

IBM MQ for z/OS Health Check

Data Collection Guide



Conducted by:

IBM Washington Systems Center

Overview	4
Scope	4
Deliverable	4
Preparation.....	5
Customer Contacts.....	5
IBM Contacts.....	5
Define Data Collection Period	5
Softcopy Reports	5
Data Collection	6
Hardware and Software Profile.....	6
IBM MQ Input.....	76
SMF Data	7
Subsystem Data	10
Other MQ Information.....	11
FTP Instructions	12
Sending Data to IBM:.....	13
APPENDIX A	15
PDSE With Program Objects Instructions	15
APPENDIX B.....	16
Sending Data via a Case.....	16
Service User License Agreement	17

By choosing to send data to IBM, you acknowledge that all information contained in your data, including source, object code, binaries, executables, comments, questions, suggestions, or the like, shall be deemed to not contain PI/SPI data as construed by the EU GDPR Regulations.

Overview

The purpose of the *MQ Health Check* is to assess the customer's production implementation and make recommendations on how to run the environment more effectively and/or to more fully exploit the Parallel Sysplex environment for MQ.

Subject Matter Experts (SMEs) from the Washington Systems Center (WSC) organization will assess the customer's production implementation and provide a list of observations and recommendations to improve the environment to better meet the customer's business requirements. The assessment will focus on the following areas of the MQ infrastructure:

- ✓ Installation setup and parameters for MQ
- ✓ MQ maintenance processes
- ✓ Operational and recovery procedures

Scope

Unless otherwise agreed, the scope of this assessment is limited to:

- ✓ One Parallel Sysplex and one data sharing group
- ✓ Previously agreed upon subsystems: MQ for z/OS

This review is initiated remotely and requires the collection of specific data. Key inputs to this review and instructions for sending the data to IBM are documented in the following sections.

All data collected will be sent to IBM and must be available as soon as possible.

Deliverable

IBM's findings and recommendations will be provided in a report and followed up with a conference call (web meeting) with the customer and IBM account team. Each finding and recommendation will provide:

- Description of the component and item analyzed
- Estimated Value and Effort - (High, Medium, Low)
- Observations
- Benefit(s) of implementing the recommendation

The draft report will be delivered and reviewed with the customer's technical team to ensure there are no inaccuracies or misunderstandings. The report will then be finalized and delivered to the customer within a week of this review, with consideration to holidays and conferences.

Preparation

Customer Contacts

Customer involvement throughout the process is important to the success of the health check.

- ☞ Designate a single contact for the data collection process.
- ☞ Designate a customer contact for each area included in the health check. The IBM SMEs may communicate with these contacts for clarification as they analyze the data and prepare the report.

IBM Contacts

The IBM health check SMEs, responsibilities, and emails:

Team Lead: Lyn Elkins

Subject Matter Experts:

MQ for z/OS: Lyn Elkins (elkinsc@us.ibm.com)
Dorothy Quincy (dorothy.quincy@ibm.com)

Define Data Collection Period

SMF/RMF data will be analyzed for selected software components. To facilitate the analysis, customer must identify a peak day that will be input to the analysis.

- SMF 115 (all classes) will be collected for the same 72-hour period.
- All other SMF data will be collected for a 2-hour duration peak on-line interval and a 2-hour peak batch interval for the day identified.

Unless discussed otherwise, the day and time selected should be a normal “peak” business day - no problems or abnormalities occurring during the interval. IBM recommends the SMF and RMF recording interval be set to 10 or 15 minutes.

Softcopy Reports

For mainframe softcopy output/reports, allocate a DOCLIB PDS for each product (MQ for z/OS) using RECFM=FB and LRECL=133 and place all reports and outputs from the display commands into these data sets. Include a member named #README that identifies all the members contained in the data set and how they are used. These data sets will be TERSED and FTPed to IBM.

Schematics and other relevant documents in Microsoft PowerPoint, Word, Excel, Visio, or Adobe Acrobat PDF format should be sent as e-mail attachments to the IBM team lead.

