

z/OS
Version 2 Release 4

Problem Management



Note

Before using this information and the product it supports, read the information in [“Notices” on page 347.](#)

This edition applies to Version 2 Release 4 of z/OS (5650-ZOS) and to all subsequent releases and modifications until otherwise indicated in new editions.

Last updated: 2019-07-12

© **Copyright International Business Machines Corporation 2006, 2019.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures.....	xiii
Tables.....	xvii
About this information.....	xix
Who should use this information.....	xix
How to use this information.....	xix
Where to find more information.....	xix
Information updates on the web.....	xx
How to send your comments to IBM.....	xxi
If you have a technical problem.....	xxi
Summary of changes for z/OS Problem Management.....	xxiii
Part 1. Problem management overview.....	1
Chapter 1. Introduction.....	3
Overview of problem resolution.....	3
Steps for diagnosing problems on z/OS.....	4
Gathering diagnosis data.....	5
Problem categories.....	6
Searching problem reporting databases.....	8
Extracting problem symptoms and search arguments.....	8
Formats for symptoms.....	9
Searching for a known problem.....	10
Steps for searching problem reporting databases.....	12
Determining the level of z/OS.....	13
Chapter 2. Common tools for problem determination.....	15
Messages.....	15
BPXMTEXT for z/OS UNIX reason codes.....	16
IPCS.....	16
Logs.....	16
Traces.....	18
Dumps.....	20
IBM Omegamon for z/OS Management Console.....	21
Sending problem documentation to IBM.....	21
IBM documentation.....	21
Chapter 3. Best practices for large stand-alone dump.....	23
Using AutoIPL for stand-alone dumps.....	23
Planning a multivolume stand-alone dump data set.....	23
Creating the multivolume SADUMP.....	24
Defining a dump directory for large stand-alone and SVC dumps.....	24
Preparing the dump for further processing with IPCS COPYDUMP.....	25
Compressing data for faster transmission and analysis.....	26
Transmitting dump data to IBM.....	26
Setting up remote access.....	26

Testing your stand-alone dump operations.....	26
Automating the SADMP process.....	27
Sample JCL for post-processing.....	27
IBM System Test example.....	28
Chapter 4. Using OPERLOG.....	31
Determining hardcopy medium settings.....	31
Setting up OPERLOG.....	31
Steps for setting up OPERLOG.....	31
Part 2. Runtime Diagnostics.....	33
Chapter 5. Runtime Diagnostics.....	35
How Runtime Diagnostics works.....	35
Setting up Runtime Diagnostics.....	36
Runtime Diagnostics output.....	38
Was the Runtime Diagnostics run successful?	40
Specific Runtime Diagnostics event reports.....	41
Syntax and parameters for the MODIFY HZR command.....	53
Chapter 6. Messages that Runtime Diagnostics analyzes.....	57
BPX message analysis.....	57
IEA message analysis.....	59
IEA230E.....	59
IEA231A.....	59
IEA359E	59
IEA360A	59
IEA404A	59
IEA405E.....	59
IEA406I	59
IEA611I.....	59
IEA793A.....	60
IEA799I.....	60
IEA911E.....	60
IEE message analysis.....	60
IEE012A	60
IEE601E.....	60
IEE711I.....	60
IEE766E.....	60
IEE767A	60
IEE786I.....	60
IEE986E.....	60
IOS message analysis.....	61
IOS078I.....	61
IOS079I.....	61
IOS431I.....	61
IOS1078I.....	61
IOS1079I.....	61
IRA message analysis.....	61
IRA100E.....	61
IRA101E	61
IRA103I.....	61
IRA200E	61
IRA201E.....	61
IRA206I.....	61
IRA210E.....	61
IRA211I.....	62

IRA220I.....	62
IRA400E.....	62
IRA401E.....	62
IXC message analysis.....	62
IXC101I.....	62
IXC102A.....	62
IXC105I.....	62
IXC244E.....	62
IXC246E.....	62
IXC255I.....	62
IXC256A.....	63
IXC259I.....	63
IXC267E.....	63
IXC402D.....	63
IXC406I.....	63
IXC409D.....	63
IXC427A.....	63
IXC430E.....	63
IXC431I.....	63
IXC432I.....	63
IXC440E.....	63
IXC446I.....	63
IXC467I.....	64
IXC500I.....	64
IXC512I.....	64
IXC518I.....	64
IXC519E.....	64
IXC522I.....	64
IXC538I.....	64
IXC552I.....	64
IXC553E.....	64
IXC573I.....	64
IXC585E.....	65
IXC615I.....	65
IXC631I.....	65
IXC633I.....	65
IXC635E.....	65
IXC700E.....	65
IXC800I.....	65
IXG message analysis.....	65
IXG051I.....	65
IXG055I.....	65
IXG056I.....	66
IXG057I.....	66
IXG058E.....	66
IXG114A.....	66
IXG121I.....	66
IXG235I.....	66
IXG257I.....	66
IXG261E.....	66
IXG262A.....	66
IXG263E.....	66
IXG266I.....	67
IXG267I.....	67
IXG268I.....	67
IXG270I.....	67
IXG271I.....	67
IXG275I.....	67

IXG285I.....	67
IXG286I.....	67
IXG287E.....	67
IXG288E.....	67
IXG289E.....	67
IXG310I.....	67
IXG311I.....	68
IXG313I.....	68
IXG314I.....	68
IXG318E.....	68
IXG324I.....	68
IXG701I.....	68
IXG733I.....	68
IXL message analysis.....	68
IXL010E.....	68
IXL013I.....	68
IXL030I.....	69
IXL040E.....	69
IXL041E.....	69
IXL044I.....	69
IXL045E.....	69
IXL049E.....	69
IXL158I.....	69
IXL159E.....	69
IXL160E.....	69
IXL162E.....	70
Test messages ignored by Runtime Diagnostics.....	70

Part 3. Predictive Failure Analysis..... 71

Chapter 7. Predictive Failure Analysis overview and installation.....	73
Avoiding soft failures.....	73
Overview of Predictive Failure Analysis.....	73
How PFA works with a typical remote check	74
How PFA interacts with IBM Health Checker for z/OS.....	75
How PFA invokes Runtime Diagnostics.....	75
Migration considerations for PFA.....	76
Running AIRSHREP.sh	76
How PFA uses the ini file.....	76
Installing PFA.....	77
Steps for installing PFA.....	78
Updating the Java path.....	82
Chapter 8. Managing PFA checks.....	83
Understanding how to modify PFA checks.....	83
MODIFY PFA, DISPLAY.....	85
MODIFY PFA, UPDATE.....	88
Configure supervised learning.....	91
Chapter 9. Predictive Failure Analysis checks.....	95
PFA_COMMON_STORAGE_USAGE.....	95
PFA_ENQUEUE_REQUEST_RATE.....	108
PFA_JES_SPOOL_USAGE.....	118
PFA_LOGREC_ARRIVAL_RATE.....	125
PFA_MESSAGE_ARRIVAL_RATE.....	133
PFA_PRIVATE_STORAGE_EXHAUSTION.....	147
PFA_SMF_ARRIVAL_RATE.....	163

