields see SCSQMACS(CSQDQESD).

Data is written to CSV files, see page 41.

SMDS has a lower CPU cost, and higher throughput than storing the messages in DB2.

When SMDS is used, each queue manager in the QSG has its own dataset for storing messages, and has read access to the data sets from the other queue managers in the QSG. There is zero or one SMDS data set per queue manager for a CF Structure.

When a message is put, and it goes to SMDS, then a buffer is used to write to the SMDS. The buffer is freed after the I/O has completed and before the put request returns to the application. So the buffer is used only for the duration of the I/O request. The number of buffers available for use can be configured. The default buffer size is 256KB. Once set, the buffer size cannot be changed. When messages are put which are larger than this, multiple buffers are used. An application may have to wait for a buffer, and will have to wait if the request requires I/O. During a get the data may already be in a buffer. In this case there is no I/O to the SMDS.

Multiple small messages can be put into a buffer (if there is enough space). In this case the buffer may actively be doing I/O to the data set for one message, and there is a request to process another message. In this case the buffer is recorded as being busy, for example *Waited for busy buffer*.

The highlighted lines in the report below are used as headings in the following sections.

```
MVCA,MQ7A,2013/02/23,15:49:21,VRM:710,
 CF manager shared message data set (SMDS) statistics
  Structure :   2,  Name APP1
    SMDS space management statistics:
    SMDS space management usage:
      Messages in data set          27827        highest      27827
      Total blocks                  22914
        Space map blocks                1
        Message data blocks         22913
        Data blocks used            11306 ( 49%) highest      11306 ( 49%)
        Data blocks free            11607 ( 51%) lowest       11607 ( 51%)
    SMDS space management activity:
      Action        Messages     4K pages
      Allocated        27827       723502
      No space             0
      Released             0            0
      Reallocated          0            0
      Cleaned up           0            0
SMDS buffer pool statistics:
  SMDS buffer pool usage:
    Buffer size (DSBLOCK)           256K
    Total buffers                     1
    Buffers in use                    1 (100%)  highest      1 (100%)
      Shared buffers                  1
    Buffers free                      0 (  0%)  lowest       0 (  0%)
      Saved buffers                   0
      Empty buffers                   0
    Waiting request queues
      For free buffer                 1         highest      1
      For busy buffer                 0         highest      0
  SMDS buffer pool activity:
    Acquired buffers              39131
      Got valid buffer                      0 (  0%)
      Got matching, empty buffer            0 (  0%)
```

```
      Got free, empty buffer                1 (  0%)
      Stole a saved buffer              39130 ( 99%)
    No buffer available           11305
    Waited for free buffer         6217 ( 16%) avg time 0.004568s
    Waited for busy buffer            0 (  0%) avg time 0.000000s
    Buffer read issued             1406
      Data already valid              0 (  0%)
      Data partly valid               0 (  0%)
      Data read from disk          1406 (100%)
    Freed valid buffer            39130
    Marked buffer deleted             0
    Buffer write issued           39130
SMDS I/O statistics:
  SMDS data set usage:
    High allocated CI           1466496
    High formatted CI           1466496
    Control interval size          4096
    Control area size             589824
  SMDS I/O activity:

    Type        Requests 4K pages pages/req avg I/O time avg wait time MB/Sec %busy
    Format             0        0     0.0       0.000000s     0.000000s      0      0
>W  Write          39130   723467   18.5       0.000676s     0.000656s    106     21%
>L  Read (local)    1406    26000   18.5       0.000438s     0.000336s    164     32%
>O  Read (Other)   39107   723034   18.5       0.000663s     0.000646s    108     22%
```

Where the records in the output file are described below. The highlighted lines in the report are headings below.

`Structure : 2, Name APP1`

This identifies the structure.

Field name used: QESDSTRN, QESDSTR.

SMDS space management statistics:

The space management statistics give you information on the usage and activity of the SMDS data set owned by this queue manager.

SMDS space management usage:

This section gives information on the number of messages in the the dataset, how many blocks are in use, and how many blocks are available.

```
   SMDS space management usage:
      Messages in data set          27827        highest      27827
```

This is number of messages in the data set, when the SMF record was created, and the highest in the interval.

Field names used: QESDSMMC, QESDSMMM.

```
      Total blocks                  22914
        Space map blocks                1
        Message data blocks         22913
        Data blocks used            11306 ( 49%) highest      11306 ( 49%)
        Data blocks free            11607 ( 51%) lowest       11607 ( 51%)
```

This is information about the number of records in the amount of space used. When 100 1MB messages were put to the queue 402 blocks were used. Each block/buffer size is 256KB (DSBLOCK attribute of CFSTRUCT) see below. 4 Blocks per 1MB message * 100 messages = 400 blocks.

Field names used: QESDSMBT; QESDSMBS; QESDSMBD; QESDSMMC, QESDSMMM; QESDSMBU,

(QESDSMBU, QESDSMBD), QESDSMMU, QESDSMMU, QESDSMBD; QESDSMBF (QESDSMBF, QESDSMBD), QESDSMMF,(QESDSMMF, QESDSMBD).

SMDS space management activity

This section gives information on the activity of the SMDS. The usage section above shows how many blocks have been used. This section tells you how many times the blocks were used.

```
SMDS space management activity:
  Action        Messages     4K pages
  Allocated        27827       723502
```

This is the number of requests to use a message or a page. A buffer and page can be reused may times.

Field names used: QESDSMAR, QESDSMA.

```
 No space             0
```

This number of times the SMDS had no space. This should be 0 in normal usage.

Field name used: QESDSMFL.

```
  Released            0           0
  Reallocated         0           0
  Cleaned up          0           0
```

Field names used: QESDSMFR, QESDSMFP; QESDSMRR, QESDSMRP;QESDSMCR, QESDSMCP.

SMDS buffer pool statistics:

The queue manager has a number of buffers to access the SMDS. This section gives information on the number of buffers and how often they were used.

SMDS buffer pool usage:

This section gives the size of the buffers (the block size) and how many buffers have been used.

```
SMDS buffer pool usage:
  Buffer size (DSBLOCK)          256K
```

This is the size of the buffers, and so the block size of the data.

Field name used: QESDBFSZ/1024.

```
  Total buffers                   1
```

This is the number of buffers allocated. One buffer was specified to produce some of the wait conditions below. You would normally have enough buffers, so that you application did not have to wait for a buffer. Using a large number of buffers can uses a lot of of MVS auxiliary storage and MVS real storage.

Field name used: QESDBFTO.

```
  Buffers in use                  1 (100%)  highest      1 (100%)
    Shared buffers                1
  Buffers free                    0 (  0%)  lowest       0 (  0%)
    Saved buffers                 0
    Empty buffers                 0
  Waiting request queues
    For free buffer               1         highest      1
    For busy buffer               0         highest      0
```

Field name used: QESDBFFS+QESDBFFE, (QESDBFFS+QESDBFFE)/QESDBFTO, QESDBFMF, 100 * QESDBFMF/QESDBFTO; QESDBFUS; QESDBFFS+QESDBFFE, (QESDBFFS+QESDBFFE)/QESDBFTO, QESDBFMF, QESDBFMF/QESDBFTO; QESDBFFS; QESDBFFE; QESDBFPW, QESDBFMP; QESDBFBW, QESDBFMB.

SMDS buffer pool activity:

This section gives information on the activity of the buffers. The buffer pool usage section above shows how many buffers were used. This section displays how many times the buffers were used.

```
SMDS buffer pool activity:
  Acquired buffers              39131
    Got valid buffer                        0 (  0%)
```

'Got valid buffer' is the number of buffers when a get was issued, and the data was in the buffer, and so there was no need to get the data from the SMDS.

Field name used: QESDBFGB; QESDBFGV, QESDBFGV/QESDBFGB.

```
    Got matching, empty buffer            0 (  0%)
    Got free, empty buffer                1 (  0%)
```

This number of times a buffer was obtained which was free, with no data in it. This occurs when the queue manager is started, and so the buffers are empty.

Field name used: QESDBFGM, QESDBFGM/QESDBFGB;QESDBFGF, QESDBFGF/QESDBFGB.

QESDBFGL, QESDBFGL/QESDBFGB.

```
    Stole a saved buffer              39130 ( 99%)
```

This number of times there were no free buffers, and so an existing buffer was used.

Field names used: QESDBFGL, QESDBFGL/QESDBFGB.

```
   No buffer available           11305
```

This number of times there were no buffers available because they were all in use. You should increase the number of buffers.

Field name used: QESDBFGN.

```
    Waited for free buffer         6217 ( 16%) avg time 0.004568s
```

This number of times there were no free buffers, and the average wait time until a buffer was available.

Field names used: QESDBFWP,(QESDBFWP, QESDBFGB), QESDBFPT/QESDBFWP.

```
  Waited for busy buffer            0 (  0%) avg time 0.000000s
```

A buffer can be in use (busy), for example if the buffer was being written to the SMDS. This number of times the processing had to wait for a busy buffer, and how long it had to wait until until this buffer was available for use.

A buffer can be used for multiple messages. If there was I/O active for a buffer, and another message also wants to use the buffer then it has to wait for the busy buffer.

Field names used: QESDBFWB,(QESDBFWB, QESDBFGB), QESDBFBT/QESDBFBW.

```
   Buffer read issued             1406
      Data already valid             0 (  0%)
      Data partly valid              0 (  0%)
      Data read from disk         1406 (100%)
```

This number of times there was a request to get a message from the queue manager that owns the SMDS.

The data shows all records were read from the SMDS, and none were already available in buffers

Field name used: QESDBFRR; QESDBFRS, QESDBFRS/QESDBFRR; QESDBFRP, QESDBFRP/QESDBFRR; QESDBFRR- QESDBFRS- QESDBFRP, (QESDBFRR - QESDBFRS - QESDBFRP)/QESDBFRR;

```
    Freed valid buffer            39130
```

This is the number of times buffer with valid content was reused.

Field names used: QESDBFFB.

```
    Marked buffer deleted             0
    Buffer write issued           39130
```

This is the number of times a buffer was written to the SMDS.

Field names used: QESDBFDB;QESDBFWR.

SMDS I/O statistics:

The section reports on the I/O statistics to the SMDS owner by this queue manager, and the read activity to the SMDS from other queue managers.

SMDS data set usage:

The section report on the highest use record in the dataset, the size of the data pages. The data page (Control Interval), size is always 4096 bytes. The Control area size is the DSBlock size on the Alter or Define CFSTRUCT command.

This is for the R/W SMDS owned by this queue manager.

```
SMDS I/O statistics:
  SMDS data set usage:
    High allocated CI           1466496
    High formatted CI           1466496
    Control interval size          4096
    Control area size            589824
```

Field names used: QESDIOHA, QESDIOHU, QESDIOCI, QESDIOCA.

SMDS I/O activity:

This section reports on the I/O activity when formatting a new extent, for writing records, and for reading records, for the R/W SMDS for this queue manager.

```
    Type         Requests 4K pages pages/req avg I/O time avg wait time MB/Sec %busy
    Format              0        0     0.0       0.000000s     0.000000s      0      0
>W  Write           39130   723467   18.5       0.000676s     0.000656s    106     21%
>L  Read (local)     1406    26000   18.5       0.000438s     0.000336s    164     32%
>O  Read (Other)    39107   723034   18.5       0.000663s     0.000646s    108     22%
```

This is information about the R/W SMDS for the queue manager. The records are

- Format- if the SMDS expanded in size, then it extends the data set, and formats it
- Write- data is written to the data set
- Read (local) - data is read from the data set owned by this queue manager
- Read (other) - data is read from the data set owner by another queue manager.

If there has been I/O then the first character is set to >. To find interesting sections in the report find '>'.

The columns are

- Requests – the number of I/O requests
- 4K pages – the number of 4K pages written
- pages/req – the number of pages per request
- avg I/O time – the average I/O time for the request.
- avg wait time – the average time the request had to wait before the I/O could be started. If this value is larger than 0, this indicates there were not enough buffers.
- MB/Sec – (number of pages * 4096)/1MB / total I/O time.
- %busy – this is the total I/O duration as a percentage of the SMF interval. Because there can be multiple concurrent task accessing SMDS this percentage can be over 100%.

Field names used: QESDIOFR, QESDIOFP, QESDIOFR/QESDIOFP, QESDIOFT/QESDIOFR, QESDIOFW, QESDIOFR; QESDIOWR, QESDIOWP, QESDIOWR/QESDIOWP, QESDIOWT/QESDIOWR, QESDIOWW, QESDIOWR; QESDIORR, QESDIORP, QESDIORR/QESDIORP, QESDIORT/QESDIORR, QESDIORW, QESDIORR;

When reading from an SMDS belonging to another queue mangers the following information is provided. Read (other). The data has the same interpretation as for the local SMDS data set. The data is accumulated for all of the other SMDS for the other systems, it is not available for individual data sets.

```
Type         Requests 4K pages pages/req avg I/O time avg wait time MB/Sec   %busy
Read (Other)   39107   723034   18.5       0.000663s     0.000646s    108      22%
```

Field names used: QESDIOOR, QESDIOOP, QESDIOOR/QESDIOOP, QESDIOOT/QESDIOOR, QESDIOFW, QESDIOOR;

Shared Message Data Set Space CSV

The data in the //SMDSSCSV file gives information about the space used by SMDS. The fields are

z/OS	The name of the z/OS image (SMFRECSID).
QM	The name of the queue manager (SMFRECSSID).
Date	The date from the SMF record (SMFRECDTE)
Time	The time from the SMF record (SMFRECTME)
Duration	The duration in seconds ((qwhsdurn+500000)/1000000).
StrNum	The structure number (QESDSTRN).
StrName	The structure name (QESDSTR).
Msgs	Current messages in data set (QESDSMMC). item[hMsgs] Current messages in data set (QESDSMMM).
DBlocks	Current data blocks used (QESDSMBUP).
DBlocks%	Current data blocks used as a percentage of maximum (QESDSMBU,QESDSMBD).
hDBlocks	Highest data blocks used(QESDSMMU).
hDBlocks%	Highest data blocks used as a percentage of maximum (QESDSMMU,QESDSMBD).