Data Collection

Hardware and Software Profile

The following information should be emailed to the IBM team lead – data can be produced in Microsoft PowerPoint, Word, or any other tool:

- **Completed hardware and software spreadsheet, attached below**



C:\Documents and
Settings\Administrato

- **Schematic of the environment - all Processors and LPARs should be marked clearly.**
 - **Hardware, software, and LPAR layout**
 - **Application mapping to systems**
- **Information regarding recent high impact outage(s)**
 - **Description of the incidents (how each was detected and recovered)**
 - **Problem description and post incident review documentation**
- **Reports and status of previous availability assessments**

IBM MQ Input

Please note these instructions have changed due to updates to IBM MQ for z/OS. Those recent changes have included changing the use of the asterisk (*) in SMF capture and production to make it behave consistently between the ZPRM macro and the START TRACE commands

SMF Data

- ✓ SMF type 115 - MQ Statistics records for at least **72-hours and up to one week**, including the interval(s) selected for the closer study. We ask for the relatively lengthy period due to the basic asynchronous nature of MQ. While the peak MQ processing is typically the same as that for the subsystems processing and producing the messages (CICS, IMS, batch, WAS, etc.), we have seen a number of instances that the peaks for MQ are not aligned, especially when there are client applications connecting directly to z/OS queue managers or queue sharing groups.
 - All classes of the statistics information (01 thru 05) should be captured.
 - To verify the current statistics gathering, issue the DISPLAY TRACE command as shown:
 - +cpf DISPLAY TRACE(STAT)
 - The result should look similar to this:


```

RESPONSE=MPX1
CSQW127I QML1 CURRENT TRACE ACTIVITY IS -
TNO TYPE CLASS DEST USERID RMID
02 STAT 01,02,03,04 SMF * *
END OF TRACE REPORT
                            
```
 - If the trace is not active, or not all the classes are being gathered, the dynamic command to turn on the correct statistics gathering, for all classes:
 - +cpf START TRACE(STAT) DEST(SMF) CLASS(*) AND
 - +cpf START TRACE(STAT) DEST(SMF) CLASS(04) AND (if at MQ V9.3 or higher)
 - +cpf START TRACE(STAT) DEST(SMF) CLASS(05)
 - Note that setting the ZPARM SMFSTAT attribute to an asterisk now only controls classes 1, 2 and 3. Turning on classes 4 and 5 MUST be done via the START TRACE command.
- ✓ SMF type 116 Class 3 & 4 data – MQ Task and Channel Accounting records for the **one to two-hour peak interval(s)** being examined.
 - The collection of SMF116 Class(3) and Class(4) records should be for the same evaluation period as the other subsystems.
 - We are not interested in the Accounting Class(1) records, aka Charge Back records, so we will ask you to remove them to save on transmission time and storage.
 - *There is a JCL sample below that will select SMF type 115 (all subtypes) and SMF 116 subtypes 1-10.*
 - To verify the accounting data gathering, issue the DISPLAY TRACE command as shown:
 - +cpf DISPLAY TRACE(ACCT)

- **The result should look similar to this:**

```

RESPONSE=MPX1
CSQW127I QML1 CURRENT TRACE ACTIVITY IS -
TNO TYPE CLASS DEST USERID RMID
03 ACCTG * SMF * *
END OF TRACE REPORT
CSQ9022I QML1 CSQWVCM1 ' DISPLAY TRACE ' NORMAL COMPLETION
    
```

- **Note that the above example shows that Accounting class 1 is also active, that data is not reviewed for a health check. We would appreciate it if that data is not sent.**
- **SMF116 data collection can be turned on and off dynamically, there is no cycle of the queue manager necessary:**
 - **The commands to dynamically turn on and off SMF 116 collection:**

```

+cpf START TRACE(ACCTG) DEST(SMF) CLASS(3)
+cpf START TRACE(ACCTG) DEST(SMF) CLASS(4)
    
```

```

+cpf STOP TRACE(ACCTG) DEST(SMF) CLASS(3)
+cpf STOP TRACE(ACCTG) DEST(SMF) CLASS(4)
    
```

- **The channel accounting data requires additional set-up to produce the SMF. The steps are:**
 - **Issue the command to turn on gathering of the information for all sender receiver and client connections**
+cpf ALTER QMGR STATCHL(HIGH)
The value of HIGH, MED, or LOW have the same effect
 - **Issue the command to turn on gathering of the information for auto-defined cluster channels if clustering is in use:**
+cpf ALTER QMGR STATACLS(HIGH)
The value of HIGH, MED, or LOW have the same effect
 - **Verify that all channels have STATCHL set to QMGR.**
+cpf DISPLAY CHANNEL(*) STATCHL
Note that even SVRCONN channels will be included in the list, even though they do not have this attribute.