Part 4. Diagnosing by problem type..... 177

Chapter 10. Diagnosing an abend.....	179
Overview of an abend.....	179
Steps for diagnosing an abend.....	182
Obtaining the abend and reason code.....	183
Steps for obtaining the abend code.....	183
Identifying the module and component.....	187
Steps for identifying the module and component.....	187
Searching the problem reporting databases.....	190
Steps for searching the problem reporting databases.....	190
Gathering additional problem data for abends.....	191
Steps for gathering additional data for abends.....	191
Chapter 11. Diagnosing a system hang or wait state.....	199
Overview of a hang or wait.....	199
Steps for diagnosing a system hang.....	200
Collecting the problem description.....	202
Steps for collecting the problem description.....	202
Diagnosing a hang or wait during IPL.....	203
Steps for diagnosing a hang or wait during IPL.....	203
Diagnosing an enabled wait state.....	205
Steps for diagnosing an enabled wait state.....	205
Diagnosing a coded disabled wait state.....	207
Steps for diagnosing a coded disabled wait state.....	208
Diagnosing a system partitioned from a sysplex because of status update missing.....	209
Steps for diagnosing a system partitioned because of status update missing.....	210
Searching the problem reporting databases.....	210
Steps for searching the problem reporting databases.....	210
Gathering additional data for hangs and waits.....	212
Steps for gathering messages and logrec for hangs.....	212
Chapter 12. Diagnosing a job or subsystem hang.....	215
Overview of a hang or wait.....	215
Steps for diagnosing a job or subsystem hang.....	215
Gathering additional data for a job or subsystem hang.....	216
Step for gathering additional data.....	217
Determining the status of a hung job or subsystem.....	217
Steps for determining the status of a hung job or subsystem.....	217
Determining if a job is waiting for resources.....	217
Steps for determining if a job is waiting for resources.....	218
Determining address space dispatchability.....	218
Steps for examining address space dispatchability.....	218
Examining the SRB status.....	220
Steps for examining the SRB status.....	220
Examining the TCB status.....	220
Steps for examining the TCB status.....	220
Examining why a job is not running.....	222
Steps for examining why a job is not running.....	223
Chapter 13. Diagnosing a loop.....	225
Overview of a loop.....	225
Steps for diagnosing a loop.....	226
Gathering additional data for a loop.....	227
Steps for gathering loop data.....	227
Analyzing the dump to determine the type of loop.....	228

Step for analyzing the dump for loop type.....	228
Diagnosing a disabled loop.....	229
Steps for diagnosing a disabled loop.....	229
Diagnosing an enabled loop.....	230
Steps for diagnosing an enabled loop.....	230
Diagnosing an excessive spin (spin loop).....	232
Steps for diagnosing an excessive spin.....	233
Analyzing a logrec error record	235
Steps for analyzing a logrec error record.....	235
Searching the problem reporting databases.....	235
Steps for searching the problem reporting databases.....	235
 Chapter 14. Diagnosing an output problem.....	 237
Overview of analyzing output problems.....	237
Steps for diagnosing output problems.....	237
Collecting problem data for an output problem.....	238
Step for collecting problem data.....	239
Analyzing data set allocation for an output problem.....	239
Steps for analyzing data set allocation.....	239
Analyzing the inputs and outputs.....	239
Steps for analyzing the inputs and outputs.....	239
Analyzing installation exits for an output problem.....	240
Steps for analyzing installation exits.....	240
Identifying the program or component.....	241
Steps for identifying the program or component.....	241
Searching the problem reporting databases for an output problem.....	242
Step for searching the problem reporting database.....	242
Gathering additional data for output problems.....	243
Steps for gathering additional information for output problems.....	243
Reporting output problems to IBM.....	244
 Chapter 15. Diagnosing a performance problem.....	 247
Overview of a performance problem.....	247
Steps for diagnosing a performance problem.....	247
Collecting data using commands.....	248
Steps for collecting data using DISPLAY.....	248
Steps for using JES2 commands to collect data.....	250
Checking for resource contention or loop.....	251
Steps for checking resource contention.....	251
Searching the problem reporting database.....	252
Steps for searching the problem reporting databases.....	252
Gathering additional data for performance problems.....	253
Steps for gathering additional information for performance problems.....	253
Analyzing a dump for performance problems.....	253
Steps for collecting and analyzing a dump for performance problems.....	254
Reporting performance problems to IBM.....	254
 Part 5. Diagnosing component-specific problems.....	 257
 Chapter 16. Catalog component operational problem determination.....	 259
Catalog component-specific problems and recovery.....	259
Hang in the Catalog address space or in the user address waiting on a request to the Catalog address space.....	259
Damaged or broken catalogs.....	260
Slow performance in various address spaces due to requests to the catalog address space taking excessive time	261

Chapter 17. Global resource serialization problem determination.....	263
GRS component-specific problems.....	264
Job hang/delay due to GRS resource contention.....	264
Job hang due to latch contention.....	265
GRS related delays across the sysplex.....	266
High CPU usage in GRS.....	267
Wait state X'0A3' for GRS.....	268
Chapter 18. JES2 operational problem determination.....	271
High CPU usage in the JES2 address space.....	271
JES2 contention.....	272
JES2 checkpoint contention.....	274
JES2 checkpoint reconfiguration.....	274
Chapter 19. PDSE operational problem determination.....	277
PDSE specific problems.....	278
ABEND0F4 failures.....	278
MSGIGW038A possible PDSE problems.....	279
PDSE data set corruption.....	281
Failure of the SMSPDSE or SMSPDSE1 address space.....	283
Chapter 20. RRS operational problem determination.....	285
Basic RRS problem determination functions.....	285
Collecting documentation for RRS.....	285
Dumping RRS information.....	285
Important RRS CTRACE information.....	286
RRS recovery options.....	286
RRS warm start.....	286
RRS cold start.....	287
RRS component-specific problems and recovery.....	287
RRS resource contention.....	287
RRS suspended, waiting for signal from system logger.....	288
RRS log stream gap condition.....	289
RRS log stream data loss condition.....	290
RRS high processor usage.....	293
RRS address space hang.....	293
RRS high storage usage.....	293
Resource manager is unable to start with RRS.....	295
Resource manager termination delay.....	295
RRS transaction hang.....	296
RRS severe error on RRS RM.DATA log stream, message ATR250E.....	297
Resolving RRS problems in a sysplex cascaded transaction environment.....	297
Collecting documentation for a sysplex cascaded transaction environment.....	298
Sysplex cascaded transaction hang.....	298
Sysplex cascaded transaction hang messages ATR246I and ATR247E.....	299
Chapter 21. System Data Mover (SDM) operational problem determination.....	301
SDM specific problems.....	301
ANTP0095I Unable to determine PPRC paths.....	301
ANTX5104E RC=0901 (XRC).....	302
ANTX5104E RC=0647 REASON=0053 (XRC).....	303
ANTX5104E RC=0647 REASON=0002 (XRC).....	303
ANTAS00* ASIDs consuming excessive storage below 2GB.....	304
Converting to IR, RC=1017.....	304
Microcode issue impacting concurrent copy.....	305
Chapter 22. System logger operational problem determination.....	307