Shared Message Data Set Buffers CSV

The data in the //SMDSBCSV file gives information about the buffers used by SMDS. The fields are

z/OS	The name of the z/OS image (SMFRECSID).
QM	The name of the queue manager (SMFRECSSID).
Date	The date from the SMF record (SMFRECDTE)
Time	The time from the SMF record (SMFRECTME)
Duration	The duration in seconds ((qwhsdurn+500000)/1000000).
StrNum	The structure number (QESDSTRN).
StrName	The structure name (QESDSTR).
DSBlock	Buffer size (QESDBFSZ).

nBuffers	Total number of buffers (QESDBFTO).
inUse	In use buffers QESDBFUS+QESDBFUP.
inuse%	In use buffers as a percentage of the total buffers ((QESDBFUS+QESDBFUP)/QESDBFTO).
hInuse	Highest in-use buffers (QESDBFMU).
hInuse%	Highest in-use buffers as a percentage of the total buffers (QESDBFMU/QESDBFTO).
wFree	Count of current waiting for free buffer (QESDBFPW).
hwFree	Highest count of waiting for free buffer (QESDBFMP).
wBusy	Count of current waiting for busy buffer (QESDBFBW).
hwBusy	Highest count of waiting for bsuy buffer (QESDBFMB).
NoBuffers	The number of times there were no buffers available (QESDBFGN)

Shared Message Data Set Activity CSV

z/OS	The name of the z/OS image (SMFRECSID).
QM	The name of the queue manager (SMFRECSSID).
Date	The date from the SMF record (SMFRECDTE)
Time	The time from the SMF record (SMFRECTME)
Duration	The duration in seconds ((qwhsdurn+500000)/1000000).
StrNum	The structure number (QESDSTRN).
StrName	The structure name (QESDSTR).
What	The request type - one of Format, Write, Read, Other, Total.
Requests	The number of requests (QESDI?R).
Pages	The total number of 4KB pages processed (QESDIO?P).
Pages/Req	The average number of pages per requests (QESDIO?R/QESDIO?P).
IOTime	The total time spend doing I/O (QESDIO?T).
WaitTime	The average time the request had to wait before the I/O could be started. If this value is larger than 0, this indicates there were not enough buffers. (QESDIO?W).
IORate	The rate of processing data (QESDIO?P * 4096)/ QESDIO?T
IOBusy	This is the total I/O duration as a percentage of the SMF interval. Because there can be multiple concurrent task accessing SMDS this percentage can be over 100%. (QESDIO?T/Stats Interval) * 100;

Additional messages produced in MESSAGE output file.

Message: MQSMDS01E . . . SMDS *s* Highest blocks used *n*%
This message reports the highest block used.
When:
If the ratio (QESDSMMU/QESDSMBD) $>$ 0.99 then the message is MQSMDS01E.
If the ratio $>$ 0.90 then the message is MQSMDS01W.

Message: MQSMDS02E . . . SMDS *s* Count of No Space *n*
This message reports the Count of No Space conditions detected.
When:
If the number of *Allocate failed, data set full* (QESDSMFL) $>$ 100 then the message is MQSMDS02E.
If the number of *Allocate failed, data set full* (QESDSMFL) $>$ 1 then the message is MQSMDS02W.

Message: MQSMDS03E . . . SMDS *s* *n*% buffers in use
This message reports the percentage buffers in use.
The calculation 100* QESDBFMU/QESDBFTO is the percentage buffers in use.

When:
If the percentage of buffers in use > 99 then the message is MQSMDS03E.
If the percentage of buffers in use > 90 then the message is MQSMDS03W.

Message: MQSMDS04E ... SMDS *s* Highest queue waiting for free buffers *n*
This message reports the number of applications waiting for a free buffer. An application needed a buffer but none was available. If this number is greater than 0, it shows applications were delayed waiting for a buffer.
When:
If the *Highest queue waiting for free buffers* > 99 the message is MQSMDS04E.
If the *Highest queue waiting for free buffers* > 1 the message is MQSMDS04W.

Action: You should increase the number of buffers by at least the reported value.

Message: MQSMDS05E ... SMDS *s* Highest queue waiting for busy buffers *n*
This message reports the number of applications waiting for a busy buffer. An application had at least one buffer, and needed another one, but this was not available.
When:
If the *Highest queue waiting for busy buffers* (QESDBFMB) > 99 the message is MQSMDS05E.
If the *Highest queue waiting for busy buffers* (QESDBFMB) > 1 the message is MQSMDS05W.
Action: You should increase the number of buffers by at least the reported value.

Message: MQSMDS06E ... SMDS *s* No buffers available *n*
This message reports the number of times an applications needed a buffer but none was available.
When:
If the *Got no buffer (conditional)* (QESDBFGN) > 99 the message is MQSMDS06E.
If the *Got no buffer (conditional)* (QESDBFGN) > 1 the message is MQSMDS06W.

Message: MQSMDS07E ... SMDS *s* Average wait for free buffer *m.n*s
The average time waiting for a free buffer is *Total pool wait time/ Count Waited for pool (free buffer)* (QESDBFPT/QESDBFWP).
When: The average time waiting > 1 millisecond.
Action: Monitor the average time, and increase the number of buffers.

Message: MQSMDS08E ... SMDS *s* Average wait for buffer *m.n*s
The average time waiting for a buffer is *Total buffer wait time/Count Buffer waiters)* (QESDBFBT/QESDBFBW).
When: The average time waiting for buffer > 1 millisecond.
Action: Monitor the average time, and increase the number of buffers.

Message: MQSMDS10E ... SMDS *s* had some format requests.
When: The number of *Format write (extend) requests*(QESDIOFR) > 0.
Action: Investigate why there was SMDS expansion.

Message: MQSMDS11E ... SMDS *s* Average write I/O time > 1ms, *m.n*s

Message: MQSMDS11E ... SMDS *s* Average read local I/O time > 1ms, *m.n*s

Message: MQSMDS11E ... SMDS *s* Average read other I/O time > 1ms, *m.n*s
When: The average time to read or write to the SMDS was > 1 ms. The duration of an I/O depends on your DASD environment, for example having mirrored DASD will increase the response time. The response time also depends on load and number of concurrent requests.
Action: If the I/O response time is higher than normal, you should review the I/O configuration with your z/OS systems programmers.

Message: MQSMDS12E ... SMDS *s* ... write time *(m.n)*s > ... I/Otime*(m.n)*s. Buffer shortage.
When: The total time waiting for a buffer > I/O time. This usually indicates a buffer shortage.
Action: Increase the number of buffers

QSUML – Queue summary information for local queues

The QSUML data is a summary of the queue usage over time, for local queues

```
Date      ,Time    ,Qmgr,Queue,Count,PS,BP,Put MB ,Get MB ,!,ValidPut,ValidGet,getpsn,MaxQDepth,
2013/01/14,15:00:00,MQPA,CP0000,2134,2,  2,1.8e+06,1.8e+06,!,1.8e+03, 1.8e+03,     0,       50
2013/01/14,15:00:00,MQPA,CP0001,  17,1,  1,1.4e+04,1.4e+04,!,     14,      14,     0,        2
2013/01/14,15:00:00,MQPA,CP0002,  17,1,  1,1.4e+04,1.4e+04,!,     14,      14,     0,        2
2013/01/14,16:00:00,MQPA,CP0000,2134,2,  2,3.3e+05,3.3e+05,!,3.3e+02, 3.3e+02,     0,       50
2013/01/14,16:00:00,MQPA,CP0001,  17,1,  1,2.7e+03,2.7e+03,!,      3,       3,     0,        2
```

Where the fields are

Date	in format YYYY/MM/DD
Time	is either the queue open time (if the open was in the SMF interval) or the interval start time from the WTAS task record. The time is then rounded down to the nearest hour.
QMGR	is the queue manager name
Queue	is the queue name
Count	is the number of queue records processed.
PS	is the page set number
BP	is the buffer pool
Put MB	is the amount of data put to the queue in MB
Get MB	is the amount of data got from the queue inMB
ValidPut	is the number of valid put requests
ValidGet	is the number of valid get requests
Getpsn	is the number of gets from the page set
MaxQDepth	is the maximum depth of the queue.

QSUMS – Queue summary information for Shared queues

The QSUMS data is a summary of the queue usage over time, for shared queues.

```
Date      ,Time    ,Qmgr,Queue,Count,Structure   ,Put MB,Get MB,!,ValidPut,ValidGet,MaxQDepth,
2013/01/15,13:00:00,MQPA,SQ1  ,    3,APP1        , 3e+03, 2e+03,!         3,     2,          1,
```

Where the fields are

Date	in format YYYY/MM/DD
Time	is either the queue open time (if the open was in the SMF interval) or the interval start time from the WTAS task record. The time is then rounded down to the nearest hour.
QMGR	is the queue manager name
Queue	is the queue name
Count	is the number of queue records processed.
Structure	is the CF Structure name
BP	is the buffer pool
Put MB	is the amount of data put to the queue in MB
Get MB	is the amount of data got from the queue inMB
ValidPut	is the number of valid put requests
ValidGet	is the number of valid get requests
Getpsn	is the number of gets from the page set
MaxQDepth	is the maximum depth of the queue.

Log statistics

The log statistics give information about the amount of data written to the log datasets.

You can use this section

1. To determine if you are reaching the limit at which the queue manager can write data to the log data sets.
2. For regular checks of the logging characteristics and workload
3. You think you have a problem with logging

The detailed information is written to ddname //LOG.

Data in CSV format is written to //LOGCSV, see Log statistics in CSV on page 51

The key throughput indicators are displayed in CSV format in Logbusy on page 51

See SCSQMACS(CSQDQJST) for the layout of the SMF record.

Has my queue manager reached the limit at which it can log data?

There are two factors which limit the rate at which MQ can write to the log data sets.

1. Pages per I/O
2. How busy the logging task is.

Pages per I/O

It is more efficient to process many pages per I/O. You can get many pages per I/O when processing large messages, large units of work, or channels with a large batch size. If you have only small messages (under 4KB) you may only get one or two pages per I/O. Increasing the number of concurrent tasks processing messages may increase this, but the number of pages per I/O may still be low.

On an IBM performance machine we have seen the upper limit of pages per I/O of about 90. This is with 1MB messages and many gets per commit.

How busy the logging task is

Within the queue manager there is a task which does the logging.

This task

1. Issues a write request to each log, this write one or more 4KB pages.
2. Waits for the write request(s) to end
3. Resumes any tasks waiting for I/O in the data just written
4. Waits for the next request to write data. In a busy system there is usually no wait time.
5. Loops to the top.

In a busy system, most of the time spent by this task is waiting for disk I/O.

The time for an I/O request breaks down into two parts.

1. Time to set up the request
2. Time to transfer the data this varies with the amount of data writen.

As more data is logged, the I/O rate can decrease. More data is written per I/O request, to the duration of the I/O request increases, so there are fewer requests per second. The I/O rate is not so important as how busy the log task is.

At a simple level the task busy is *rate of I/O requests * average time for I/O request.*

This "task busy" is calculated as part of the log statistics displayed below.

To improve throughput you need to reduce the I/O time. The time for an I/O request is very dependant on the workload and the I/O configuration, including the path to the DASD (channels) as well as the DASD subsystem itself. For example on a performance machine in IBM

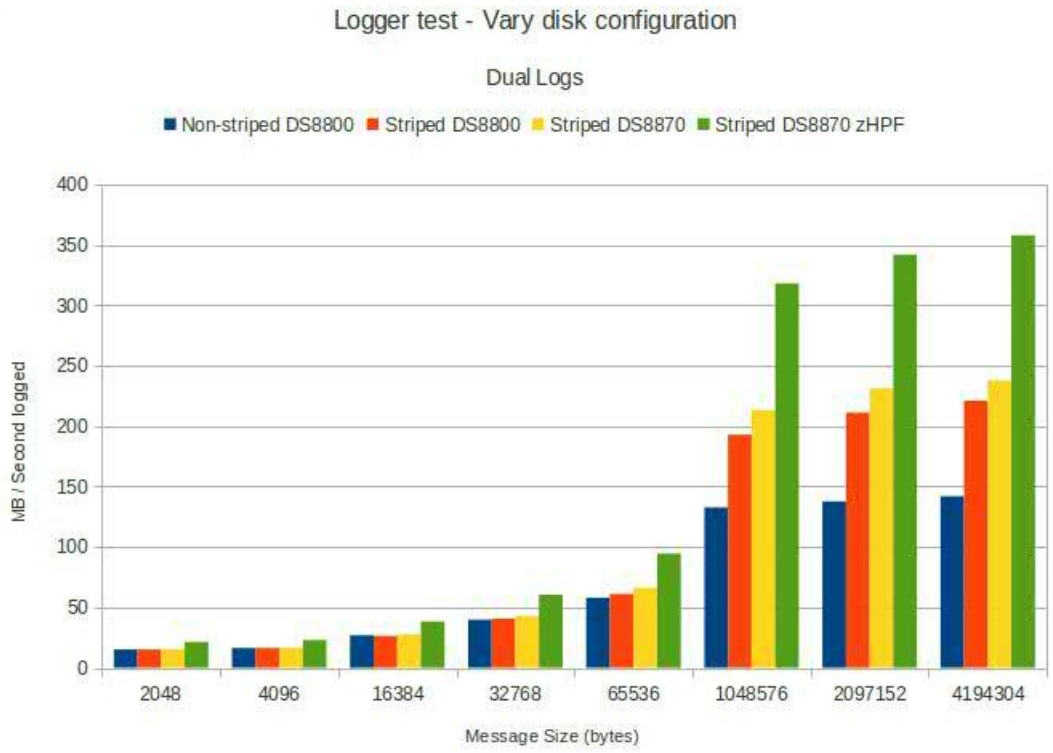

The chart shows the effect of different hardware configurations using MQ V8. This is the logging rate for single log in a dual log environment.

1. Message sizes from 2KB to 4MB were used.
2. Striping only makes a difference with larger messages
3. Faster disks (DS8870) gave a bit more uplift (at the top end) but little difference with smaller messages (because the internal logger task was 99% busy for the 2KB messages)
4. zHPF gave a big increase especially with large messages.

Using mirrored DASD will decrease throughput depending on the distance.

To see if your queue manager is approaching its logging limit see Logbusy on page 51.

Regular checks

You should review the log statistics regularly to get a profile of the logging statistics by hour and by day. This will allow you to compare a bad day with a good day to see if there is a significant change in the logging rate.

This data is reported in Log statistics in CSV on page 51.

You need to review

1. Wait for buffer. Ideally this should be zero or a very small percentage of the number of CIs offloaded. A value greater than zero means you have run out of buffers. Monitor this to see if the value increases over time.
2. Log task busy.
3. Pages per I/O

Depending on your workload you may have log busy > 90% and be limited when

1. pages per I/O is small. Typically lots of short messages are being processed
2. pages per I/O is large. Typically large messages are being processed, or many concurrent tasks processing short messages

You need to know the profile of pages per I/O for your environment.