- ✓ For all versions of MQ, please check the following:
 - Verify the STATIME value for the queue managers being evaluated:
 - If this evaluation is for a Queue Sharing Group, please make sure that the STATIME is consistent across the QSG. If they are not, please set them to the same interval before starting the data gathering process.
 - Display the system settings using the DISPLAY SYSTEM command
+cpf DISPLAY SYSTEM
 - If the STATIME value is 0, meaning that the statistics interval is set to the default LPAR value, then this is typically fine. We have seen instances of very high volume LPARs creating varying durations of the SMF data production. If in your evaluation of the data there have been widely varying durations, please set the STATIME to a nonzero value for all queue managers.
 - If the STATIME is set to 30 or above, alter it to be no greater than 15 via the SET SYSTEM command. The sample is setting a 5 minute interval.
+ cpf SET SYSTEM STATIME(05) or if at MQ V9.2.4 or higher
+cpf SET SYSTEM STATIME(05.00)
NOTE: The new interval will not take effect until the current interval has elapsed, so this change needs to be done prior to the start of data collection.
 - The various IBM groups and business partners working with SMF evaluation tools that include machine learning and or AI components are recommending an SMF interval of NO MORE than 5 minutes. We are now recommending the 5 minute or lower interval setting to comply with new tools.
 - Verify the ACCTIME value for queue managers that are at release level 9.2.4 or higher:
 - If ACCTIME is set to -1, it will be the same as the STATIME attribute.
 - If this evaluation is for a Queue Sharing Group, please make sure that the ACCTIME is consistent across the QSG. If it is not, please set the values to the same interval before starting the data gathering process.
 - If the ACCTIME is set to 30 or above, alter it to be no greater than 15 via the SET SYSTEM command. The sample is setting a 5 minute interval.
+ cpf SET SYSTEM STATIME(05.00)
NOTE: The new interval will not take effect until the current interval has elapsed, so this change needs to be done prior to the start of data collection.
 - Verify ACCTQ setting on the queue managers
 - Display the ACCTQ setting for the queue managers in the study
+cpf DISPLAY QMGR ACCTQ
 - If the value is ACCTQ(ON), no further action is needed

- **Otherwise, please use the following command:**
+cpf ALTER QMGR ACCTQ(ON)
- **If clustering is in use, make sure that accounting is turned on for the SYSTEM.CLUSTER.TRANSMIT.QUEUE and any other named cluster transmit queues hosted by the queue manager. To determine whether accounting is enabled for the cluster transmission queues:**
 - **Display the ACCTQ setting for the queue(s)**
+cpf DISPLAY QL(SYSTEM.CLUSTER.TRANSMIT.QUEUE) ACCTQ
 - **If the value of ACCTQ is ACCTQ(OFF), then it should be turned on**
+cpf ALTER QL(SYSTEM.CLUSTER.TRANSMIT.QUEUE) ACCTQ(ON)
Display the ACCTQ value again, to verify the change.
- **For all other high-volume queues, please also verify that the ACCTQ is ON or set to QMGR for the interval under examination.**

Subsystem Data

The following MQ data is to be collected and can be emailed to IBM as .txt files; files can be zipped. Or include these job outputs in the softcopy TERSED dataset to be FTPed to IBM.

- ✓ **MQ topology, showing queue sharing groups, queue manager cluster(s), queue manager, and client connections**
 - **If the queue manager(s) is (are) part of a queue manager cluster:**
 - **Please provide the location of the full repositories.**
 - **Are multiple cluster transmission queues in use (MQ V8 and above)?**
- ✓ **Queue manager and channel initiator JES logs:**
 - **This must include the time period under evaluation**
 - **This must include the queue manager and channel initiator start ups OR contain the output from the following commands for each queue manager:**
 - **+cpf DISPLAY QMGR ALL**
 - **+cpf DISPLAY SYSTEM**
 - **+cpf DISPLAY USAGE**
 - **+cpf DISPLAY LOG**
 - **+cpf DISPLAY ARCHIVE**
 - **+cpf DISPLAY SECURITY**
 - **+cpf DISPLAY CHINIT**
- **If the queue managers are in a Queue Sharing Group, please include:**
 - **+cpf DISPLAY CFSTRUCT(*) ALL**
 - **+cpf DISPLAY CFSTATUS(*)**
 - **+cpf DISPLAY SMDS(*) all**
 - **+cpf DISPLAY SMDSCONN(*) CFSTRUCT(STRUCTURE NAME) for each structure**

Other MQ Information

- ✓ List any open IBM MQ Cases.
- ✓ Definitions of MQ log files. If dual logging is used, both copy 1 and copy 2.