Basic system logger problem determination functions.....	307
Collecting documentation for system logger.....	307
System logger component-specific problems.....	308
Delay in offload or log stream data set allocation.....	308
System logger high CPU usage.....	309
System logger allocation errors.....	309
 Chapter 23. VSAM component operational problem determination.....	 311
VSAM specific problems.....	311
VSAM Index Trap.....	311
Hang in VSAM record management code.....	313
Loop in VSAM record management code.....	313
Unexpected return codes from VSAM record management.....	314
Issues opening, closing, extending VSAM data sets.....	314
 Chapter 24. VSAM record-level sharing (RLS) operational problem determination.....	 317
VSAM record-level sharing (RLS) specific problems.....	318
ABEND0F4 failures.....	318
HANG/WAIT in RLS/TVS.....	319
SMSVSAM will not start up.....	320
Share Control Datasets not specified.....	320
 Chapter 25. XES and XCF operational problem determination.....	 321
Basic XES and XCF problem determination functions.....	321
XCF and XES component-specific problems.....	321
Sysplex sympathy sickness.....	322
XCF group member stall and signaling sympathy sickness.....	322
Hangs in structure connectors using XES services	324
Unable to bring new or upgraded coupling facility into use	326
High CPU usage in the XCF address space (XCFAS).....	328
 Part 6. Diagnosis reference material.....	 331
Chapter 26. Diagnosis information for z/OS base elements and features.....	333
Chapter 27. Reporting problems to IBM.....	335
Software support service checklist.....	335
Automatic problem reporting.....	337
Invoking IPCS as a background job.....	337
Step for invoking IPCS as a background job.....	337
Chapter 28. Problem diagnostic worksheet.....	339
 Appendix A. Accessibility.....	 343
Accessibility features.....	343
Consult assistive technologies.....	343
Keyboard navigation of the user interface.....	343
Dotted decimal syntax diagrams.....	343
 Notices.....	 347
Terms and conditions for product documentation.....	348
IBM Online Privacy Statement.....	349
Policy for unsupported hardware.....	349
Minimum supported hardware.....	350
Trademarks.....	350

Index..... 351

Figures

1. IBMLink example.....	13
2. No events found message.....	38
3. Runtime Diagnostics output - summary information.....	38
4. Runtime Diagnostics output where home and target systems are different.....	39
5. Status message that reports Runtime Diagnostics success.....	40
6. Status message that reports qualified success with events.....	41
7. Runtime Diagnostics critical WTO message analysis.....	42
8. Runtime Diagnostics HIGHCPU and LOCK report.....	44
9. Runtime Diagnostics LOOP and HIGHCPU report.....	44
10. Runtime Diagnostics ENQ contention report.....	45
11. HZR event for the GRS latch contention event record.....	46
12. HZR event for the z/OS UNIX file system latch contention and waiting threads record.....	47
13. HZR event for JES2	47
14. HZR event for server address spaces.....	48
15. Deadlock event between 3 tasks in an address space across 3 ENQs.....	49
16. Deadlock event between 2 tasks in an address space between an ENQ and a LATCH.....	49
17. Reading correlated events.....	51
18. Correlated ENQ contention found between three jobs across three systems.....	52
19. 2 Correlated events across two address spaces.....	53
20. LPAR view of the PFA components.....	74
21. Runtime Diagnostics report within the PFA message arrival rate check.....	75
22. Runtime Diagnostics report within the SMF arrival rate check.....	76
23. Defining the file system.....	81

24. Common storage usage prediction report.....	103
25. Prediction report for enqueue request rate higher than expected - tracked job.....	112
26. Prediction report for enqueue request rate lower than expected - tracked jobs.....	113
27. No problem report: enqueue request rate higher than expected.....	114
28. Total system exception report: low enqueue request rate.....	114
29. JES spool usage exception report.....	122
30. JES spool usage no problem report.....	122
31. LOGREC arrival rate prediction report.....	129
32. Message arrival rate prediction report: tracked jobs higher than expected	138
33. Message arrival rate prediction report: tracked jobs lower than expected	139
34. Message arrival rate prediction report: other persistent jobs with high arrival rate.....	140
35. Message arrival rate prediction report: other persistent jobs with low arrival rate.....	141
36. Message arrival rate prediction report: non-persistent jobs higher than expected.....	142
37. Message arrival rate prediction report: total system higher than expected.....	142
38. Message arrival rate prediction report: total system lower than expected.....	143
39. Calculating the cause of the private virtual storage exhaustion.....	149
40. Private storage exhaustion prediction no problem report.....	157
41. An example of extended user private area at high capacity but low usage.....	158
42. Private storage exhaustion prediction no problem report (1 of 2).....	159
43. Private storage exhaustion prediction no problem report (2 of 2).....	159
44. SMF arrival rate: tracked jobs higher than expected.....	168
45. SMF arrival rate: tracked jobs lower than expected.....	169
46. SMF arrival rate: other persistent jobs higher than expected.....	170
47. SMF arrival rate: other persistent jobs lower than expected.....	170
48. SMF arrival rate: non-persistent jobs with high counts.....	171

49. SMF arrival rate: no problem.....	171
50. SMF arrival rate; Total system exception report and total system lower than expected.....	172
51. Flowchart for abend analysis.....	183
52. RSMDATA SUMMARY report.....	197
53. Flowchart for system hang analysis.....	201
54. IPL statistics example.....	203
55. IPCS STATUS WORKSHEET report from a stand-alone dump	208
56. Flowchart for job or subsystem hang analysis.....	216
57. Flowchart for diagnosis of a loop.....	227
58. System trace table entry.....	230
59. Flowchart for output problem analysis.....	238
60. Flowchart for performance problem analysis.....	248
61. Output from the DISPLAY SLIP command.....	249
62. JCL for invoking IPCS as a background job.....	337

Tables

1. Problem indicators by type.....	7
2. Obtaining search arguments from SVC dump, stand-alone dump or SYSMDUMP using IPCS commands	9
3. Standardized symptom keyword list.....	10
4. ENQ checking: IBM-supplied address spaces.....	45
5. PFA EXCLUDED_JOBS file format.....	93
6. PFA_COMMON_STORAGE_USAGE check parameters.....	97
7. Multiplier notation used in values for PFA reports.....	103
8. PFA_ENQUEUE_REQUEST_RATE check parameters.....	109
9. Multiplier notation used in values for PFA reports.....	115
10. PFA_JES_SPOOL_USAGE check parameters.....	119
11. Multiplier notation used in values for PFA reports.....	122
12. PFA_LOGREC_ARRIVAL_RATE check parameters.....	126
13. Multiplier notation used in values for PFA reports.....	129
14. PFA_MESSAGE_ARRIVAL_RATE check parameters.....	134
15. Multiplier notation used in values for PFA reports.....	144
16. PFA_PRIVATE_STORAGE_EXHAUSTION check parameters.....	150
17. PFA_SMF_ARRIVAL_RATE check parameters.....	164
18. Multiplier notation used in values for PFA reports.....	172
19. Summary of IPCS dump subcommands by problem.....	196
20. Common wait states that occur during IPL.....	204
21. Common disabled wait states	209
22. Selecting the problem type for STATUS CPU output.....	211

23. Example of using the message prefix to identify the component.....	241
24. Severity levels and examples.....	336
25. What is the impact of your problem?.....	339
26. How is your system configured?.....	339
27. What are the external symptoms.....	340
28. What symptom information did you collect?.....	340
29. Which type of documentation did you obtain?.....	341
30. What recovery actions did you attempt?.....	341

About this information

This information is designed to help you avoid potential problems and diagnose problems on z/OS®, its subsystems, its components, and problems in applications running under the system. Using this information, you can:

- Identify a potential problem
- Identify the problem type
- Determine the failing subsystem, component, job, or application
- Collect the correct data needed to diagnose the problem
- Develop a search argument and use it to search problem reporting databases
- Know the correct problem data to collect before reporting the problem to IBM® or the independent software vendor.

This information can help you determine why a problem occurred and where a problem occurred; it does not describe how to fix program instructions in your own code.