You think you may have a problem with logging.

If you think you have a problem with the logging, you may need to look at the detailed records.

Key fields to check. See the text following to identify the specific fields you need to check.

Wait for buffers should always be zero. If this number is greater than zero, the internal buffer filled up and there was no space to store any more data. Applications processing persistent messages will be delayed until buffers are available.

Space in this buffer is freed when the I/O using the space has completed. This problem can be caused by

1. Active logs filling up – perhaps due to a problem with archiving.
2. More data is being logged than the I/O system can handle. Improve the I/O rate, perhaps by striping the logs, or moving the log data sets to low use volumes, or reduce the work on the queue manager.
3. It can be caused by putting large messages, or many smaller messages, and then doing a commit. If the application has put more data than available buffers, you will get a *wait for buffers* condition.

The number of pages per I/O. If this is small you may have many small transactions. Having more transactions running concurrently should increase the pages per I/O and so the throughput

There have been some unusual disk response times. With MQ V7.1 APAR PM61284 and MQ V800 the log statistics report the average I/O response time, and the maximum response time seen in the interval. This may indicate a problem with the DASD subsystem.

If the I/O response time for writing one page is significantly different from what you expect this may indicate a problem with the I/O subsystem.

Interpreting the output

Records always produced

```
MVCA MQ7A 2013/01/04 14:29:37 VRM:710
   Wait for buffers(should be 0):             0
   Total Number of pages written:          5078
   Total Number of write requests:          172
   Pages written per I/O:                    29
```

Records produced if detail >= 10.

```
MVCA MQ7A 2013/01/04 14:29:37 VRM:710

   Wait for buffers(should be 0):             0   out of          2526
```

```
   Total Number of pages written:          5078
   Total Number of write requests:          172
   Pages written per I/O:                    29
   Write_Wait         0, Write_Nowait       391, Write_Force          1, WTB         0
   Read_Stor          0, Read_Active          0, Read_Archive         0, TVC         0
   BSDS_Reqs          3, CIs_Created       2526, BFWR               113, ALR         0

   ALW                0, CIs_Offload          0, LLCheckpoints        0

   Read delayed       0, Tape Lookahead       0, Lookahead Mount      0
   Write_Susp        13, Write_Reqs         172, CI_Writes         5078
   Write_Serl         0, Write_Thrsh        100, Buff_Pagein          0
```

Where LLCheckpoints shows the number of checkpoints due to the LogLoad value, and does not include checkpoints due to the logs filling up.

Records produced at V8 or with APAR PM61284

With MQ V7.1 APAR PM61284 and V8 the format is

```
From 2014/06/06,19:01:00.017286 to 2014/06/06,19:02:00.016951, duration    60 seconds.

 Wait for buffers(should be 0):             0   out of          2526
 Total Number of pages written:       2783396
 Number of pages written/sec:           46389
 Amount of data written/sec:              181 MB/Sec
 Total Number of write requests:       161826
 Number of write requests/sec:           2697
 Pages written per I/O:                    17
 Total number of read requests:             0
 _____,__ write requests,        CIs, Average I/O , After I/O   , pages/IO
                                      time in uSec, time in uSec,

 Log 1, 1 page      45373,      45373,         275,             5,        1
 Log 1,>1 page      35540,    1346300,        1192,             4,       38

 Log 2, 1 page      45373,      45373,         274,            18,        1
 Log 2,>1 page      35540,    1346300,        1174,            50,       38
 Standard deviation of first log, 1 page per I/O, response time +-        3
 Log 1, 1 page Longest I/O      149742 at 2014/06/06,18:01:37.439014 UTC
 Log 1, 1 page Longest Request  149746 at 2014/06/06,18:01:37.439014 UTC
 Log 1,>1 page Longest I/O      331246 at 2014/06/06,18:01:37.095495 UTC
 Log 1,>1 page Longest Request  331249 at 2014/06/06,18:01:37.095495 UTC
 Log 2, 1 page Longest I/O       46341 at 2014/06/06,18:01:15.381098 UTC
 Log 2, 1 page Longest Request  149742 at 2014/06/06,18:01:37.439018 UTC
 Log 2,>1 page Longest I/O       52811 at 2014/06/06,18:01:39.381000 UTC
 Log write rate      92MB/s per copy
 Logger I/O busy : 95.11%
 Logger task busy: 97.99%           (V800)
```

This information shows log 1 had 45373 write requests where it wrote one page per I/O. The average I/O time of these requests was 275 microsecond.

Log 1 also had 35530 requests to write more than one page – the average was 38 pages per I/O. The average elapsed time for the I/O was 1192 microseconds. Similarly for log 2.

Over time you should expect the average time to write just one page to vary by a small amount.

When more than one page is written per I/O, the duration of the I/O will depend on the amount of data written.

When an I/O operation has finished the logging task has to be resumed. The duration between the I/O complete and the time the logging task processed the next instruction is recorded as the '*After I/O time in uSec'* This value should be small. If there is a shortage of CPU, this task may be delayed, so a significant increase in this value may indicate a shortage of CPU.

Typically the time for log 2 is greater than log 1, because it includes time for processing log 1.

In V8 the line

```
Standard deviation of first log, 1 page per I/O, response time +-        3
```

reports how consistent the response time was for writing one page. This should be close to 0. If this is larger than your average response time this may indicate your I/O subsystem is not giving you consistent response time.

The time of data when the longest I/O occurred, and the longest time the logging task was suspended is also recorded.

So in the report where it has

```
Standard deviation of first log, 1 page per I/O, response time +-        3
Log 1, 1 page Longest I/O      149742 at 2014/06/06,18:01:37.439014 UTC
```

The longest I/O time was 149742 microseconds and this occurred at 18:01:37. You can use this time to correlate with any out of line events, such as transactions taking longer than usual.

The log write rate is the number of pages written per log per second * 4096 bytes per page.

The logger I/O busy is the sum of

- Maximum of (time for log 1 to write 1 page, time for log 2 to write 1 page)
- Maximum of (time for log 1 to write >1 page, time for log 2 to write >1 page)

So the value of 95.11% shows that there was I/O active for most of the time. The queue manager is close the the limit of what it can log. To improve throughput you will need to improve the I/O response time to the DASD.

The response time for writing one page should be consistent over time. The response time when multiple pages are written will depend on the number of pages written and whether the log datasets are striped. So monitoring the response time of writing one page gives you a good indication of the behaviour of your DASD over time.

The *logger task busy* field is new in V8 and shows how busy the log task was. The value in the SMF record is the time the log task was idle. So the time the log task was active is (duration – qjstslptu).

The example data shows that the queue manager is close to the limit at what it can log.

This can be seen from

- The majority of the requests processed multiple pages 1346300 (> 1 page per I/O) compared to 45373(1 page per I/O)
- High number of pages per I/O (38)
- Logger I/O % busy (95%.11)
- **Logger task busy: 97.99**

With 10 batch jobs putting and getting 1KB persistent messages the key statistics were

```
_____,__ write requests,        CIs, Average I/O , After I/O   , pages/IO
                                      time in uSec, time in uSec,
 Log 1, 1 page     209425,     209425,         218,             3,        1
 Log 1,>1 page          1,          2,        2078,             2,        2
 Log 2, 1 page     209425,     209425,         221,             7,        1
 Log 2,>1 page          1,          2,        2066,             3,        2
 Log write rate      13MB/s per copy
```

```
Logger task busy: 82.53%
Logger I/O busy : 79.86%
```

This shows the logging task was busy, but the majority of the I/O had just one page per I/O.

The queue manager can log more data, but more concurrent jobs will be needed (or more data written before the commit request).

Log statistics in CSV

The log statistics give information about the amount of data written to the log datasets.

The data reported is for one log, so if you have dual logs you will have similar data for the second logs.

The records are written to ddname //LOGCSV.

```
'z/OS','QM  ','Date    ','Time','MB Written','MB/SEC','MB Used','Pages per
I/O','LLCheckpoints',read_buf,read_act,read_arc,r_delay
MVCA,MQ7A,2013/02/09,16:50:00,51749, 172,25675,  31,1
MVCA,MQ7A,2013/02/09,16:55:00,48389, 161,24009,  32,  1
```

Where

z/OS	Is the name of the z/OS image
QM	Is the queue manager name
Date	Date in YYYY/MM/DD formatting
Time	The time in HH:MM:DD format
MB Written	This is field (qjstioci for 1 page and > 1 page for the first log)converted to MB. This field is the number of pages written to DASD. A page may be written multiple times if it was not full, and a commit was issued
MB/SEC	This is field *MB Written/ SMF duration*
MB Used	This is the amount of log data used qjstbfflin MB. This can be smaller that the MB written, as a page may get written multiple times
Pages per I/O	As rate of data written to the increases, then this value increases. The maximum value depends on your environment, especially the rate at which your can write data to your DASD.
LLCheckpoints	The number of checkpoints that occurred in the interval due to logload.
read_buf	The number of read requests satisfied from buffers
read_act	The number of read requests satisfied from active logs
read_arc	The number of read requests satisfied from archive logs.
r_delay	

Logbusy

With WMQ V7.1 APAR PM61284 and WMQ V8 information is provided on log data set I/O response times as seen by the logger.

The duration of the interval is in qwhsdurn.

Information is reported when one page is written, and when multiple pages are written, and for each log.

Qjstiototsus is the total time spent suspended for I/O.

The busy time is calculated as

max1page time = max(qjstiototsus for 1 page log1,qjstiototsus for 1 page log 2)

maxNpagetime = max(qjstiototsus for >1 pages log1,qjstiototsus for > 1 pages, log2)

busy time = max1page time + maxNpagetime

100 * Busy time/ duration is the logBusy% time.

The logbusy data set has the following data

```
z/OS,QM,Date,Time,logBusy%,MB/Sec, pages/IO
```

Where

z/OS	Is the name of the z/OS image
QM	Is the queue manager name
Date	Date in YYYY/MM/DD formatting
Time	The time in HH:MM:DD format
logBusy%	See logBusy% above
MB/Sec	This is the same as in Log statistics in CSV above
Pages/IO	This is the overall average pages/IO.

Additional messages produced in MESSAGE output file.

Message: MQQJST00I ... Log read log buffers from storage n > 0

Detail: 15

When: If qjstrbuf > ReadLogBuffers

Reason: The field Read_stor (QJSTRBUF) is the number of requests which were satisfied from the log buffers, and did not require the log datasets to be read. The value of this field is greater than the MQSMF parameter ReadLogBuffers.

Reading data from the log can often occur, for example when an application processing persistent messages rolls back the work. If the amount of roll back activity is large, then this may indicate a problem in the application or the environment.

Action:

Monitor the number of log buffers read from storage and specify a suitable value in the ReadLogBuffers parameter.

Message:MQQJST01W ... QJST read log buffers from active logs n > 0

Detail: 10

When: qjstract > 0

Reason: Message MQQJST00I is for the number of log buffers read from memory. If the data is not in memory, then the data is read from a log data set.

This message is produced when data was read from an active log dataset.

In normal operation you do not expect messages to be read from the active log datasets.

Action:

Investigate why data was being read from the active log datasets. For example, this could be due to transactions processing lots of data in an unit of work, or taking a long elapsed time before work is committed; and then rolling back.

Message: MQQJST02S ... Log read log buffers from archive logs n > 0

Detail:4

When: qjstrarh > 0

Reason: This message is produced if log data is read from an archive log dataset, which should not happen.

Action:

See MQQJST01W on page 52.

You may need to increase the number of active log datasets. Check the log datasets are large enough.

Message: MQQJST03E ... Log Number of checkpoints n > m

Detail:10

When: qjstllcp > 10

Reason: This message is produced when there are more checkpoints in the SMF interval than expected.

Action:

This may indicate a peak in activity.

The active logs may be too small.

Monitor the qjstllcp and specify a suitable value in the parameters

Message: MQQJST04E ... Log Number of buffer paged in n > 0

Detail:10

When: qjstbpag > 0

Reason: Some log requests were delayed because the log buffers had to be paged in. This value should be 0

Action: This indicates that there may be a real storage shortage. Investigate and resolve this.

Message: MQQJST05E ... Log Number of read accesses delayed n

Detail:10

When: qjstwur > 0

Reason: Some log read requests were delayed due to an archive log not being available. This field should be 0

Action: Investigate why the archive log was being used.

Message: MQQJST06E ... QJST Number of look ahead tape mounts attempted n

Detail:10

When qjstlama > 0

Reason: An archive log on tape was needed.

Action: Investigate why the archive log was being used.

Message: MQQJST07E ... Log Number of look ahead tape mounts performed n > 0

Detail: 10

When: qjstlams > 0

Reason: An archive log on tape was needed.

Action: Investigate why the archive log was being used.

Message: MQQJST07E ... QJST Number of reads delayed for tape contention n

Detail: 10

Reason: There was contention for an archive log on tape

Action: Investigate why the archive log was being used.

Message: MQJST09W ... Log % requests waiting for buffer n > m

Detail:15

When: pcwtb >= 1

Reason: There were no log buffers available to be used, they were either being written to the active log datasets, or waiting to be written to the active log datasets.

This can be caused by

- increased application traffic,
- writes to the active log datasets are slower usual (this may be a hardware related problem)
- or all of your active logs are full.

Action: Investigate your active log DASD and check this is not being constrained.

In the long term you may need to move work to a different queue manager, as you are reaching a limit to the rate at which you can log data.

Monitor how frequently this message occurs. If it only occurs infrequently, this may be acceptable. If it starts to occur more frequently then you need to take action to improve the rate at which you can log data (stripe the datasets, or move the logs to faster DASD) or try to reduce the amount of data logged by this queue manager.

Message: MQQJST10E ... Log High logging rate n > m MB/Sec

Message: MQQJST11W ... OK logging rate n > m MB/Sec

Message: MQQJST11I ... logging rate is low n < m MB/Sec

Detail:15

When: the number of CIs written (qjstciwr) in the specified time exceeds the parameter value

Reason: The number of 4K pages written to the log is qjstciwr, This value in MB is qjstciwr * (4*1024) / (1024*1024=1MB) = qjstciwr * 4/1024 MB

The SMFTime parameter is the time you have specified between SMF records, so the rate at which data is logged is (qjstciwr * 4/1024 MB)/SMFTime (in seconds).

If this data rate exceeds the HighLogRateMB parameter, then message

MQQJST10E ... High logging rate n > m MB/Sec is produced

else if this data rate exceeds the OKLogRateMB parameter, then message

MQQJST11W ...OK logging rate n > m MB/Sec is produced

else message

MQQJST11I ... logging rate is low n < m MB/Sec is produced.