Sample JCL for IBM MQ SMF data selection (uses IFASMFDP) to be sent to WSC. If there are multiple LPARs being sent, please divide the data by LPAR if possible.

```
//SMFDMPA1 EXEC PGM=IFASMFDP
//DUMPIN DD DISP=SHR,DSN=MQHC.CUSTNAME.++LPAR.MQSMF.IN
//DUMPOUT DD DISP=(NEW,CATLG,DELETE),
//      MGMTCLAS=LARGTEMP,DATACLAS=COMT,RECFM=VB,
//      SPACE=(CYL,(5000,2000)),UNIT=SYSDA,
//      DSN=MQHC.CUSTNAME.++LPAR.MQSMF.OUT
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
      INDD(DUMPIN,OPTIONS(DUMP))
      OUTDD(DUMPOUT,TYPE(115,116(1:10)))
/*
//
```

FTP Instructions

By choosing to send data to IBM, you acknowledge that all information contained in your data, including source, object code, binaries, executables, comments, questions, suggestions, or the like, shall be deemed to not contain PI/SPI data as construed by the EU GDPR Regulations.

As of August 31st, 2022, the use of the anonymous id to upload to Testcase and ECuRep servers will no longer be permitted. This means all individuals needing to upload or retrieve data from the servers will need a valid IBMid.

Get an IBMid at <https://www.ibm.com/account/us-en/signup/register.html>

Clients will then use that IBMid to obtain a Support File Transfer ID and password.

Get a Support File Transfer ID at <https://www.secure.ecurep.ibm.com/transfersids/>

The password is displayed only once and never expires. If it is lost, a new Support File Transfer ID and password can be obtained and the old one removed. The Support File Transfer ID and password are used in the JCL to upload documentation.

Sending Data to IBM:

Note: Raw SMF records have a variable block record format. To accurately transfer the SMF data, the customer must convert the datasets to fixed block record format using the TRSMAIN Program.

1. Please compress your dataset using AMATERSE which is shipped with z/OS (MIGLIB)

Please make sure that the input dataset that you terse is on DASD, and not on tape. Tensing a tape dataset produces an output dataset with attributes that the unterse function is unable to process and we will not be able to read the data.

The following are samples of IFASMFDP and AMATERSE jobs supplied for your reference:
(Please don't specify DCB information in these jobs. This will help to avoid TERSE problems.)

```
//SMF1      EXEC PGM=IFASMFDP
//DUMPIN    DD DSN=your-gdg-smf-file(-x),DISP=SHR
//DUMPOUT  DD DSN=FTP.SMFDATA,UNIT=3390,DISP=(,CATLG),
//          VOL=SER=(volser1,volser2),SPACE=(CYL,(3000,200),RLSE)
//SYSPRINT DD SYSOUT=*
//SYSIN     DD *
            INDD(DUMPIN,OPTIONS(DUMP))
            OUTDD(DUMPOUT,TYPE(70:75,78,89,113))
/*
//STEP2     EXEC PGM=AMATERSE,PARM=PACK
//SYSPRINT DD SYSOUT=*
//SYSUT1   DD DISP=SHR,DSN=FTP.SMFDATA
//SYSUT2   DD DISP=(NEW,CATLG,DELETE),VOL=SER=volser,
//          SPACE=(CYL,(2000,100),RLSE), DSN=ZHC.SMFDATA.TERSED
```

Instructions for PDSE Data Sets with Program Objects

Unlike sequential files or most PDSE data sets, you cannot terse PDSE's that include program objects with TRSMAIN. See Appendix B for detailed instructions on sending PDSE data sets with program objects.

2. FTP the compressed dataset to the /systems/toibm/techdata subdirectory at testcase.boulder.ibm.com.

Please prefix the file name to be FTP'ed with the following: ZHC.custname.

The value ZHC allows data on the FTP site to be identified as IBM Z Health Check data. An example of a dataset name would be:

ZHC.ACME.SYSD.D030107.T70.TRS

3. **Connect to the FTP site: testcase.boulder.ibm.com**
(if your installation needs to use an IP address to connect, you should ping testcase.boulder.ibm.com for the current IP address).
4. **Specify BINary transfer mode for the compressed dataset.**
5. **Place the compressed dataset in the /systems/toibm/techdata directory using the put command.**

Sample JCL to FTP a dataset to testcase.boulder.ibm.com

```
//JOBNAME JOB (????,????),'TESTCASE FTP',MSGCLASS=O
//FTPSTEP EXEC PGM=FTP,REGION=4096K,PARM='-r TLS'
//SYSDUMP DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//OUTPUT DD SYSOUT=*
//INPUT DD *
testcase.boulder.ibm.com
<userid>
<password>
cd systems/toibm/techdata
binary
LOCSite FWFriendly
put <input.filename> ZHC.custname....
quit
/*
```

6. **When the transfer is complete, send an email to the IBM WSC team lead which includes:**
 - **File name(s) at the FTP site**
 - **MVS dataset attributes (prior to compression) including LRECL, RECFM BLKSIZE, and Primary and Secondary space**

If there are any questions regarding data collection, please contact or send an e-mail to the IBM team lead.