Who should use this information

This information is for anyone who diagnoses software problems that occur while running the operating system. This person is typically a system programmer for the installation. This information is also for application programmers who are testing their programs.

The level of detail at which this information is written assumes that the reader:

- Understands basic system concepts and the use of system services
- Codes in Assembler language, and reads Assembler and linkage editor output
- Codes Job Control Language (JCL) statements for batch jobs and cataloged procedures
- Understands the commonly used diagnostic tasks and aids, such as message logs, dumps, and Interactive Problem Control System (IPCS)

How to use this information

Use the procedures in this information to properly collect problem data, avoid potential problems, and diagnose failures.

If your installation does not want to debug the problem or does not have the source code involved in the problem, use the diagnosis procedures to collect the problem data needed for reporting the problem to IBM or other software vendors. The techniques described in this information are also relevant to non-IBM problems.

If your installation wants to debug the problem and has the source code, use the procedures to collect problem data and debug the problem. If the problem is in IBM code, report the problem to IBM. Where possible, IBM will debug the problem and provide a fix.

Where to find more information

Where necessary, this information references information in other documents, using cross-document links that go directly to the topic in reference using shortened versions of the document title. For complete titles and order numbers of the documents for all products that are part of z/OS, see [*z/OS Information Roadmap*](#).

This information also references diagnosis books for specific components, see [Chapter 26, “Diagnosis information for z/OS base elements and features,”](#) on page 333.

Information updates on the web

For the latest information updates that have been provided in PTF cover letters and Documentation APARs for z/OS, see the [z/OS APAR book \(publibz.boulder.ibm.com/cgi-bin/bookmgr_OS390/Shelves/ZDOCAPAR\)](http://publibz.boulder.ibm.com/cgi-bin/bookmgr_OS390/Shelves/ZDOCAPAR).

This information is updated weekly and lists documentation changes before they are incorporated into z/OS publications.

How to send your comments to IBM

We invite you to submit comments about the z/OS product documentation. Your valuable feedback helps to ensure accurate and high-quality information.

Important: If your comment regards a technical question or problem, see instead [“If you have a technical problem”](#) on page xxi.

Submit your feedback by using the appropriate method for your type of comment or question:

Feedback on z/OS function

If your comment or question is about z/OS itself, submit a request through the [IBM RFE Community](#) (www.ibm.com/developerworks/rfe/).

Feedback on IBM Knowledge Center function

If your comment or question is about the IBM Knowledge Center functionality, for example search capabilities or how to arrange the browser view, send a detailed email to IBM Knowledge Center Support at ibmkc@us.ibm.com.

Feedback on the z/OS product documentation and content

If your comment is about the information that is provided in the z/OS product documentation library, send a detailed email to mhvrcfs@us.ibm.com. We welcome any feedback that you have, including comments on the clarity, accuracy, or completeness of the information.

To help us better process your submission, include the following information:

- Your name, company/university/institution name, and email address
- The following deliverable title and order number: z/OS Problem Management, SC23-6844-40
- The section title of the specific information to which your comment relates
- The text of your comment.

When you send comments to IBM, you grant IBM a nonexclusive authority to use or distribute the comments in any way appropriate without incurring any obligation to you.

IBM or any other organizations use the personal information that you supply to contact you only about the issues that you submit.

If you have a technical problem

If you have a technical problem or question, do not use the feedback methods that are provided for sending documentation comments. Instead, take one or more of the following actions:

- Go to the [IBM Support Portal](http://support.ibm.com) (support.ibm.com).
- Contact your IBM service representative.
- Call IBM technical support.

Summary of changes for z/OS Problem Management

Summary of changes for z/OS V2R4

New and changed:

This document contains terminology, maintenance, and editorial changes. Technical changes are indicated by a vertical line to the left of the change.

- Additional information added to Runtime Diagnostics Chapter 5, “Runtime Diagnostics,” on page 35.
- New Section for Outstanding WTOR message analysis “[Specific Runtime Diagnostics event reports](#)” on page 41.
- Additional options for MODIFY HZR Command “[Syntax and parameters for the MODIFY HZR command](#)” on page 53.
 - CRITICALWTO
 - NOCRITICALWTO
 - CRITICALWTOR
 - NOCRITICALWTOR
 - OTHERWTOR
 - NOOTHERWTOR
- Adjustments made to SCOPE=FOLLOW|LOCAL “[Syntax and parameters for the MODIFY HZR command](#)” on page 53.

Summary of changes for z/OS V2R3

New and changed:

This document contains terminology, maintenance, and editorial changes. Technical changes are indicated by a vertical line to the left of the change.

- Updates to “[RRS log stream data loss condition](#)” on page 290.
- The system now starts Runtime Diagnostics by default at IPL time - see “[Setting up Runtime Diagnostics](#)” on page 36.

For information on migration actions required, see [z/OS Upgrade Workflow](#).
- New examples of Runtime Diagnostics output, see “[Runtime Diagnostics output](#)” on page 38 and “[Specific Runtime Diagnostics event reports](#)” on page 41.
- Runtime Diagnostics detects and displays information about both TCB enabled and TCB recovered loops. See “[LOOP - Loop detection](#)” on page 44.
- New Runtime Diagnostics event reports:
 - DEADLOCK event: Runtime Diagnostics now gathers information about enqueues and latches on the system where it is running in the new DEADLOCK event. It also gathers enqueue and latch information on each system where Runtime Diagnostics invokes itself internally to follow contention blockers. See “[DEADLOCK - Deadlock events](#)” on page 48.
 - CORRELATED event: As Runtime Diagnostics detects events, it analyzes the detected events to determine if any of them are for the same address spaces. When it finds events for the same address space, Runtime Diagnostics correlates the events into a new CORRELATED event. See “[CORRELATED - Correlated events](#)” on page 49.
- New parameters for the MODIFY HZR,ANALYZE command:

- New DEADLOCK and NODEADLOCK values added to the DEBUG parameter, see [“Syntax and parameters for the MODIFY HZR command”](#) on page 53.
- New SCOPE parameter, see [“Syntax and parameters for the MODIFY HZR command”](#) on page 53.
- Runtime Diagnostics now analyzes system logger messages. See [“IXG message analysis”](#) on page 65 for a list of the system logger messages analyzed.

Summary of changes for z/OS V2R2

This document contains terminology, maintenance, and editorial changes. Technical changes are indicated by a vertical line to the left of the change.