You can use this to monitor your logging rate.

Action: Monitor your logging rate and select a high and OK values.

If you find the logging rate is often larger than your HighLogRateMB value, then this may indicate an increasing amount of data being logged. You should investigate moving your logs to faster DASD, striping the logs, or moving work from the queue manager.

Message: MQQJST14S ... Log max I/O time n:=m t uSec at tt UTC

Message: MQQJST15S ... Log max Suspend time n:=m t uSec at tt UTC

For example

```
MQQJST14S ..., Log max I/O time     1:=1 0234839 uSec at 2014/09/01,09:01:26....
MQQJST15S ..., Log max Suspend time 1:=1 0234860 uSec at 2014/09/01,09:01:26....
```

When: Detail >= 10 and the duration of the request > MQSMF parameter LongLogIO

- n can be 1 or 2 for the log copy 1 or log copy 2
- m can be 1 or 1> for one page or multiple pages per I/O
- t is the duration in microseconds
- tt is the time when it occurred.

The logging task has the following logic (for a single log)

- get time of day t1
- start I/O request to log
- wait
- when the I/O completes, get time of day t2, and resumes the logging task
- logging task is resumed
- get time of day -time t3
- ...

The I/O takes t2-t1 to complete

The logger task was suspended for t3 - t1

The interval between the I/O completing and the logger task being woken up is t3 - t2.

Although the average Log I/O response time is often under 1000 uSeconds, we have seen the maximum I/O time being 10ms to 100 ms, so a high value may not necessarily be a problem.

Action: Monitor this value and see how it varies over a period of a week.

If the time of the longest I/O corresponds to a slow down in your applications, then this indicates a problem and you should investigate why the DASD is slow.

Usually the time the between the I/O completing and the logger task being woken up is about 20 microseconds. A large (eg > 1000) difference can indicate a CPU problem. Check the WLM classification of the Queue manager, and use tools like RMF Work Delay Monitor (RMFWDM) to see if the queue manager is delayed for CPU.

Message MQQJST16W ... Log Log write task nn% busy

Reason: The internal logger task was more than 50% busy

Action: Investigate to see if the I/O response time can be improved.

Action: This gives you early warning that the logger is busy. If the busy time continues to increase you should consider moving the log data sets to faster DASD, or reducing the persistent workload on the queue manager.

Message MQQJST17S ... Log Log write task very busy at nn % busy

Reason: The internal logger task was more than 50% busy

Action: This internal logger task is very busy and may be close to the throughput limit.

Message MQQJST18S ... Log % of interval spent doing disk I/O nn %

Reason: The percentage of the duration there was I/O active was more than 90% busy

Action: Investigate to see if the I/O response time can be improved.

Message MQQJST19W ... Log % of interval spent doing disk I/O nn %

Reason: The percentage of the duration there was I/O active was more than 50% busy

Action: This gives you early warning that the queue manager is spending a significant amount of time doing log I/O. If the busy time continues to increase you should consider moving the log data sets to faster DASD, or reducing the persistent workload on the queue manager.

Buffer pool Statistics

The buffer pool statistics give information about the buffer pools.

In version 8 buffer pools can be in 64 bit storage and can be page fixed.

The records are written to ddname //BUFF.

See SCSQMACS(CSQDQPST) for the layout of the SMF record.

Records always produced.

```
MV4C,MQ03,2014/10/27,11:26:00,VRM:900,
= BPool   0, Size  2000,%full now  0, Highest %full  0, Disk reads     0
= BPool   1, Size  2000,%full now 31, Highest %full 31, Disk reads     0
```

Where the fields are

BPool	The buffer pool number
Size	Number of buffers in this buffer pool when the SMF record was created
%full now	How full the buffer pool was when the SMF record was created. This value is the number of buffers in the buffpool – number of free buffers.
Highest %full	This is the maximum usage of the buffer pool during the period, as a percentage of the total size of the buffer pool.
Disk reads	The number of pages read from disk. You get best performance if all of your data is in the buffer pool. Reading from the pageset usually indicates the buffer pool had filled, and so buffers had to be written out to the page set

Record produced if detail >= 5

```
MV4C,MQ03,2014/11/20,13:19:55,VRM:900,

  From 2014/11/20,13:19:51.363707 to 2014/11/20,13:19:55.276668, duration     4 seconds.

= BPool   0, Size  2000,%full now  0, Highest %full  6, Disk reads   261

< BPool   0, Pages written/sec    33, Pages read/sec                  65

 > 00 Buffs     2000  Low     1866  Now     1996  Getp      537  Getn        0
   00 Rio        261  STW      373  TPW      135  WIO        12  IMW         2
   00 DWT          0  DMC        0  STL      261  STLA        0  SOS         0
   00 Below the bar   PAGECLAS 4KB
```

The line

```
  From 2014/11/20,13:19:51.363707 to 2014/11/20,13:19:55.276668, duration     4 seconds.
```

is available with MQ V7.1 APAR PM61284 and MQ V800, and provides the interval start and end times, and the interval duration.

With MQ V7.1 APAR PM61284 and MQ V800 the line

```
< BPool   0, Pages written/sec    33, Pages read/sec                  65
```

uses the interval duration to covert the number of pages read and written to a rate of pages processed per second. If the both the pages read and written per second are zero this line is ommitted.

Where `00 Above the bar PAGECLAS 4KB` is new in V800.

Buffer pool Statistics CSV

```
MVS,QM,Date,Time,BP,size,lowest_free,highest_used,used_now,% hfull,SOS,...
```

```
...
#_sync_write,DWT,#_get_new_pg,#_get_old_pg,#_read_I/Os, ...
... #_pg_writes,#_write_I/Os,Location,PageClas

MV41,MQ13,2015/08/10,11:45:20,  3,50000,49994,  6,    6,    0,    0,    0,    0,...
    0,    0,    0,    0,BELOW,4KB
MV41,MQ21,2015/08/10,11:45:52,  0, 5000, 4997,  3,    3,    0,    0,    0,    0,...
    0,    1,    0,    0,ABOVE,FIXED4KB
```

Additional messages produced in MESSAGE output file.

Message: MQQPST00W . . . BP n Many(m) buffers written immediately. Buffer pool may be too small.

When: If qpstimw $>$ 0

Reason: The field Number of synchronous write operations (qpstimw) is the number of requests where a page had to be written directly to the page set. The value of this field is greater than the MQSMF parameter BPIMW.

During a checkpoint, pages which were in use during the checkpoint, may get flagged to be written out to the page set when the page is freed. This is normal. Typically this value is under 100 pages.

If the buffer pool is more than 95% full of pages then when a page is released it is immediately written to disk, and so the page is made available to other requests.

A high value of qpstimw indicates the buffer pool was very full.

Action: Review why the buffer pool is filling up. It may be because getting applications are not running, and so there is increasing queue depths.

Message: MMQQPST01S . . . BP n Short on storage occurred m times. Buffer pool may be too small.

When: If qpstsos $>$ 0

Reason: Buffer pool n had no available buffers. This occurred m times.

The buffer pool was critically short of buffers

Action: Review why the buffer pool is filling up. It may be because getting applications are not running, and so there is increasing queue depths.

Add more buffers to the buffer pool.

Message: MQQPST02S . . . BP nu Filled many(m) times. This is typical of long lived messages. Buffer p

When: If qpstdwt $>$ DeferredWriteTaskS

Reason: When the number of free buffers reached 15% this triggers the Deferred Write Processor to start up to move pages from the buffer pool to the page set, and so make more space available. The number of times this happened (qpstdwt) was greater than the MQSMF parameter DeferredWriteTaskS.

Action: Consider making the buffer pool bigger.

Message: MQQPST03W . . . BP n Filled (DWT) m times. Buffer pool may be too small or you had a pea

When: If qpstdwt > DeferredWriteTaskS

Reason: When the number of free buffers reached 15% this triggers the Deferred Write Processor to start up to move pages from the buffer pool to the page set, and so make more space available. The number of times this happened (qpstdwt) was greater than the MQSMF parameter DeferredWriteTaskW.

Action: If the number of times the DWP started is low, this could be a temporary increase in workload (or depth of queues). As the number increases this is an indication that the buffer pool is too small. With increasing workload or depth of queues, this may lead to the buffer pool getting over 95% busy (less than 5% free pages) when there will be a significant performance drop when pages are written synchronously to the page set.

Message: MQQPST04E . . . BP n Many (m) pages read from disk. This is typical of long lived messages.

When: If qpstrio > BPReadIOS

Reason: An application needed a page but it was not in the buffer pool, so the page was synchronously read in from the page set. The number of times this happened (qpstrio) was greater than the MQSMF parameter BPReadIOS.

Action: This can occur if the buffer pool has filled up, and so buffer pool pages were written to the page set, and the buffer page reused. It can also occur if this is the first time the queue has been used since the queue manager was started, and messages had to be read from disk.

If the queue has long lived messages (for example written to an audit queue for processing overnight) then having the messages read from the page set is normal behaviour.

if the messages are short lived messages (less than a couple of minutes) this indicates the buffer pool is too small. Try making the buffer pool bigger, or move the queues with deep, or long lived messages to a different buffer pool.

Message: "MQQPST05I . . . BP n Some (m) pages read from disk. Buffer pool may be too small.

See Message: MQQPST04E above.

The number of times this happened (qpstrio) was greater than the MQSMF parameter BPReadIOW.

Message: MQQPST06I . . . BP n get old to get new page ratio > m Queues not indexed ? Could be a lot o

When: If qpstgetp > BPGetPGetN * qpstgetn

Reason: qpstgetn is the number of requests where a new, empty page was requested. qpstgeto is the number of requests where an existing page was requested. Where BPGetPGetN is MQSMF parameter BPReadIOS

Action: If qpstgeto >> qpstgetn then this can indicate unusual behavior, such as a queue being scanned sequentially and the queue is not indexed.

Review the queues in the buffer pool with deep queues, and check the queues are indexed appropriately.

Message: "MQQPST07I . . . BP n Write rate m pages per second.

When: If qpsttpw / statsDuration > 0

Reason: This is the highest average rate of pages written to the page set. This is written to the TASKSUM output. Use this to locate the entry in buffer pool statistics report.

Message: "MQQPST08I . . . BP n Read rate m pages per second.

When: If qpstrio / statsDuration > 0

Reason: This is the highest average rate of pages read from the page set. This is written to the TASKSUM output. Use this to locate the entry in buffer pool statistics report.

Subsystem statistics

The subsystem statistics give information on how many jobs ended normally, and how many jobs ended abnormally.

The records are written to ddname //EOJ.

Records always produced.

```
MVS  QMGR Date       Time     VRM          Jobs EOT     Jobs EOM
MVCA MQPA 2013/01/06 11:36:46 VRM:701            10            0
MVCA MQPA 2013/01/06 11:37:31 VRM:701             2            0
```

Jobs EOT is the count of TCBs that ended normally. Jobs EOM is the number of jobs that ended abnormally.

Message manager

The message manager statistics gives information on how many API requests there were in the interval.

The records are written to ddname //MSGM.

See SCSQMACS(CSQDQMST) for the layout of the SMF record.

```
MVCA MQ7A 2013/02/09 01:00:00 VRM:710 : no data found
MVCA MQP0 2013/02/09 01:00:00 VRM:710
   MQOPENs         0,  MQCLOSEs         0,  MQGETs         6,  MQPUTs         6
   MQPUT1s         0,  MQINQs           0,  MQSETs         0,  C ALL H        0
   MQSUBs          0,  MQSUBRQs         0,  MQCBs          0
   MQCTLs          0,  MQSTATs          0,  Publish        0
```

If all the value are zero, *no data found* is displayed.

The field names are

MQOPENs	MQOPEN
MQCLOSEs	MQCLOSE
MQGETs	MQGET
MQPUTs	MQPUT
MQPUT1s	MQPUT1
MQINQs	MQINQ
MQSETs	MQSET
C ALL H	Close all handles – issued at the end of the task
MQSUBs	MQSUB
MQSUBRQs	MQ Subscription request
MQCBs	MQ call Back
MQCTLs	MQ control call back
MQSTATs	MQ retrieve status information
Publish	The total number of messages published.

Message manager CSV

The message manager statistics give information on how many API requests there were in the period.

The records are written to ddname //MSGMCSV with one line per record, suitable for importing into a spread sheet.

```
'MVS','QM','   Date','  Time','    Puts','  Put1s','      Gets
MVCA,MQP0,2013/02/09,01:00:00,         0,         0,         6
MVCA,MQPG,2013/02/09,01:00:00,         0,         0,       326
```

Topics

```
MVCA MQ7A 2013/02/12 18:02:35 VRM:710
 Total Subs            2 Durable Subs          0 Expired Subs        0
 Total messages        6 Single publish        0
  API     sub HW        0 Sub LW               0 Tot Pub             6
  ADMIN  :Sub HW        0 SUB LW               0 Tot Pub             0
  PROXY  :Sub HW        0 SUB LW               0 Tot Pub             0
  Single PUB HW:        2 Pub LW               2 Pub Nosub           0
  Max Pub time        259 Avg pub time        76
```

The Topic section gives information on Publish Subscribe

The records are written to ddname //TOPIC.

See QTST in SCSQC370(CSQDSMFC)) for the layout of the SMF record.

Data manager

The data manager manages the links between messages and queues. It calls the buffer manager to process pages with messages on them.

The records are written to ddname //DATA.

See SCSQMACS(CSQDIST) for the layout of the SMF record.

```
MVCA MQ7A 2013/02/09 10:20:00 VRM:710
   Obj Cre         1, Obj Puts        0,  Obj Dels        0,  Obj Gets         3
   Locates         4, Stgclass        0,  Enum           25
   Msg Gets     1000, Msg Puts     1000
   Lock MM         0, Rel MM         32,  Delete MM       0
   Read Ahead:IO  29,:Buffer         22,  Gets disk      17,  Gets BP          3
```

The fields names are

Obj Cre	The number of object creates – such as DEFINE QLOCAL
Obj Puts	The number of times an object was changed, such as ALTER QL(Z)DESCR(Comment)
Obj Dels	The number of times an object was deleted
Obj Gets	The number of times an object definition was got, for example for a display, or for an alter
Locates	The number of times a locate object was issued. This could be for displaying or altering an object, or application activity opening a queue.
Stgclass	The number of requests to alter a STGCLASS
Enum	The number of requests to find an object
Msg Gets	The number of MQGET requests
Msg Puts	The number of MQPUT or MQPUT1 requests
Lock MM	The number of Lock Marked Message requests
Rel MM	The number of Release Marked Message requests
Delete MM	The number of Delete Marked Message requests

If messages are being got in order, and the messages are being read from the page set, then an internal task may be started to perform read ahead to read messages from the disk, so that they are in the buffer pool when the application next does an MQGET.