APPENDIX A

PDSE With Program Objects Instructions

Unlike sequential files or most PDSE data sets, you cannot terse PDSE's that include program objects with TRSMAIN. First, you must XMIT them into a FB, LRECL=80 format.

Prior to tersing the file(s) use this syntax:

XMIT X.Y DS('data.set(member)') OUTDSN(xmitname)

X.Y is a required placeholder for system.userid (but is ignored) and may be coded as -is (the letter X then a dot then the letter Y). The member name is optional. If omitted, the entire PDSE is used as input. This will convert your PDSE into a sequential file. The output file, "xmitname" will be RECFM=FB, LRECL=80, BLKSIZE=3120. You can then terse your OUTDSN data set. We will do the RECEIVE at our end after untersing the file.

APPENDIX B

Sending Data via a Case

In the event customer is unable to send data using the standard FTP process outlined above, data may be submitted via a Case.

Using a Case to submit data to the WSC will cause delays in the processing of the client data by the WSC team.

Instructions for sending data via a case:

- 1. Open a single case against the Product z/OS and specify the Product Area to be zPP.**
- 2. In the description field, please specify "This case is for the WSC. Please notify <focal point> this case was opened."**
- 3. Upload all data to this case as requested in the data collection guide.**

Service User License Agreement

International Business Machines Corporation
Internet Data Exchange Services

Service User License Agreement

This is a valid agreement to which you are consenting to be bound by use of the Service. If you do not agree to all of the terms of this license, do not use the Service. International Business Machines Corporation (IBM) grants you non-exclusive access to the Data Exchange Service via the Internet. IBM makes no representations about the suitability of the Service or about any content or information made accessible by the Service, for any purpose. The Service is provided "as is" without express or implied warranties, including warranties of merchantability and fitness for a particular purpose or non-infringement. The Service is provided gratuitously and, accordingly, IBM shall not be liable for any damages suffered by you or any user of the Service. You also acknowledge that sending information over the public Internet includes the risk that this information data might be kept / modified / read by others and that IBM through the Service does not provide secured data transmission over the Internet. Furthermore, you will defend IBM and its affiliates against any claim that is brought by a third party and is related, directly or indirectly to your use of the Service. While IBM intends to maintain the Service, IBM reserves the right at any time to alter access, features, capabilities and functions of this Service. IBM may terminate the Service at any time without notice to you.

You are responsible as sole data controller for ensuring the compliance with all applicable data protection or similar laws regulating the processing of any personal data which is contained in the information provided by you to IBM. Specifically, Customer understands and acknowledges that IBM may use its global resources, including but not limited to, IBM Affiliates and personnel located in various countries, for the delivery of Services. Customer warrants that none of the Customer Data exported to, or otherwise accessible by, IBM under this Agreement is controlled as a defense article under the U.S. International Traffic in Arms Regulation (ITAR) or under any other country's laws or regulations. Customer further warrants that none of the Customer Data exported to or otherwise accessible by IBM under this Agreement requires an export license or is otherwise restricted from export under applicable export control laws to any IBM Affiliate.

You may not download or upload or otherwise export or re-export images or files from or to systems providing the Service except in full compliance with all applicable export and import laws and regulations. In particular, but without limitation, none of the images or files provided or received by the Service may be downloaded or uploaded or otherwise exported or re-exported: into or received from (or to a national or resident of) Cuba, Iran, North Korea, Sudan, or Syria; or to anyone on (i) the US Treasury Department's list of Specially Designated Nationals or (ii) the US Commerce Department's Table of Denial Orders. By downloading or uploading images or files from or to systems providing the Service, you are agreeing to the foregoing and you are representing and warranting that you are not located in, under control of, or a national or resident of any such country on any such list.

Statement on Confidentiality and EU GDPR Regulations

By choosing to send data to IBM, you acknowledge that all information contained in your data, including source, object code,

binaries, executables, comments, questions, suggestions, or the like, shall be deemed to be non-confidential.

By choosing to send data to IBM, you acknowledge that all information contained in your data, including source, object code,

binaries, executables, comments, questions, suggestions, or the like, shall be deemed to not contain PI/SPI data as construed by the EU GDPR Regulations. IBM shall be free to use, disclose, and distribute the information within IBM without limitation for the purposes of rendering the services contracted for between you and IBM.

Further, IBM shall be free to use in its business activities any ideas, concepts, know-how or techniques contained in such information which are retained in the memories of IBM employees

who have had access to the information in connection with supporting its products and services.

These terms and conditions are in addition to those contained in other agreements in place between you and IBM.