New and Changed:

- Runtime Diagnostics now identifies server health issues. As part of RTD's normal diagnostic reporting, it obtains the health status of all address spaces providing health information and include SERVERHEALTH information in system message HZR0200I RUNTIME DIAGNOSTICS RESULT for address spaces with health scores under 100. See [“Specific Runtime Diagnostics event reports”](#) on page 41.
- New SERVERHEALTH and NOSERVERHEALTH values are added to the DEBUG parameter of MODIFY HZR,ANALYZE. See [“Syntax and parameters for the MODIFY HZR command”](#) on page 53.
- Easier set up and installation of PFA including:
 - You can now define an installation-specific PFA directory path using the **PFADIR=** parameter in the `/etc/PFA/ini`. See [“How PFA uses the ini file”](#) on page 76.
 - `/etc/PFA/ini` is now the only ini file now used by PFA. PFA automatically creates `/etc/PFA/ini` if it doesn't exist or copies from an existing check's ini file if one already exists. See [“How PFA uses the ini file”](#) on page 76.
 - With z/OS V2R2, there is no need to run AIRSHREP.sh when installing a new release. PFA will create all directories required by all checks when PFA starts, creating them in the PFA user's home directory or the directory defined on PFADIR= parameter in `/etc/PFA/ini`. See [“Migration considerations for PFA”](#) on page 76.
- ALL and INCLUDED_JOBS parameters added to MODIFY PFA,UPDATE. ALL is the default, giving you both included and excluded jobs. EXCLUDED_JOBS is no longer the default . See [“MODIFY PFA, DISPLAY”](#) on page 85.
- INCLUDED_JOBS information is now displayed in MODIFY PFA,DISPLAY DETAIL output. See [“MODIFY PFA, UPDATE”](#) on page 88.
- Changes to the supervised learning service to support INCLUDED_JOBS files. See [“Configure supervised learning”](#) on page 91.
- PFA check output provides new guidance for existing STDDEV and STDDEVLOW check parameters to avoid check exceptions. See [Chapter 9, “Predictive Failure Analysis checks,”](#) on page 95.
- New FORCEMODEL parameter added to all PFA checks to specify base time for modeling to occur (in 24-hour clock format). This allows installations to set up a preferred, stable time to do modeling, avoiding peak periods. See [Chapter 9, “Predictive Failure Analysis checks,”](#) on page 95.
- You can now have PFA automatically delete check EXC timestamp directories by age using the new EXCDIRDAYS check parameter added to all checks. See [Chapter 9, “Predictive Failure Analysis checks,”](#) on page 95.
- PFA_COMMON_STORAGE_USAGE now allows exceptions to be issued using dynamic severity based on the number of minutes before the exception is predicted to occur. See [“PFA_COMMON_STORAGE_USAGE”](#) on page 95.
- The PFA_JES_SPOOL_USAGE check now tracks address spaces that start more than one hour after IPL. See [“PFA_JES_SPOOL_USAGE”](#) on page 118.
- New PFA check to track potential exhaustion of private virtual storage. See [“PFA_PRIVATE_STORAGE_EXHAUSTION”](#) on page 147.
- New component-specific operational problem determination sections, including:

- [Chapter 17, “Global resource serialization problem determination,” on page 263](#)
- [Chapter 18, “JES2 operational problem determination,” on page 271](#)
- [Chapter 22, “System logger operational problem determination,” on page 307](#)
- [Chapter 25, “XES and XCF operational problem determination,” on page 321](#)

Summary of changes for z/OS Problem Management as updated June 2014

New information about component-specific problem determination, [Part 5, “Diagnosing component-specific problems,” on page 257](#), was missing from the online version of this information. This is now corrected thanks to a thoughtful reader!

Part 1. Problem management overview

Before you begin diagnosing problems, or using PFA or Runtime Diagnostics, it is important to understand the basics and best practices of problem management covered in this section.

Chapter 1. Introduction

If a problem occurs, this information can help you determine what happened, why it happened, and how to find the fix or report the problem to IBM.

This information can also help you avoid problems and soft failures, find specific information about tools and service aids, and locate diagnosis information in the z/OS library. For details, see:

- [Chapter 2, “Common tools for problem determination,” on page 15](#)
- [Part 3, “Predictive Failure Analysis,” on page 71](#)
- [Chapter 26, “Diagnosis information for z/OS base elements and features,” on page 333.](#)

This chapter covers these topics:

- [“Overview of problem resolution” on page 3](#)
- [“Gathering diagnosis data” on page 5](#)
- [“Problem categories” on page 6](#)
- [“Searching problem reporting databases” on page 8](#)
- [“Extracting problem symptoms and search arguments” on page 8](#)
- [“Formats for symptoms” on page 9](#)
- [“Determining the level of z/OS” on page 13](#)

Overview of problem resolution

Typical z/OS problems are classified by the following symptoms:

- **Abend** – an error or abnormal end of a program or job.
- **Wait or Hang** – a coded wait state is loaded or the system or a job appears hung or does not complete.
- **Loop** – the system or program executes infinitely typically using large or higher amounts of processor resource.
- **Incorrupt** – there is incorrect or missing output from a program or job.
- **Performance** – processing is using too much system resource and impacting other parts or users of the system, or processes are taking too long.
- **Message** – an error is reported through a message to the operator or in a log.

When an error occurs, z/OS provides various forms of diagnosis information that contains symptoms. These symptoms can help you with diagnosis and be used in problem source identification (PSI). PSI is the determination of what caused the error based on answers to these questions:

- Why was an abend issued?
- What caused storage to be exhausted?
- Why is the job hung?
- What is causing the job to loop?

This document is designed to help answer these questions and others and make efficient use of your time when diagnosing, searching, and reporting a problem to IBM.

PSI is useful even if the root cause of the problem is not identified. During the process, information and symptoms are gathered to check for a known problem or report a new problem. To ease and expedite problem identification, it is important to provide all the background information available. This includes information about:

- Hardware involved
- System and application software levels
- External symptoms
- Problem impact
- Diagnostic data produced

By providing sufficient information during the first call to IBM or the individual software vendor, you might avoid having to re-create the problem.

The problem diagnostic worksheet contains key information needed to expedite problem resolution. If you are an experienced z/OS system programmer, use the [Chapter 28, “Problem diagnostic worksheet,”](#) on page 339 as a reminder of the important information to gather and report. For example:

- Extract the diagnostic data and symptoms
- Build a search argument
- Search for a known problem
- Gather available diagnostic information
- Report a new problem.

Steps for diagnosing problems on z/OS

To diagnose a problem, follow these steps:

1. When the problem occurs, gather all the available diagnosis information for problem. Use the [Chapter 28, “Problem diagnostic worksheet,”](#) on page 339 as a template for recording data. This might also include your internal problem report describing external symptoms, what might have triggered the problem, and what was done to recover, including the following types of diagnostic information:
 - Dumps
 - Traces
 - Error messages
 - SYS1.LOGREC entries
 - External symptoms
 - Hardware devices
 - Processor models
 - Any other information

These topics can help you collect the data more effectively:

- [“Gathering diagnosis data”](#) on page 5
 - [“Problem categories”](#) on page 6
2. After the problem type is identified, see these diagnosis procedures to identify the source and extract symptoms:
 - [Chapter 10, “Diagnosing an abend,”](#) on page 179
 - [Chapter 11, “Diagnosing a system hang or wait state,”](#) on page 199
 - [Chapter 12, “Diagnosing a job or subsystem hang,”](#) on page 215
 - [Chapter 13, “Diagnosing a loop,”](#) on page 225
 - [Chapter 14, “Diagnosing an output problem,”](#) on page 237
 - [Chapter 15, “Diagnosing a performance problem,”](#) on page 247
 3. While using the procedure, build a search argument from the data collected. See [“Extracting problem symptoms and search arguments”](#) on page 8 for more information.

4. Perform the search. Keep in mind that you might refine your search with more data from the problem. See [“Searching problem reporting databases”](#) on page 8 for more information.
5. If the problem is not found in a database, report it as a new problem providing the documentation and information collected in [Chapter 28, “Problem diagnostic worksheet,”](#) on page 339. See [Chapter 27, “Reporting problems to IBM,”](#) on page 335 for more information.

Tip: Sometimes information is found that is useful in routing the problem to the right place. For example, if the problem is an ABEND0C4 in module XYZ and your search shows multiple hits for ABEND0C4 and XYZ and information about product from another company, contact that company or search that company’s problem reporting databases.