The read ahead data for messages in MQ V6, and in all releases for getting objects from page set 0.

```
Read Ahead:IO   0,:Buffer          3,  Gets disk       1,  Gets BP       2660
```

Read Ahead:IO	The read ahead task got pages which required I/O to a page set
Buffer	The read ahead got a page – but it was already in the buffer pool and did not require an I/O to the page set
Gets:disk	An MQGet was suitable for a read ahead. This is the count of pages the application got which required I/O to a page set
Gets BP	An MQGet was suitable for a read ahead. This is the count of pages the application got which were already in the buffer pool.

MQQIST01W ... QIST read ahead message count n > m When: qistraio> ReadAheadMsgCount

Reason: The number of messages which had read ahead(n) is greater than the user specified value ReadAheadMsgCount(m).

Action: If the number of messages which had read ahead is consistently higher than the ReadAheadMsgCount count, consider increasing the value of ReadAheadMsgCount. If the value of messages which had read ahead is higher than ReadAheadMsgCount, this could be due to a change in the workload patterns, and may indicate buffer pools are filling up, so causing pageset I/O

] MQQIST02W ... QIST Message read from disk n > m

When: qistgetd > kw.MsgsReadFromDisk

Reason: The number of messages which were got from disk(n) had read ahead is greater than the user specified value MsgsReadFromDisk(m).

Action: If the number of messages which were read from disk was consistently higher than the MsgsReadFromDisk count, consider increasing the value of MsgsReadFromDisk. If the value of messages which had read ahead is higher than MsgsReadFromDisk, this could be due to a change in the workload patterns, and may indicate buffer pools are filling up, so causing pageset I/O

Lock manager

The queue manager uses locks to prevent concurrent updates to resources.

This information is reported in ddname //LOCK. It is usually useful only to IBM.

See SCSQMACS(CSQDQLST) for the layout of the SMF record.

```
                                                Gets Already Held      Releases
MVCA MQPA 2013/01/06 09:00:00 VRM:701           6673            0          1934
MVCA MQPA 2013/01/06 09:15:00 VRM:701           4566            0          1359
```

Storage manager

The storage manager is responsible for managing virtual storage within the queue manager. The SMF records report on the real and virtual storage used by the queue manager. You can use these to monitor your usage. To convert from n 4KB pages to MB, the calculation is n *(4 *1024)/(1* 1024 * 1024)= n/256.

Overall storage usage

A summary of the overall storage usage for the queue manager is reported in //STGCSV

```
"z/OS",QM,Date,Time,VRM,">16MB",Bar,Real
MVCA,MQPV,2016/03/01,13:00:00,VRM:800,   391,  329,   82
MVCA,MQPV,2016/03/01,13:30:00,VRM:800,   389,  160,   62
```

These values are taken from z/OS control blocks when the SMF record is created. They are in MB.

">16MB"	Is the amount of storage used > 16 MB and below the bar. (qsrseloal + qsrsehial) /1MB.
Bar	Is the amount of 64 bit storage used above the bar. qsrsabytes/1MB.
Real	This is the real storage used by the address space. qsrsfmct/256.

Storage summary

A summary of the storage information is repored in ddname //STGSUM. Some fields contain information which is for IBM use only, these are grayed out.

See qsrs in member csqdsmfc.h for the layout of the SMF record.

```
Reason:  Statistics interval
>16MB Used   390 MB Free  1080 MB %used  26 delta    0 MB
Real Used                              68 MB
Pool global fixed,        200704,     0 MB
Pool global variable,      61440,     0 MB
Pool local fixed    ,     237568,     0 MB
Pool local variable,      528384,     0 MB
Total stack storage      3125248,     2 MB
Getmained          ,    11287880,    10 MB
Avail 64 storage   ,  2962227200,     2 GB
Cushion 64 bit     ,   321912832,     0 GB
<16MB Allocated user,   00014000,     0 MB
<16MB Allocated system  00068000      0 MB
<16MB Max size user  ,  009fa000,     9 MB
<16MB Low start     ,   00006000,
<16MB current user high,0001a000,
<16MB not used          ,            9 MB
Region >16MB size        ,         1471 MB
>16MB Allocated user  , 009a5000,     9 MB
```

```
>16MB Allocated system, 17d04000,   381 MB
>16MB Max limit        ,5bf00000,  1471 MB
>16MB Start of region  ,24100000,
>16MB Current High     ,24aac000,
>16MB not used              ,      1080 MB
QM available                    ,  1080 MB
Reserved for MVS       ,   1048576,   1 MB
Critical level         ,   1048576,   1 MB
SOS Cushion            , 104857600, 100 MB
ASID mem limit                  ,     0 GB
Object storage                      307 MB
Object storage hwater               247 MB
Number of objects           ,        78
Why limited:Set in the JCL
```

Where

```
Reason:  Statistics interval
```

Field:Reason.

The SMF record can be produced on normal statistics records are produced(reason = 1), typically half hourly, or when there has been a short on storage condition (reason = 2).

```
>16MB Used   390 MB Free  1080 MB %used  26 delta    0 MB
```

Storage > 16MB and below 31 bit bar.

1. used (qsrseloal+qsrsehial)/1MB
2. free (qsrsesiza - (qsrseloal + qsrsehial))/1MB
3. %used (100 * qsrsesiza) / 1MB
4. delta, the change since the previous record ((qsrsoeloal + qsrsoehial) -used)/1MB

```
Real Used                           68 MB
```

The real storage used. qsrsfmct/256.

The following fields are for IBM use only.

```
Pool global fixed,        200704,     0 MB
Pool global variable,      61440,     0 MB
Pool local fixed    ,     237568,     0 MB
Pool local variable,      528384,     0 MB
Total stack storage      3125248,     2 MB
Getmained          ,    11287880,    10 MB
```

```
Avail 64 storage   ,  2962227200,     2 GB
```

This is the amount of 64 bit storage available about the bar. qsrsaval64 /1GB

```
Cushion 64 bit     ,   321912832,     0 GB
```

qsrscush64/1GBThe records <16MB are for "24 bit storage".

```
<16MB Allocated user,   00014000,     0 MB
```

The address in hex of the start of the user area.

```
<16MB Allocated system  00068000      0 MB
```

The address in hex of the start of the system area.

```
<16MB Max size user  ,  009fa000,     9 MB
```

The maximum size of the user area in hex and in MB.

```
<16MB Low start     ,   00006000,
```

The start address of the user area in hex

```
<16MB current user high,0001a000,
```

The end of the user area in hex

```
<16MB not used          ,             9 MB
```

The fields >16MB are for 31 bit storage

```
Region >16MB size       ,          1471 MB
```

This is the total region size for 31 bit storage. qsrsesiza/1MB

```
>16MB Allocated user  , 009a5000,     9 MB
>16MB Allocated system, 17d04000,   381 MB
>16MB Max limit       ,5bf00000,  1471 MB
>16MB Start of region ,24100000,
>16MB Current High    ,24aac000,
>16MB not used              ,      1080 MB
```

These values are taken from z/OS control blocks. The fields used are qsrseloal/1MB, qsrsehial/1MB, qsrselim/1MB, qsrsestra, qsrsergtp, qsrsesiza - qsrseloal - pSS -> qsrsehial;

```
QM available                    ,  1080 MB
Reserved for MVS       ,   1048576,   1 MB
Critical level         ,   1048576,   1 MB
SOS Cushion            , 104857600, 100 MB
ASID mem limit                  ,     0 GB
```

QM Available. Field qsrsaval/1MB.

ASID mem limit is the address space memlimit value. Field qsrsmemlim_in_MB/1024 in GB.

```
Object storage                     307 MB
Object storage hwater              247 MB
Number of objects          ,        78
```

z/OS 64-bit virtual memory is organized as memory objects. The memory objects are allocated in 1 MB increments and are on a 1 MB boundary. Object storage is the size allocated. hwater is the High Water mark of used storage.

```
Why limited:Set in the JCL
```

This field is taken from a z/OS control block and has the following values. qsrsmemlims

1. Set by SMF or SMF default
2. Set in the JCL
3. Unlimited Region=0
4. Set by IEFUSI

Internal queue manager usage

This information is reported in ddname //STG. It is usually useful only to IBM.

See SCSQMACS(CSQDQSST) for the layout of the SMF record.

```
MVCA MQPA 2013/01/06 09:00:00 VRM:701
   Fixed pools   : Created         0,  Deallocated         0
   Fixed segments: Freed           0,  Expanded            2,  Contracted         1
   Varbl pools   : Created         0,  Deallocated         0
   Varbl segments: Freed           0,  Expanded            6,  Contracted         5
                   Getmains        1,  Freemains           1,  Non-zero RCs       0
                   SOS             0,  Contractions        0,  Abends             0
```

If Contraction (QSSTCRIT) is non zero, this indicates a severe problems with lack of storage within the queue manager.

If SOS (QSSTCONT) is non zero this indicates a problem was detected and recover actions were taken. This indicates a severe problem within the queue manager.

Internal queue manager storage pool usage

This output is for IBM use only. It describes the various pools of storage, and how much space they use.

```
          Storage       delta number    Chunk bsize #4k ! DSpace,Pool
 GFNN+,     57344,       8192,     1,    4096, 672,   6,!       ,ccb for ssam
```

Storage	Is the amount of storage used in the pool
delta	Is the change in storage since the previous SMF record
pool	The name of the storage pool

DB2 statistics

The DB2 statistics gives information about the DB2 tasks running in the queue manager.

The information is reported in ddname //DB2.

See SCSQMACS(CSQD5JST) for the layout of the SMF record.

The DB2 manager manages the interface with the DB2 database that is used as the shared repository.

When using shared queues, object definitions and other information are stored in DB2 tables.

DB2 requests are made from the queue manager by passing a request to a pool of server tasks that issue the DB2 request on behalf of the applications.

The figure below shows how DB2 requests are issued

Flow of a request for a DB2 service from a thread to server task

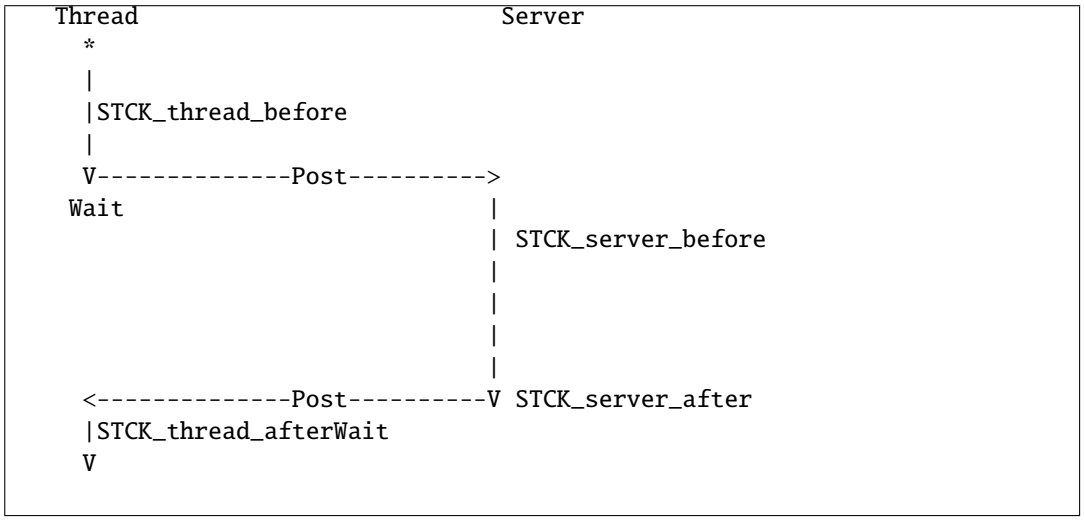

The processing for a thread wanting to issue a read request is as follows:

1. The thread puts a request onto a server work list.
2. The thread determines the current time (`STCK_thread_before`).
3. The thread posts a server task.
4. The thread waits.
5. The server task wakes up, and determines the current time (`STCK_server_before`).
6. The server takes the first request off the server work list and issues the DB2 request.
7. When the request has ended it posts the thread task.
8. The server task determines the current time (`STCK_server_after`) and updates the statistics:

   1. It increments the number of read requests `READCNT`.

   2. It calculates the time taken it took to process the request, `STCK_server_after-STCK_server_before` and adds this to the cumulative time `READSCUW`.

   3. If the time for the request was larger than the previous maximum if replaces the READSMXW with the delta.

Note: For other request, other counters are updated. These are `LIST*`, `UPDT*`, `DELE*`, and `WRIT*`.

1. The original thread wakes up and determines the current time (`STCK_thread_after`) and updates the statistics:

1. It calculates the time spent waiting (`STCK_thread_after - STCK_thread_before`) and adds this to the cumulative time `READTCUW`.
2. If the time spent waiting for the request was greater than the previous maximum it replaces the `READTMXW` with the larger value.

Note: For other request, other counters are updated. These are `LIST*`, `UPDT*`, `DELE*`, and `WRIT*`.

1. The thread continues processing.

The processing is similar for update, write, and delete requests. The list request is more complex and can result in reads being done from the server task issuing the list request.

Shared-channel-status and shared-sync-key tables

If you are using shared channels, shared-channel-status information and information about the shared-sync-queue are stored in DB2 tables. The fields with names starting SCS* are for DB2 selects, inserts, updates, and deletes from the shared-channel-status table. The fields with names starting SSK* are for DB2 selects, inserts, updates, and deletes for information about the shared-synch-key table.

The shared-sync-key table is used to locate the message id for messages on the shared sync queue. The Shared Channel Sync queue is used when the channel NPMSPEED(NORMAL) is used. Messages on the queue have information about the status of messages in a batch. The Shared Sync Key table, provides a mapping from channel name, XMITQ name, and remote queue manager name to the messages for the channel in the Shared Channel Sync queue.