Gathering diagnosis data

It is important to gather the external symptoms and know the impact to the system or sysplex to define the scope of the problem. There can be many symptoms.

For example: Shortly after JOB A started, a dump was produced for an ABEND0C4, the system went into a WAIT064, and was partitioned from the sysplex.

1. Start with diagnosis of the ABEND0C4, which appears to be the trigger, but also understand the cause of the WAIT064 and why the job failure resulted in a system outage. It is important to check for known problems for both symptoms.
2. Next, gather all the diagnosis data available from the time frame the problem or failure occurred.

To identify a system problem, look at the diagnostic data such as:

- External symptoms and the initial problem report. Look for indications, which can include:
 - Messages
 - Job hang
 - System hang
 - High processor usage
 - Incorrect output
 - Dumps produced
 - System slowdown
 - Jobs not starting
- SVC dumps produced as indicated by these messages:

IEA794I

```
IEA794I SVC DUMP HAS CAPTURED: DUMPID=dumpid REQUESTED BY JOB (*MASTER*)
      DUMP TITLE=dump-title
```

IEA911E

```
IEA911E {COMPLETE|PARTIAL} DUMP ON SYS1.DUMPnn

      DUMPid=dumpid REQUESTED BY
      JOB (jobname)
      FOR ASIDS(id,id,...)
      [REMOTE DUMPS
      REQUESTED | REMOTE
      DUMP FOR SYSNAME: sysname]
      INCIDENT TOKEN:incident-token
      [SDRSN =
      vvvvvvvv wwwwww xxxxxxxx
      zzzzzzzz]
      [reason-text]
      [ERRORID = SEQyyyyy]
      CPUzz ASIDasid
      TIMEhh.mm.ss.f]
```

```
[TSOID = tsoid]
[ID = uuuuuuuuuu]
```

IEA611I

```
IEA611I {COMPLETE|PARTIAL} DUMP ON dsname
      text
```

- SYS1.LOGREC data set, which is a repository for information about hardware and system-level software errors.
- Logs from the time frame the problem occurred. This can include SYSLOG, OPERLOG, job log(s), and others.
- Traces associated with the problem.

Tip: After a problem has been isolated to a particular component, query using the TRACE command to see if detailed component traces or GTF was active at the time. For example, if the error is announced by ISGxxxx messages, then check for SYSGRS CTRACE. The message prefix (three or more characters) determines the component owner. See the topic about identifying modules, components, and products in *z/OS MVS Diagnosis: Tools and Service Aids*.

Problem categories

The problem indicator table contains examples of indicators. Some problems might need to be investigated using more than one diagnostic procedure to find the cause. If there are several indicators, look for the earliest problem that caused the other problems.

For example, you find several abends and a wait state, look for the earliest abend code and begin diagnosis there.

Table 1. Problem indicators by type. Where to find problem indicators by type

Problem type	Indicator	System action	System programmer action
Abend See Chapter 10, “Diagnosing an abend,” on page 179.	SVC dump taken and a record of the error (in Logrec)	Produces dump	Review dump to determine if further diagnosis is required
	Message received indicating a system or user abend	Produces record of error	Review response of system message to determine the impact of the abend on the installation.
	An ABEND dump is produced	Continue processing	Review the dump to determine if further diagnosis is required.
	<ul style="list-style-type: none"> • SVC dump produced • Error recorded to SYS1.LOGREC • Error message issued • SYSUDUMP, SYSABEND or CEEDUMP produced 	<p>System actions are the same as the indicators listed previously.</p> <p>Note: The system might also initiate recovery actions. See SYSLOG and component trace to determine what these recovery action were. Some recovery actions can cause data to be lost, requiring the installation to resubmit jobs or transactions.</p>	<ol style="list-style-type: none"> 1. Use IPCS to do problem diagnosis on the dump. 2. Look up abend and reason code recorded for more information about error. 3. Look up the message to gather more information about cause of the error and the system programmer action to correct. 4. Review the dump to do problem diagnosis.
Job hang/wait or loop Chapter 12, “Diagnosing a job or subsystem hang,” on page 215	Job does not end, no further output is produced, and the job can or cannot be CANCEL'ed or FORCE'd	No response	Use the DUMP command to obtain an SVC dump of the hung job. If the DUMP is not successful, consider taking a stand-alone dump.
System hang or wait Chapter 11, “Diagnosing a system hang or wait state,” on page 199	Disabled wait indicated on the HMC and wait state message issued	The system issues a wait state message and loads a disable wait state PSW. The system might load the following into the PSW: X'070E0000 00000000'	Take a stand-alone dump.
	Many jobs are hung in the system	Resource contention	<p>Enter the DISPLAY GRS,C command to check for ENQ resource and latch contention and take a dump of the holder of the resource including SDATA=GRSQ.</p> <p>Note: Use the DISPLAY GRS,ANALYZE command to aid in the discovery of blockers in the system.</p>
	No response to system or subsystem commands entered	No response	Partition the system from the sysplex and take a stand-alone dump.

Table 1. Problem indicators by type. Where to find problem indicators by type (continued)

Problem type	Indicator	System action	System programmer action
Loop Chapter 13, “Diagnosing a loop,” on page 225	High processor resource being consumed locking out other work; Excessive spin detected with IEE178I or ABEND071 issued, or both.	ABEND071 issued in an attempt to stop the looping program	Use an online monitor, such as Resource Measurement Facility RMF or IBM OMEGAMON® z/OS Management Console, to determine whether the problem originates from a high priority job in normal processing or from a problem.
	A job is using a high percentage of central processor storage	Processing degrades	Use an online monitor, such as RMF, to determine whether the problem originates from a high priority job in normal processing or from a problem.
Enabled wait or performance degradation Chapter 15, “Diagnosing a performance problem,” on page 247	System processing slows.	Processing degrades	Use an online monitor, such as RMF, to determine where the problem originates.
	There is a series of WAIT messages followed by a burst of activity	Processing continues	Use an online monitor, such as RMF, to determine where the bottleneck is occurring.
Output problem Chapter 14, “Diagnosing an output problem,” on page 237	Job output is missing or is incorrect.	Processing continues	Use GTF or SLIP to trace input and output.

Searching problem reporting databases

While you are diagnosing a system problem, you will collect data about that problem:

- What was the abend code?
- What did the registers and PSW contain at the time of error?
- What is the failing module or CSECT?
- What components or products were involved with the error?

The answers to these questions are the material for a search argument. A search argument is a list of symptoms for a problem. A search argument is also called a symptom string.

This section contains these topics:

- [“Extracting problem symptoms and search arguments” on page 8](#) describes how to develop a search argument while you are performing diagnosis.
- [“Formats for symptoms” on page 9](#) distinguishes between the types of symptom formats.
- [“Searching for a known problem” on page 10](#) lists the symptoms used in search arguments.
- [“Steps for searching problem reporting databases” on page 12](#) explains the steps to begin your search.

Extracting problem symptoms and search arguments

Obtain search arguments from an SVC dump, SYSMDUMP dump, or stand-alone dump by using IPCS subcommands.

For most problems, use three to five symptoms in the search argument. If the first search produces no matches, remove some symptoms and search again. If the first search produces too many matches, add one or more symptoms and search again. Also, try different symptoms. Searching is an iterative process.

The following are suggestions for selecting symptoms:

- Start with the symptom keyword, for example ABEND0C4, WAIT or LOOP and the module or CSECT name and component ID. Add or remove symptoms from the search argument depending on the number of matches produced.
- Symptoms about data areas are useful for identifying a problem. Use the names of a data area and the incorrect field in the data area as symptoms.
- If searching does not produce a match, remove some symptoms or use different symptoms and try again.

Table 2. Obtaining search arguments from SVC dump, stand-alone dump or SYSMDUMP using IPCS commands

IPCS subcommand	Dump output heading
STATUS FAILDATA	Search Argument Abstract
VERBEXIT DAEDATA	DUMP ANALYSIS AND ELIMINATION (DAE)
VERBEXIT LOGDATA	SEARCH ARGUMENT ABSTRACT
VERBEXIT SYMPTOM	Primary Symptom String

Build a free-format search from the IPCS reports by extracting:

- CSECT name
- Abend code
- Reason code
- Component id

Ensure use of the standardized symptom keyword. For example, system abend code presented as S005C in the IPCS ST FAILDATA report is converted to ABEND05C and reason code PRCS/00000214 is converted to RSN00000214. [Table 3 on page 10](#) contains a list of the standardized symptoms.

Formats for symptoms

Symptom strings or search arguments are presented in several different formats. They include:

- Free-Format symptom: is commonly used to search on the Internet and in IBMLINK for a known problem. The symptoms in the freely formatted string are standardized (see [Table 3 on page 10](#)).

For example:

- ABEND0C4 5752SCXCF IXCS2STB
- A module CSECT name: IEAABCD

- RETAIN symptom string: Use RETAIN symptoms:
 - With a tool such as Info/Management to search the RETAIN database
 - When reporting a problem to IBM
 - In descriptions of problems in APARs and program temporary fixes (PTF)

RETAIN symptoms are also called *structured symptoms* and *failure keywords*. An example of a module CSECT name as a RETAIN symptom is: RIDS/IEAABCD

The table of RETAIN and MVS™ symptoms is in the topic on specifying symptoms in [z/OS MVS Diagnosis: Reference](#).

- MVS symptom, is used by dump analysis and elimination (DAE) when determining if a dump is a duplicate of a previous dump; MVS symptoms are not used for searching problem databases. These symptoms are contained in the DAE data set. An example of a module CSECT name as an MVS symptom

is: CSECT/IEAABCD. For a complete example, see the topic on dump analysis and elimination (DAE) in *z/OS MVS Diagnosis: Tools and Service Aids*.

Searching for a known problem

Use the following search argument of standardized symptoms when performing a search for a known problem in the Technical Support database, IBMLink or when reporting a problem to IBM:

- Always concatenate a number to a word when it modifies that word (for example, SVC13, ABEND0C4)
- Use the word "missing" whenever messages do not appear as expected
- Never abbreviate system commands
- Include = in a search argument with no blanks on either side (for example, DISP=MOD is correct)
- Do not use hyphens in search arguments (for example, SC231234 is the correct way to enter a publication number).

The standardized symptom table describes common symptom keywords to use while doing a search for a known problem or reporting a problem to IBM:

<i>Table 3. Standardized symptom keyword list</i>	
Free-format symptom	Problem data
ABENDxxx	Any system abend except JES3; where xxx is the hexadecimal value of the abend code, always 3 digits including leading zeros
ABENDDMxxx	JES3 abend; where xxx is the hexadecimal value of the abend code, always 3 digits including leading zeros
ABENDUxxxx	User abend; where xxxx is the user abend code
AMODE31	Program running in AMODE 31 (31-bit mode)
AMODE64	Program running in AMODE 64 (64-bit mode)
ARnn	Access register; where nn is the decimal register number without leading zeros
CRnn	Control register; where nn is the decimal register number without leading zeros
D/Txxxx	Device type; where xxxx is the device number
DATASET	Data set
DEQ	Dequeue
DESCCODEnn	WTO descriptor code; where nn is the decimal value of the code, 1-13, without leading zeros
ENQ	Enqueue
ERRNO2n...n	Where n...n is the 4 byte hexadecimal value of the errno2
ERRNOJRn...n	Where n...n is the 4 byte hexadecimal value of the errnojr
ERRNOnnn	Where nnn is the errno in decimal
HANG	Always include this form of the word
I/O	Input Output
KEYn	PSW Key or Storage Key (in hex)
KEYnn	PSW Key or Storage Key (in dec)

Table 3. Standardized symptom keyword list (continued)

Free-format symptom	Problem data
LATCH#nn	Where nnn is decimal latch number without leading zeros (for example: LATCH#2)
LOOP	Always include this form of the word
LPAR	Logical Partition (PR/SM)
LU62	Logical Unit 6.2 protocol
MIH	Missing interrupt handler
MSGxxxx	Any message except JES2 messages; where xxxx is the complete message id of any length
MSGHASPxxx	JES2 messages; note that the '\$' prefix has been removed and xxx is the message id of any length
OVERLAY	Storage overlay; always include this form of the word
PAGEFIX	Page-Fix
PICxx	Program Interrupt Code associated with ABEND0Cx; where xx is the interrupt code, always 2 digits with leading zeros
Rxxx	Release level; where xxx is the product release level
RCnn	Return code; where nn is decimal or hexadecimal and at least two digits
REGnn	General purpose register; where nn is the decimal register number without leading zeros
ROUTCODEnnn	WTO route code; where nnn is the decimal value of the code, 1-128, without leading zeros
RSNxxx	Reason code; where xxx is the hexadecimal reason code of any length
SADMP	Stand-alone dump
SIGxxxx	Where xxx is the name of the signal (for example: SIGTERM)
SIO	Start Input Output
SMFTYPEnnn	SMF type records; where nnn is the decimal value of the record, 0-255, without leading zeros
SUBTYPEnnn	SMF subtype records; where nnn is the decimal value of subtype, 0-255, without leading zeros. Also make sure the SMFTYPEnnn is included
SPnnn	Subpool number; where nnn is the decimal value of subpool, 0-255, with no leading zeros
SVCnnn	Supervisor Call; where nnn is the decimal value of the SVC, 0-255, with no leading zeros
VOLSER	Volume serial
WAIT	Always use this form of the word
WAITxxx	System wait state; where xxx is the hex value of the wait code, always 3 digits including leading zeros
Z/ARCHITECTURE	64-bit mode

Related information:

- See *z/OS MVS IPCS Commands* for the subcommands.
- See *z/OS MVS Diagnosis: Reference* for logrec record formats.
- For formatting of logrec records, see the topic on recording logrec error records in *z/OS MVS Diagnosis: Tools and Service Aids*.

Steps for searching problem reporting databases

About this task

Often the problem has already been reported and fixed. Using the symptom string or search argument extracted, you can do a search of the technical database associated with the product identified.

1. Go to one of these Web sites:
 - [IBM Support Portal \(support.ibm.com\)](http://support.ibm.com).
 - [IBM Techdocs \(www.ibm.com/support/techdocs\)](http://www.ibm.com/support/techdocs).
2. Select the documents you want to search for problem related information. For example, select APARs, FAQs, Technotes, or Flashes.
3. If the Internet is not available at your installation, call the IBM support center and ask them to do the search for you.

Search arguments are used to search problem reporting databases. If the problem being diagnosed was already reported and the symptoms entered into the database, the search will produce a match.

[IBMLINK/ServiceLink \(www.vm.ibm.com/service/zservice.pdf\)](http://www.vm.ibm.com/service/zservice.pdf) contains an overview of IBMLink services.