Information is inserted into the Shared Sync Key table, when a channel processes messages with NPMSPEED(normal) for the first time. Both of these have times for the thread and the server, as described above.

```
MVCA MQPA 2013/01/15 16:30:00 VRM:701
 Tasks  : Servers    8, Active    9, Conns       0, Discs       0
          HighMax    3, Abend        0,  Requeue       0
               Count  Task avg   Task max   DB2 avg  DB2 max(ms)  (Task-DB2)Avg    Max
 List       :    180         2          6         2          5                0      0
 Read       :      1         3          3         3          3                0      0
 SCS Select :     15         2          4         2          4                0      0
```

The first column is the request type. These can be

List	Interpretation
List	This is when a query is done to the DB2 database, for example as a result of a display command
Read	Read of the definition of an object stored in DB2
Update	Update of the definition of an object stored in DB2
Write	Insert the definition of an object stored in DB2
SCS Select	Shared-channel-status table Select
SCS Insert	Shared-channel-status table Insert
SCS Update	Shared-channel-status table Update
SCS Delete	Shared-channel-status table Delete
SSK Select	Shared-synch-key Select
SSK Insert	Shared-synch-key Insert
SSK delete	Shared-synch-key Delete
Blob Select	Shared large message Select
Blob Insert	Shared large message Insert

List	Interpretation
Blob Update	Shared large message Update
Blob Delete	Shared large message Delete
Blob List	Shared large message List

Messages

MQQ5ST01S ... Q5ST Abend count n > 0 When: abndcnt > 0

Reason: The DB2 servers tasks have abended.

Action: There will be messages on the job log about why the server task abended. Review these and resolve the problem.

MQQ5ST02E ...Q5ST Retry count n > 0 When: requcnt > 0

Reason: A DB2 server task failed and the request was retried.

Action: There will be messages on the job log about why the server task abended. Review these and resolve the problem.

MQQ5ST04E Q5ST SCS Maximum rows returned on query When: scsmaxr > 0

Reason: More rows were returned on a query than fitted into a buffer

Action: None

MQQ5ST10S ... Q5ST Number of deadlock conditions n When: deadcnt > 0

Reason: Some deadlocks occurred in DB2, and a request was rolled back.

Action: If this happens frequently contact your IBM service representative.

MQQ5ST04W ... Q5ST DB2 Average read time n> m When: The number of DB2 requests is greater than 10, and readscuw/readcnt > DB2Time parameter

Reason: The average DB2 time for a read (n) is greater than the DB2 time passed as a parameter(m).

Action: Investigate any DB2 delays. You may need to set DB2Time to a more suitable value.

Accounting data describing the task

Information is displayed about the task.

For summary information see

1. TaskElapsed Time on page 88
2. TaskCSV on page 89.

See wtid in SCSQC370(CSQDSMFC) for the layout of the task ID information, wtas in SCSQC370 (CSQDSMFC) for the task layout, and wq in SCSQC370 (CSQDSMFC) for the queue specific layout.

For batch job:

```
MQ06 Batch Jobname:PAICECC  Userid:PAICE
```

This shows the data is for batch job PAICECC and userid PAICE and queue manager name MQ06.

For CICS transaction:

```
MQ01 CICS IYFFC001 opid:PAICE    userid:SCENSTC  Tran:CN15 task:0001664c
```

This shows the data is for CICS region IYFFC001, terminal userid PAICE CICS region userid SCENSTC, transaction name CN15, CICS task number 1664c on queue manager name MQ01.

For a channel:

```
MQ01 MOVER Jobname:MQ01CHIN Userid:POC005

Channel COLIN61A             ::ffff:9.20.5.21
```

This shows the data is for the mover MQ01CHIN with userid POC005 on queue manager name MQ01.

The channel name was COLIN61A with IP address ::ffff.9.20.5.21

Detailed information common to all types of applications

```
Start time Feb  1 13:24:13 2013 Started in a different time interval
Interval   Feb  1 13:27:13 2013 - Feb  1 13:27:13 2013 : 112.913245 seconds
Other reqs : Count               119
Other reqs : Avg elapsed time     36 uS
Other reqs : Avg CPU              22 uS
Other reqs : Total ET              0.004345 Seconds
Other reqs : Total CPU             0.002628 Seconds
> Latch 11, Total wait          0 uS, Waits          1, Name DMCSEGAL|SSSCONN
> Latch 12, Total wait      37806 uS, Waits         13, Name DMCNMSPC|XMCHASH
> Latch 16, Total wait       8113 uS, Waits        182, Name BMXL2   |RMCRMST |RLMARQC
> Latch 19, Total wait      74869 uS, Waits        196, Name BMXL3   |CFXML2  |SRH1_L19
> Latch 21, Total wait       6661 uS, Waits       2045, Name RLMLWRT
> Latch 32, Total wait          0 uS, Waits          2, Name SMCPHB
Longest latch wait at 0000000020b48598 74869 uS
Avg Latch time per UOW          2770 uS
Commit count                      44
Commit avg elapsed time         6061 uS
Commit avg CPU time               13 uS
Backout count                      2
Backout avg elapsed time      260912 uS
Backout avg CPU time          260912 uS
Log write count                19312
Log write avg et                 765 uS
Log I/O Bytes             1015644601
Log Force Count                19312
Log Force Avg elapsed time       765 uS
Suspend Count                     46
```

```
Suspend Avg elapsed time       17129 uS
Total suspend time                 0.787944 Seconds
Pages old                        46146
Pages new                      1275923
Log wait/Commit                   1023 uS
Log force/Commit                  1023 uS
Grand total CPU time           0.000117 S
Grand Elapsed time             0.004464 S
```

Interpretation of accounting data

Task related

The value in () at the start of the line is the minimum Detail level required to display the record.

```
( 1)  Grand total CPU time              n uS
```

This is the sum of all the CPU time for all MQ requests, used by the task.

```
( 1)  Grand Elapsed time                n uS
```

This is the sum of all of the elapsed times or all MQ requests, used by the task.

```
( 3) Start time s Started this interval
```

is the start time of the transaction in the format *Month day hh:mm:ss yyyy*. If it is in the current SMF period, then is also displays *Started this interval*. If it started in a different SMF interval it displays *Started in a different time interval.*

```
( 3) Interval   ss - se: n seconds
```

ss is the start time of this record, se is the end time of the interval. The format is Month day hh:mm:ss yyyy. N is the number of seconds in the interval.

```
(15) Avg Latch time per UOW             n uS
```

This is the total latch wait time / (count of commits + count of backouts).

```
( 5) Commit count                       n
```

This is the number of commits.

```
( 5) Commit avg elapsed time            n uS
```

This is the average time per commit. This is the elapsed total time for commits/count of commits.

```
( 5) Commit avg CPU time                n uS
```

This is the average CPU on the application thread for the CPU time. This is the application CPU time for commits/count of commits. Note, most of the commit activity is done on a SRB running in the queue manager address space, and the CPU used by this SRB is not recorded in this thread value.

```
( 5) Backout count                      n
```

This is the number of back out requests.

```
( 5) Backout avg elapsed time           n uS
```

This is the average time per backout. This is the elapsed total time for backout/count of backout requests.

```
( 5) Backout avg CPU time    %12.1u uS
```

This is the average CPU on the application thread for the CPU time. This is the application CPU time for backouts/count of backouts. Note,most of the backout activity is done on a SRB running in the queue manager address space, and the CPU used by this SRB is not recorded in this thread value.

```
(10) Log write count                    n
```

This is the number of writes to the log

```
(10) Log write avg et                    n uS
```

This is the average time to write to the log. This is the total write time/count of log requests. This is the total logging including the logging reported under PUT,PUT1, GET and SET.

```
(10) Log I/O Bytes                       n
```

This is the number of bytes logged by the application. This is the total logging including the logging reported under PUT,PUT1, GET and SET.

```
(10) Log Force Count                     n
```

```
(10) Log Force Avg elapsed time          n uS
```

```
(10) Suspend Count                       n
```

This is the number of times an MQ request was suspended within the queue manager. This may be for internal latches (used to serialise activity), waiting for logging or waiting for other activity.

```
(10) Suspend Avg elapsed time            n uS
```

This is the total suspend time/count of suspends.

```
(10) Total suspend time       n.nnnnnn Seconds
```

This is the total suspend time.

```
(10) Pages old                           n
```

This is the total number of requests to get a page from a buffer pool where the contents were needed. Put and gets both increment this number.

```
(10)Pages new                            n
```

This is the total number of requests to get a page from a buffer pool when a new page was needed. Only put requests increment this number.

```
(10) SMDS: Blocks written n, blocks read n
```

This is the number of SMDS blocks written or read for all of the queues.

```
(10) SMDS: Pages written n, pages read  n
```

This is the number of SMDS blocks written or read for all of the queues.

```
(10) SMDS: Pages read from cache         n
```

When a get request was issued, the data was in the cache and did not need to be read in.

```
(10) SMDS: Total wait time for I/O  n.n Seconds
```

```
(10) SMDS: I/O wait time per block  n.n Seconds
```

```
(10) SMDS: I/O wait time per page   n.n Seconds
```

```
( 4) == SRB CPU time used             n.n Seconds
```

When using publish/subscribe some work is run on an SRB task on behalf of the application. This is the total amount of SRB CPU time used.

```
(6) Other reqs : Count                   n
```

MQ Requests, MQOPEN, MQCLOSE, MQPUT, MQPUT1, MQGET have their own sections in the Queue accounting record. This field is the count of any other requests – including internal request, and for any requests to queues which cannot be recorded against a queue.

```
(6) Other reqs : Avg elapsed time     n uS
```

```
This is the total elapsed time for other/count of other requests.
```

```
(6) Other reqs : Avg CPU              n uS
```

```
This is the total CPU time for other/count of other requests.
```

```
(5) Other reqs : Total ET             n.n Seconds
```

This is the total elapsed time for other requests

```
(5) Other reqs : Total CPU            n.n Seconds
```

This is the total CPU time for other requests.

```
( 4) == DB2 activity : n requests
```

This is the number of requests to DB2 as part of offloading shared queue messages to DB2

```
( 6) > Average time per DB2 request-Server : n uS
```

This is the average time of the requests to DB2 as part of offloading shared queue messages to DB2, as seen by the server.

```
( 6) > Average time per DB2 request-Thread : n uS
```

This is the average time of the requests to DB2 as part of offloading shared queue messages to DB2, as seen by the thread. If this is significantly larger than the average time as seen by the server, then there has been some delay due to a shortage of DB2 threads within the queue manager.

```
( 6) > Maximum time per DB2 request-Server : n uS
```

```
( 6) > Maximum time per DB2 request-Thread : n uS
```

```
( 6) > Bytes put to DB2    : n
```

This is the number of bytes written to DB2 when offloading large shared queue messages to DB2.

```
( 6) > Bytes read from DB2 : n
```

This is the number of bytes read from DB2 when getting shared queue messages that have been offloaded to DB2.

```
( 6) == CF activity : Requests - Single n, Multiple m
```

When using shared queues there are requests which update one entry in the coupling facility. There are other requests which update multiple entries within the CF – for example a commit with several message in the unit of work.

So the length of time for the update of a single entry should be consistent, but the time to update multiple entries will depend on the number of entries being updates.

```
( 6) >                Retries  - Single %10.1u,  Multiple %
```

If the CF was 'busy' then a request to the CF may need to be retried. This is a count of the number of retries.

```
( 6) > Average time per single requests : n uS
```

This is a measure of the CF response time when updating single entries. It is the total time spent doing IXLLSTE requests/count of IXLLSTE requests.

```
( 6) > Average time per multiple requests : n uS
```

This is a measure of the CF response time. when updating multiple entries. It is the total time spent doing IXLLSTM requests/count of IXLLSTM requests. This time will depend on the number of entries updated.

```
(10) == Page set 0 activity  : Count n,  Avg elapsed n
```

This is the number of puts and gets which mapped to page set 0.

```
(15) > Latch n Total wait n uS, Waits n, Name sss
```

This is information useful to IBM about the latches used within the queue manager to serialise work.

```
(15) Longest latch wait at xxx %lld uS
```

This is information useful to IBM about the latches used within the queue manager to serialise work.

Accounting data for a task using local queues

Example data for put and get to a local queue.

```
  Open name                              CP0000
  Queue indexed by MSG_ID                CP0000
  First Opened   Jan  8 13:11:36 2013     CP0000
  Last Closed    Jan  8 13:11:36 2013     CP0000
  Page set ID                    2       CP0000
  Buffer pool                    2       CP0000
  Current opens                  0       CP0000
  Total requests                 6       CP0000
  Open Count                     2       CP0000
  Open Avg elapsed time         10 uS    CP0000
  Open Avg CPU time             10 uS    CP0000
  Close count                    2       CP0000
  Close avg elapsed time         3 uS    CP0000
  Close avg CPU time             3 uS    CP0000
  Get count                      1       CP0000
  Get avg elapsed time          29 uS    CP0000
  Get avg CPU time              29 uS    CP0000
  Get skipped message count      2       CP0000
  Get TOQ average             3103 uS    CP0000
  Get TOQ maximum             3103 uS    CP0000
  Get TOQ minimum             3103 uS    CP0000
  Get valid count                1       CP0000
  Get size maximum            1000 bytes CP0000
  Get size minimum            1000 bytes CP0000
  Get size average            1000 bytes CP0000
  Get Dest-Specific              1       CP0000
  Get Persistent count           1       CP0000
  Put count                      1       CP0000
  Put avg elapsed time          33 uS    CP0000
  Put avg CPU time              33 uS    CP0000
  Put + put1 valid count         1       CP0000
  Put size maximum            1000       CP0000
  Put size minimum            1000       CP0000
  Put size average            1000       CP0000
  Put num persistent             1       CP0000
  Curdepth maximum               3       CP0000
  Total Q CPU used              92 us    CP0000
  Total Queue elapsed time      92 us    CP0000
```

Interpretation of the data

The value in () is the minimum value of Detail to display the data.

Opening a queue

`( *) Open name` `queueName`

This is the name of the queue opened by the application.

`( *) Base name` `queueName`

This is the name of the base queue if different to the opened name - for example an alias queue.

`(10) Queue indexed by MSG_ID` `queueName`

How the queue is indexed.

```
(10)  First Opened   Jan  8 13:11:36 2013      queueName
(10)  Last Closed    Jan  8 13:11:36 2013      queueName
( 5)  Page set ID                   n         queueName
( 5)  Buffer pool                   n         queueName
(14)  Current opens                 n         queueName
(14)  Total requests                n         queueName
(14)  Generated messages            n         queueName
```

Page set ID and *Buffer pool* are only displayed if the number of valid puts or the number of valid gets are > 0. A valid put or valid get sets these values in the SMF records.

How many trigger or event messages were generated. The cost of producing these message will be reported by this queues.

```
(10) Open Count                     n         queueName
```

How many times this queue was opened by this application.

```
(10) Open Avg elapsed time          n uS      queueName
```

The average elapsed time to open the queue. This is calculated as total open elapsed time/count of open requests.

```
(10) Open Avg CPU time              n uS      queueName
```

The average CPU time to open the queue. This is calculated as total CPU time for open/count of open requests.

```
(10) Open avg suspend time          n uS      queueName
```

If the open request was suspended, the average time the requests was suspended by the queue manager. This is calculated as opensuet/count of open requests.

```
(10) Open avg topic srb time        n uS      queueName
```

When a topic is opened it may attach an SRB to do some processing. This is the average SRB time per open request. It is calculated at the total SRB for open/count of open requests.

Closing a queue

```
(14) Close count                    n          queueName
(14) Close avg elapsed time         n uS      queueName
```

The average elapsed time to close the queue. This is calculated as total close elapsed time/count of close requests.

```
(14) Close avg CPU time             n uS      queueName
```

The average CPU time to close the queue. This is calculated as total CPU time for close/count of close requests.

```
(10) Close avg suspend time         n uS      queueName
```

If the open request was suspended, the average time the requests was suspended by the queue manager. This is calculated as closesuet/count of open requests.

```
(15) Close CF access                n         queueName
```

The number of times a close caused a request to the CF to update the status in the Coupling Facility. This occurs when this is the last close on the queue manager, and the queue manager updates the CF to indicate that it does not have the queue open.

```
(15) Close No CF access             n         queueName
```

The number of closes which did not cause a CF update, because other applications on this queue manager have the queue open.

```
(15) Close topic srb CPU time       n         queueName
```

When a topic is closed it may attach an SRB to do some processing. This is the average SRB time per close request. It is calculated at the total SRB for close/count of close requests.

Getting from a queue

```
(5) Get count                          n          queueName
```

The total number of get requests from the queue.

```
( 5) Get avg elapsed time              n uS       queueName
```

The average elapsed time getting messages from the queue. This is calculated as the total elapsed time for gets/count of get requests.

```
( 5) Get avg CPU time                  n uS       queueName
```

The average CPU time getting messages from the queue. This is calculated as the total CPU time for gets/count of get requests.

```
( 5) Get suspended time                n uS       queueName
```

The average time per message that the request was suspended within MQ. This includes logging of out of syncpoint gets, waits for internal latches etc.

```
( 5)  Get skipped message count        n        queueName
```

If a get request is made for a specific message, and the queue is not indexed, the queue is searched for the specified message. This field is the number of messages skipped over. This can also occur if another application is getting a message, so the message is unavailable. In this case the number of skipped messages per get should be small.

```
(10) Get pageset total count           n        queueName
```

The number of get requests which had to read from the page set.

```
(10) Get pageset elapsed time          n        queueName
```

If the get pageset count > 0, this is the average time spent getting the messages from the page set.

```
(15) Get log force count               n        queueName
```

Some gets can cause a log force, for example, if the buffer pool is critically short of buffers. This value is the number of times a log force occurred.

```
(15) Get log force elapsed time        n        queueName
```

If a log force occurred during a get, this is the average time per message doing a log force.

```
(17) Get log write count            n       queueName
```

This is the number of log write requests.

```
(17) Get log write elapsed time     n       queueName
```

If the application had to log data during the get request, this is the time taken for this write request. For requests within sync point, this just writes to log buffers.

```
(14) Get total empty pages          n       queueName
```

If all the messages on a page have been got, but the page has not been freed up, the page counts as empty. A large value can indicate an application getting a specific message from a deep queue, and the queue is not indexed.

```
(14)  Get TOQ average                  n uS    queueName
```

This is the average Time On the Queue. This is the time from when the message was put to the queue on this queue manager, to the time it was destructively got.

```
(14)  Get TOQ maximum                  n uS    queueName
```

This is the maximum time on the queue. This is the time from when the message was put to the queue on this queue manager, to the time it was destructively got.

```
(14)  Get TOQ minimum                    n uS    queueName
```

This is the minimum time on the queue. This is the time from when the message was put to the queue on this queue manager, to the time it was destructively got.

```
(14)  Get valid count                    n       queueName
```

This is the number of gets requests with return code zero, or truncated message accepted.

```
(10)  Get size maximum                   n bytes queueName
```

This is the maximum size of message processed.

```
(10)  Get size minimum                   n bytes queueName
```

This is the minimum size of message processed.

```
(10)  Get size average                   n bytes queueName
```

This is the total number of bytes processed/count of valid gets(rc equal 0 or MQRC_truncated_msg_accepted)

```
(10)  Get Dest-Specific                  n       queueName
```

The number of times a destructive get for a specific message id or correld-id was issued.

```
(10)  Get Dest-Next              n          queueName
```

The number of times a destructive get for the first or next message was issued.

```
(10)  Get Browse-Specific                n       queueName
```

The number of times a browse for a specific message id or correldid was issued

```
(10) Get QSUBrowse-Next             n          queueName
```

The number of times a browse for the first or next message was issued

```
(15) Get log force elapsed time          n       queueName
```

If a log force occurred during a get, this is the average time per message doing a log force.

```
(10) Get errors                             n      queueName
```

This is for get requests which had reason codes MQRC_OPTIONS_ERROR,MQRC_GMO_ERROR, or MQRC_HOBJ_ERROR

```
(10) Get persistent count                n       queueName
```

The number of persistent messages processed

```
(10) Get non persistent count            n       queueName
```

The number of non persistent messages processed. This is calculated as count of valid messages, count of persistent messages

Putting to a queue

```
( 5) Put count                           n       queueName
```

The number of puts to the queue.

```
(5)  Put avg elapsed time                n uS    queueName
```

The average elapsed time putting messages to the queue. This is calculated as the total elapsed time for puts/count of put requests.

```
(5)  Put avg CPU time                    n uS    queueName
```

The average CPU time putting messages to the queue. This is calculated as the total CPU time for puts/count of put requests.

```
(15) Put suspended time              n uS       queueName
```

The average time per message that the request was suspended within MQ. This includes logging of out of syncpoint puts, waits for internal latches etc.

```
(15) Put pageset count               n          queueName
```

The number of put requests which put directly to the page set. If this value is non zero then the buffer pool is likely to have been constrained.

```
(15) Put pageset elapsed time        n          queueName
```

If the put pageset count > 0, this is the average time spent putting the messages to the page set.

```
(15) Put log force count             n          queueName
```

Some puts can cause a log force, for example, if the buffer pool is critically short of buffers. This value is the number of times a log force occurred.

```
(15) Put log write total count       n          queueTime
```

This is the number of log write requests.

```
(15) Put log write elapsed time      n          queueName
```

If the application had to log data during the put request, this is the time taken for this write request. For requests within syncpoint, this just writes to log buffers.

```
(15) Puts direct to a SQ             n          queueName
```

With messages coming in from the CHINIT, messages can be routed to a queue manager in the QSG using IGQ, or can be put directly. This field has the number of messages which were put directly to the shared queue and did not use IGQ. This avoids the IGQ double-hop where the message is put to a queue, IGQ processes it, and puts it on another queue manager. See *Puts to IGQ XMITQ* below.

```
(15) Puts to IGQ XMITQ               n          queueName
```

With messages coming in from the CHINIT, messages can be routed to a queue manager in the QSG using IGQ, or can be put directly. This field has the number of messages using IGQ to route to another queue manager. See *Puts direct to a SQ* above.

```
(15) Put + put1 valid count          n          queueName
```

This is the number of MQPUT +MQPUT1 requests which were successful, and had return code MQCC_OK.

```
(20) Put + put1 not valid count      n          queueName
```

This is the number of MQPUT +MQPUT1 requests which were not successful, and had return code $>$ MQCC_OK.

```
(10) Puts with put delay             n          queueName
```

With JMS 2.0, you can specify a delivery delay when sending a message. The queue manager does not deliver the message until after the specified delivery delay has elapsed. This is the number of messages put specifing message delay.

```
(15) Put waiting getter              n          queueName
```

This is the number of put requests which satisfied a waiting getter.

```
(15) Put topic srb CPU time          n          queueName
```

When putting to a topic, it may attach an SRB to do some processing. This is the average SRB time per put request. It is calculated at the total SRB for put/count of valid put requests.

```
( 5) Put1 count                      n          queueName
```

The number of puts to the queue.

```
( 5) Put1 avg elapsed time           n uS       queueName
```

The average elapsed time putting messages to the queue. This is calculated as the total elapsed time for puts/count of Put1 requests.

```
( 5)  Put1 avg CPU time                n uS    queueName
```

The average CPU time putting messages to the queue. This is calculated as the total CPU time for puts/count of Put1 requests.

```
(15)  Put1 suspended time              n uS    queueName
```

The average time per message that the request was suspended within MQ. This includes logging of out of syncpoint gets, waits for internal latches etc.

```
(15) Put1 pageset count                n       queueName
```

The number of Put1 requests which Put1 directly to the page set. If this value is non zero then the buffer pool is likely to have been constrained.

```
(15) Put1 pageset elapsed time         n       queueName
```

If the Put1 pageset count > 0, this is the average time spent putting the messages to the page set.

```
(15) Put1 log force count              n        queueName
```

Some puts can cause a log force, for example, if the buffer pool is critically short of buffers. This value is the number of times a log force occurred.

```
(15) Put1 log force elapsed time       n        queueName
```

If a log force occurred during a Put1, this is the average time per message doing a log force.

```
(15) Put1 log write total count       n       queueTime
```

This is the number of log write requests.

```
(15) Put1 log write elapsed time       n        queueName
```

If the application had to log data during the Put1 request, this is the time taken for this write request. For requests within syncpoint, this just writes to log buffers.

```
(15) Put1 + put1 valid count           n       queueName
```

This is the number of MQPUT +MQPUT1 requests which were successful, and had return code MQCC_OK.

```
(15) Put1 waiting getter               n        queueName
```

This is the number of put1 requests which satisfied a waiting getter.

```
(15) Put1 topic srb CPU time           n        queueName
```

When putting to a topic, it may attach an SRB to do some processing. This is the average SRB time per Put1 request. It is calculated at the total SRB for put1/count of valid put1 requests.

```
(10) Put size maximum                  n       queueName
```

This is the maximum message size from a put or PUT1 request.

```
(10)  Put size minimum                 n       queueName
```

This is the minimum message size from a put or PUT1 request.

```
(10) Put size average                 n       queueName
```

This is the total number of bytes processed/count of valid puts or put1s (rc = 0).

```
(10) Put num persistent                n       queueName
```

This is the number of persistent messages put to the queue.

```
(10) Put num not persistent             n       queueName
```

This is the number of non persistent messages put to the queue. It is calculated as number of puts, number of persistent messages.

```
(10 Published msgs                    n       queueName
```

This is the number of messages which resulted in a publish. A put or a put1 can result in 0 or more messages published.

MQINQ on a queue

```
 5) Inq count                          n          queueName
```

The number of inquires to the queue.

```
(5)  Inq avg elapsed time              n uS       queueName
```

The average elapsed time doing MQINQ to the queue. This is calculated as the total elapsed time for MQINQ/count of MQINQ requests.

```
(5)  Inq avg CPU time                  n uS       queueName
```

The average CPU time doing MQINQ to the queue. This is calculated as the total CPU time for MQINQ/count of MQINQ requests.

MQSET to a queue

```
( 5) Set count                         n          queueName
```

The number of MQSET requests to the queue

```
( 5) Set avg elapsed time              n uS       queueName
```

The average elapsed time doing MQSET to the queue. This is calculated as the total elapsed time for MQSET/count of MQSET requests.

```
( 5) Set avg CPU time                  n uS       queueName
```

The average CPU time doing MQSET to the queue. This is calculated as the total CPU time for MQSET/count of MQSET requests.

```
(15) Set log force elapsed time        n          queueName
```

This is the average time per message doing a log force.

```
(15) Set log write elapsed time     n        queueName
```

This is the average time per message doing a log write.

Other information

```
(10) Curdepth maximum                  n          queueName
```

This is the maximum depth of the queue found during the interval.

```
(10) Number selector calls             n          queueName
```

This is the MQSUB requests where a selector string was specified.

```
(10) Max selector length               n          queueName
```

This is the maximum selector length specified when there was an MQSUB requests and where a selector string was specified.

```
(10) MQCB requests MQCB requests       n          queueName
```

This is the number of MQCB requests issued.

```
(10) MQCB avg elapsed time             n          queueName
```

This is the average elapsed time of the MQCB requests.

```
(10) MQCB avg CPU time                 n          queueName
```

This is the average CPU time of the MQCB requests.

```
( *) Total Q CPU used                  n us       queueName
```

This is the total of the CPU used to process this queue. It is the sum of the CPU used in the API requests.

```
( *) Total Queue elapsed time         n us      queueName
```

This is the total elapsed time processing this queue. It is the sum of the Elapsed time used in the API requests.

Additional information for shared queues

```
Open CF access                    2       APP1
Open No CF access                 0       APP1
Close CF access                   2       APP1
Close No CF access                0       APP1
```

Interpretation

```
(15) Open CF access                    2      APP1
(15) Open No CF access                 0      APP1
(15) Close CF access                   2      APP1
(15) Close No CF access                0      APP1
```

If the queue manager does not currently have a shared queue open, then when an application opens the queue, the queue manager has to update the CF to indicate it has the queue open. This is recorded as *Open CF access.* If there is an application with the queue open, then with another open request, the queue manager does not need to update the CF. This is recorded as *Open No CF access.* The Open CF access has a higher cost than the Open No CF access.

When the queue is closed, if there are other applications with the queue open, then there is no CF update. When the last application closes a queue, the queue manager updates the CF to indicate it does not have the queue open. There are two fields *Close CF No access* and *Close CF access.* A close with the *Close CF access* updates the CF and this has a slightly higher cost than a close with no CF access.

```
(15) Open CF access                    2      APP1
(15) Open No CF access                 0      APP1
(15) Close CF access                   2      APP1
(15) Close No CF access                0      APP1
```

If the queue manager does not currently have a shared queue open, then when an application opens the queue, the queue manager has to update the CF to indicate it has the queue open. This is recorded as Open CF access. If there is an application with the queue open, then with another open request, the queue manager does not need to update the CF. This is recorded as Open No CF access. The Open CF access has a higher cost than the Open No CF access.

When the queue is closed, if there are other applications with the queue open, then there is no CF update. When the last application closes a queue, the queue manager updates the CF to indicate it does not have the queue open. There are two fields Close CF No access and Close CF access. A close with the Close CF access updates the CF and this has a slightly higher cost than a close with no CF access.

Explanation of CF time/Verb If there were 100 get requests of which 5 of these accessed the CF. The total time doing the CF requests was 50 microseconds. Using this data

- CF time/verb is 50 uSeconds/100 = 0.5 uSeconds.
- CF ... number of requests = 5
- CF ... CF response time = 50 uSeconds/5 = 10 uSeconds.

If the CF request was synchronous then the duration of the request contributes to the CPU used.

```
Get count                        1000       APP1
Get avg elapsed time             1096 uS    APP1
Get avg CPU time                   56 uS    APP1
CF time/verb    966
 CF Avg Sync elapsed time/verb      41 us
```

```
 CF Avg Sync number of request     140
 CF Avg Sync CF response time      294 us
 CF Avg Async elapsed time/verb    925 us
 CF Avg Async number of request    861
 CF Avg ASync CF response time    1074 us

  StartMon Avg Async elapsed time/verb      1 us
  StartMon Avg Async number of request      1
  StartMon Avg ASync CF response time    1253 us
  Move     Avg Sync elapsed time/verb      41 us
  Move     Avg Sync number of request     140
  Move     Avg Sync CF response time      294 us
  Move     Avg Async elapsed time/verb    924 us
  Move     Avg Async number of request    860
  Move     Avg ASync CF response time    1074 us
```

Interpretation

```
(10) CF time/verb    966
```

The Get avg elapsed time above was 1096 microseconds. Of the average elapsed time of 1096 microseconds per get request, 966 microseconds were spent in a coupling facility request

```
(11) CF Avg Sync elapsed time/verb      41 us
(11) CF Avg Sync number of request     140
(11) CF Avg Sync CF response time      294 us
(11) CF Avg Async elapsed time/verb    925 us
(11) CF Avg Async number of request    861
(11) CF Avg ASync CF response time    1074 us
```

A coupling facility request can be synchronous or asynchronous. The times and counts is displayed for each type of request. The total CF requests were 140 + 861 = 1001. Of these 140 requests were synchronous. The average time spent doing synchronous requests was 294 microseconds.

Looking at the total number of get requests and allocating the synchronous and asynchronous time, on average each get request had 41 microseconds synchronous time, and 925 microseconds asynchronous. In reality a get was either synchronous or asynchronous, not a mixture. But this shows that most of the time of the get request 925 out of 1096 was spent in asynchronous requests.

```
(15)   StartMon Avg Async elapsed time/verb      1 us
(15)   StartMon Avg Async number of request      1
(15)   StartMon Avg ASync CF response time    1253 us
(15)   Move     Avg Sync elapsed time/verb      41 us
(15)   Move     Avg Sync number of request     140
(15)   Move     Avg Sync CF response time      294 us
(15)   Move     Avg Async elapsed time/verb    924 us
(15)   Move     Avg Async number of request    860
(15)   Move     Avg ASync CF response time    1074 us
```

This information is likely to be of use only to IBM personnel. The information displays the time at the detailed CF request level. The information shows there was 1 StartMonitor request. This was an asynchronous request taking 1253 microseconds. There were Move requests, some of which were synchronous and some were asynchronous.

This information is likely to be of use only to IBM personnel. The information displays the time at the detailed CF request level.

The information shows there was 1 StartMonitor request. This was an asynchronous request taking 1253 microseconds.

There were Move requests, some of which were synchronous and some were asynchronous.

QALL

This file has the same information as described above, but summarised by queue name. If there were 100 CICS transactions processing messages on a queue. There will be one section within this file for the queue, and it will report 100 opens and 100 gets. You can use this file to identify the queues which use a lot of CPU, or have a high activity.

See page 77 above.

Rules for accounting data

MQTASK01x Queue not indexed Detail: 5

When: There are gets or browse from the queue where a message is got by msgid or correlid, but the queue is not indexed. If the queue is not indexed, this causes a sequential scan of the queue which can be expensive, the deeper the queue, the more expensive the request.

MQTASK01I is produced if the max queue depth is less than 10

MQTASK01W is produced if the max queue depth is greater equal than 10, and less than 100

MQTASK01E is produced if the max queue depth is greater equal than 100

Action: Review the queue definition, work with the applications team to determine how the queue should be indexed.

MQTASK02W High percent of no msg found Detail: 12

When: The number of gets which did not return a message is more than 3 times the number of gets which returned a message.

When waiting for a particular message there will be a get for the message, and the message is not there, so the application waits. When the message arrives the application is posted and the get is reissued and retrieves the message-id.

Having a high proportion of requests not returning a message can indicate

1. A get next message is wanted, but the message-id or correlid field is not being cleared, and so the requests is a get for a specific message.
2. There are many application instances getting from the queue. A message arrives, all the applications rush to get the message. One instance is successful, the other instances are unsuccessful. This may not be a problem but may indicate you may need to look at the way the applications are set up.

Action: Review the applications.

MQTASK03 long open Detail 10:

When: The average open time is greater than the Long_open parameter in the MQSMF formatting program.

Action: This may alert you to possible application problems. You should investigate the other data for the record, for example to see if there is a high latch wait – which would indicate contention with other applications, or with a long CF response time.

MQTASK04I long open ET >2 CT Detail:10

When: Usually the elapsed time of an open request is close to the CPU used. This reports if the elapsed time is much longer than the CPU time used for the open.

Action: Investigate the delays.

MQTASK05I No msg returned Detail: 8

When: There were gets from the queue, but none were successful. This is valid, but may indicate a problem with the application setup.

Action: Review the application.

MQTASK07I High browse rate/get ratio When: There were gets and browses from the queue, and the number of browse requests is 3 times the number of get requests. This is valid, but may indicate a problem with the application setup.

Action: Review the application.

MQTASK08E Long Put time due to logging Detail: 12

When: The average time for a put was longer than the Long_Put parameter in the MQSMF formatting program, and the average message size was less than 10,000 bytes. More than 20% of the time was spent logging.

Action: Review the logging statistics and the logging datasets to see if the system is constrained by the DASD.

MQTASK09E Long Put time not due to logging Detail: 15

When: There average time for a put was longer that the Long_Put parameter in the MQSMF formatting program, and the average message size was less than 10,000 bytes. Less than 20% of the time was spent logging.

Action: You should examine the detailed record and find why the put took so long.

MQTASK10W Max depth > n and puts Detail: 10

When: There maximum depth of the queue was greater than MaxDepth parameter in the MQSMF formatting program, and there were puts to the queue.

Action: Investigate why the queue depth is so large.

MQTASK12x Get Specific and Get Next Detail: 5

When: There is a mixture of Get Specific and Get Next requests from an application. This is valid, but may not be as designed.

MQTASK12I - the max depth is less than 10 messages

MQTASK12I -the max depth is greater equal to 10 and less than 100

MQTASK12I -the max depth is greater equal 100

Action: Investigate the application.

MQTASK16S long latch wait n Name s Detail: 15

When: A latch wait time was found where the latch time was greater than the 10* `LongLatchWait` parameter in the MQSMF formatting program.

Action: This is usually of interest only to IBM personnel when investigating performance problems

MQTASK16E long latch wait n Name s Detail: 10

When: A latch wait time was found where the latch time was greater than the `LongLatchWait` parameter in the MQSMF formatting program.

Action: This is usually of interest only to IBM personell when investigating performance problems

MQTASK18W Multiple put1s Detail: 10

When: Multiple MQPUT1 Request were made within an application. If multiple messages have to be put, it is more efficient to use MQOPEN, MQPUT multiple times and MQCLOSE.

Action: Review the application.

MQTASK35W Waiting for long shunt, task id Detail: 1

When: The task record has waiting for log shunt.

This occurs when there are many puts or gets before a commit, or a long time interval between a put or a get and the commit. The time is the time waiting for the log shunt activity.

Action: Review the application.

MQTASK36W Did BP n fill up for puts to queue Detail: 10

When:There is writes to pagesets during an MQPUT or MQPUT1. This usually indicates the buffer pool was $>$ 95% full so had syncronous writes to the page set and the log.

Action: Review the buffer pool size.

MQTASK37W Gets from local queue are more expensive than puts queueName Detail: 10

When: Normally an MQGET from a queue uses less CPU than an MQPUT. When displaying the summary of the queue usage (QALL) the average CPU time for gets $>$ average CPU time for puts. The puts and gets are typically done in different transactions.

Action: Investigate why the gets are using more CPU. This can be caused by get requests having to read from the page set, a CF request took much longer, or if multiple messages had to be scanned to locate the particular message.

MQTASK38W Task had long DB2 request over n ,

Detail: 10

When:The task had average response time for DB2 requests larger than the DB2Time value.

Action: Depending on the average elapsed time and the number of requests you may want to investigate the DB2 accounting data.

TaskElapsed Time

This section has a summary of the tasks records, and where the time is spent.

This is written to ddname //TASKET.

The fields columns are

CPU This is the sum of the CPU recorded.

LogLatch This is the time waiting for the log latch

Latch This is the time waiting for other latches

Log This is the time waiting for log I/O activity – typically commits and out of syncpoint puts of persistent messages

PutPS The time putting to a page set

GetPS The time getting from a page set

Delta The difference between the total time and the sum of the CPU, Log latch etc.. This could be application time, or time spent in MQ which is not recorded in the class 3 accounting because ACCTQ is off

TotalElapsedTime,

Duration

This allows you to quickly see where the time is spent for an application, for example logging or doing page set I/O

TaskCSV

This section has a summary of the tasks records, one per task/transaction name and CPU used, amount of data logged, put and got in MB. Data for multiple instances of a transaction or a job are summarised into one record.

This is written to ddname //TASKCSV.

The time is the time the SMF record was created, rounded down to the hour boundary. So if a transaction ran at 10:04:02.20 it would be recorded as 10:00:00

```
Date','Time','Type','Tran1','Tran2','Count','CPU S','logBytesMB','put MB','Get MB',

2013/02/15,10:00:00,B,'PAICEP4','         ',    2,4.3,2096.2,  2000,    0,
2013/02/15,10:00:00,C,'CN15    ','IYFFC000',32091,2.3,    0,    20.1,30.2,
2013/02/15,10:00:00,C,'CP15    ','IYFFC000',  496,0.1,1.1697,   1.0,    0,
2013/02/15,10:00:00,M,'P1.TO.P2','MQP2<1414>', 10,0.9,23.57,    23.0,   0,
```

Interpretation

Date	
Time	
Type	B for Batch, RB for RRS batch, C for CICS, M for Mover, CS for command server, I for IMS
Tran	Batch job name, CICS transaction, Channel name
Tran2	CICS region name, IP address
Count	Number of task records found for this row
CPU S	CPU used in seconds
logBytesMB	Total amount of data logged in MB
put MB	Total amount of data put, in MB
Get MB	Total amount of data get, in MB

Task Sumary

In ddname //TASKSUM is a summary of messages produced when looking at the task and queue records.

This has data like

```
Record#    Count    Value Message
   2202       25    98908 MQTASK13E long commit time C,'CP15','IYFFC000',
     38        1   106347 MQTASK13E long commit time B,'PAICEP7A',' ',
```

This has the following meaning.

- There was a message `MQTASK13E long commit time C,'CP15','IYFFC000',`
  - It was produced 25 times
  - The largest value (of the commit time) was 98908 microseconds. This was at record 2202 in the input file
- There was a message `MQTASK13E long commit time B,'PAICEP7A',' ',`
  - It was produced 1 time,
  - the largest (only) value was 106347 in record 38 of the input file.

To investigate these in more detail you can use StartRecord=2202, LastRecord=2202 and Detail(20). This will give all maximum level of detail for the one record.

Put CSV information

The PUTSCSV reports information about puts for each task and queue.

The PUTSCSVS reports information about puts summarised for tasks using the queue.

The headers for the PUTSCSV are

MVS	The name of the z/OS LPAR
QM	The name of the queue manager
SMF_Date	The date from the SMF header
SMF_Time	The time from the SMF header
TranT	The type of the transaction. M for mover, B for Batch, C for CICS, I for IMS
Tran1	Task identification, such as channel name
Tran2	Additional information about the task, such as IP address

The following headings are common to both PUTSCSV an PUTSCSVS

Queue	The name of the queue
Puts	The number of put requeusts
Put1s	The number of put1 requests
TotBytes	The total number of bytes put to the queue
MaxMsgSz	Maximum message size
MinMsgSz	Minimum message size

Get CSV information

The GETSCSV reports information about gets for each task and queue.

The GETSCSVS reports information about gets summarised for tasks using the queue.

The headers for the GETSCSV are

MVS	The name of the z/OS LPAR
QM	The name of the queue manager
SMF_Date	The date from the SMF header
SMF_Time	The time from the SMF header
TranT	The type of the transaction. M for mover, B for Batch, C for CICS, I for IMS
Tran1	Task identification, such as channel name
Tran2	Additional information about the task, such as IP address

The following headings are common to both GETSCSV an GETSCSVS

Queue	The name of the queue
TotalGets	
ValidGets	The number of gets which returned a message
TotBytes	The total number of bytes put to the queue
MaxMsgSz	Maximum message size
MinMsgSz	Minimum message size
MaxLat(us)	Maximum time on queue (latency) in microseconds
MinLat(us)	Minimum time on queue (latency) in microseconds
AvgLat(us)	Average time on queue (latency) in microseconds

Address space level storage usage

The address space storage usage statistics give information about virtual storage and real storage used by the queue manager.

The output is displayed in the ddname //STGSUM.

For the details about the fields see the layout for CSQDQSRS in SCSQC370(CSQDSMFC).

```
>16MB Used   397 MB Free  1080 MB %used  26 delta    4 MB
 Bar   Used   237 MB               %used   0
 Real  Used                         1010 MB
 why limited:Set in the JCL
```

`>16MB` shows 397 MB of the region's virtual storage has been used. There is 1080MB left in the region. The amount used is 26% of the total virtual storage. Since the last storage SMF record there has been a 4MB increase in virtual storage usage.

`Bar` shows the storage usage above the bar.

`Real` shows the real storage usage.

`why limited`: This what limits the 31 bit storage. The value (raxlvmemlims) is extracted from a z/OS control block. The common values of this are:

- Set by SMF or SMF default
- Set in the JCL
- Unlimited Region=0
- Set by IEFUSI
- Set by Auth interface

Other information displayed is likely to be of use to IBM only.

Additional messages produced in MESSAGE output file.

MQVSTG01E virtual storage usage > 95%

When: 31 bit storage use is > 95% of the available storage use.

Action: Review storage usage. Consider making buffer pools smaller.

Using detail(20) will provide additional information.

MQVSTG02W virtual storage usage > 90%

When: 31 bit storage use is > 90% of the available storage use.

Action: Review storage usage. If start to get this message, you should monitor the storage usage. If your storage continues to increase, you should start reviewing how you are using your buffer pools and perhaps reduce the size of them.

-End of document-