Example

If your installation has access to IBMLink, an interactive online database program, you can:

- Search for an existing authorized program analysis report (APAR) that is similar to your problem.
- Search for an available program temporary fix (PTF) for the existing APAR.
- Order the PTF if it is available.
- Create an Electronic Technical Response (ETR) problem report to get assistance from a service representative.

The following IBMLink example uses a free-format symptom value of ABEND0C4 IEFJRASP in the search entry.

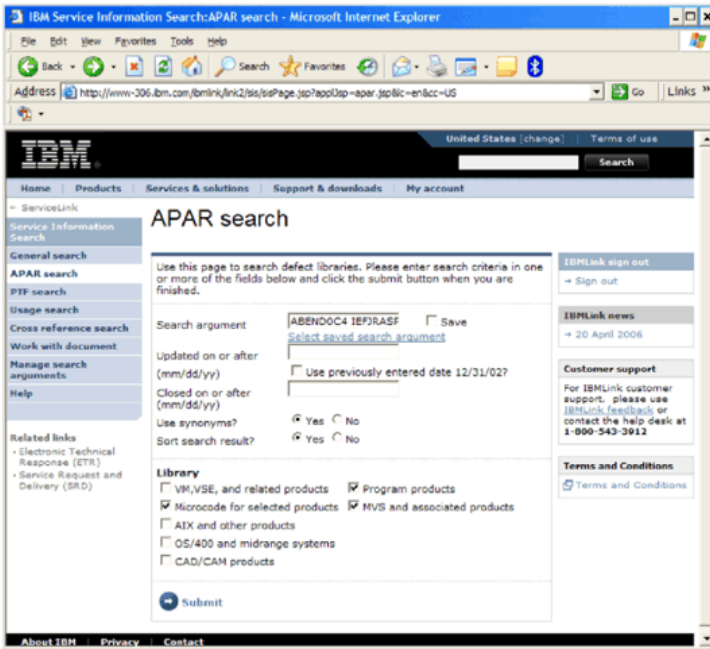


Figure 1. IBMLink example

Determining the level of z/OS

When you report problems to the IBM Support Center, you must provide the name and level of the operating system or systems. If you have communication with your console, you can use the DISPLAY command or you can query the dump using IPCS.

Use the console command DISPLAY IPLINFO or the IPCS command IPLINFO in a dump to obtain the following information:

- The date and time of the IPL
- The release level of the system
- The license value for the system
- The contents of parmlib members IEASYSxx and IEASYMxx
- LOADxx information used for the IPL
- The architecture level of the IPL
- The IODF (input/output definition file) device
- The IPL device and volume serial
- The status of MTL (manual tape library) tape devices.

For example:

```
D IPLINFO
  IEE254I 11.14.07 IPLINFO DISPLAY 350
  SYSTEM IPLED AT 01.15.39 ON 11/01/2007
  RELEASE z/OS 01.09.00 LICENSE = z/OS
  USED LOAD08 IN SYS0.IPLPARM ON ACB2
  ARCHLVL = 2 MTLSHARE = N
  IEASYM LIST = (X6,U6,0L,R8)
  IEASYS LIST = (ST,LN) (0P)
  IODF DEVICE ACB2
  IPL DEVICE 3C2A VOLUME D83EL
```

If you cannot communicate through the console, use IPCS to perform the following steps to determine which system or systems you are using:

1. Use the IPCS subcommand `CBFORMAT` with the communications vector table (CVT) control block to determine the product level.

In the `CBFORMAT CVT` example output, the `PRODN` field indicates an MVS operating system level of `SP7.0.7` and the `PRODI` field indicates the FMID as `HBB7720`.

```
CVT: 00FD48A0
-0028  PRODN.... SP7.0.7  PRODI.... HBB7720  VERID....
-0006  MDL..... 2084    RELNO.... 038
+0000  TCBP.... 00000218  0EF00.... 00FEA3EC  LINK..... 00FD481C
+000C  AUSCB.... 00FD57E8  BUF..... 00000000  XAPG..... 00FE0380
```

2. Determine if the system is running as a uniprocessor or multiprocessor. In the IPCS STATUS WORKSHEET output `PROCESSOR RELATED DATA`, find the MVS Diagnostic Worksheet:

```
MVS Diagnostic Worksheet

Dump Title: W059 SLIP TRAP

CPU Model 2084 Version 00 Serial no. 220CBE Address 00
Date: 09/15/2006      Time: 09:33:32.124515 Local

CSD Available CPU mask: FFC0 Alive CPU mask: FFC00000 00000000
Number of active CPUs: 0000000nn
```

In `0000000nn`, the variable `nn` indicates the number of processors running.

In this output example, there are ten active processors.

```
CSD Available CPU mask: FFC0 Alive CPU mask: FFC00000 00000000
No. of active CPUs: 0000000A
```

- See [z/OS SMP/E User's Guide](#) for using SMP/E.
- See [z/OS MVS IPCS Commands](#) for more information about the `CBFORMAT` subcommand.
- See [z/OS MVS Data Areas](#) in the [z/OS Internet library \(www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary\)](http://www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary) for the format of the SCCB.

Chapter 2. Common tools for problem determination

z/OS contains many tools and service aids to assist you when a problem does occur. The more you know about these tools and service aids, the easier it is for you to diagnose problems and send data to IBM. This chapter provides an overview of the some commonly used tools and where to find more information about each of them.

This chapter covers:

- [“Messages” on page 15](#)
 - [“BPXMTEXT for z/OS UNIX reason codes” on page 16](#)
- [“IPCS” on page 16](#)
- [“Logs” on page 16](#)
- [“Traces” on page 18](#)
- [“Dumps” on page 20](#)
- [“IBM Omegamon for z/OS Management Console” on page 21](#)
- [“Sending problem documentation to IBM” on page 21](#)
- [“IBM documentation” on page 21](#)

Messages

z/OS issues messages from z/OS elements, features, program products, and application programs running on the system. The system issues messages in different ways and to different locations:

- Automated messaging services automatically react to certain messages.
- Most messages are issued through WTO and WTOR macros to one of these locations:
 - Console
 - Hard-copy log
 - Job log
 - SYSOUT data set

Routing codes determine where the messages are displayed or printed. The routing codes for messages issued by the operating system are included with each message.

- Unless specified otherwise, messages, in general, go to the system log (SYSLOG).
- Dump messages are issued through the dumping services routines and are found in:
 - SVC dumps, stand-alone dumps, or SYSMDUMP ABEND dumps formatted by the interactive problem control system (IPCS)
 - Trace data sets formatted by the interactive problem control system (IPCS)
 - ABEND dumps or SNAP dumps produced by the dumping services

In dump or trace data sets formatted by IPCS, the messages are shown on a terminal or in a printed dump.

- Some messages are issued through DFSMS/MVS access methods directly to one of these locations:
 - Output data set
 - Display terminal

BPXMTEXT for z/OS UNIX reason codes

BPXMTEXT is shipped in SYS1.SBPXEXEC and must be in SYSEXEC or SYSPROC to use. It can run from TSO, IPCS, or the z/OS UNIX Shell. You can use BPXMTEXT to interpret errnojr values from zFS (reason code qualifier=EFxx), TCP/IP (reason code qualifier=7xxx), and the C/C++ run-time library (reason code qualifier=Cxxx).

To determine the meaning of reason codes for z/OS UNIX and zSeries File System (zFS), use BPXMTEXT.

From TSO, enter TSO BPXMTEXT xxxxxxxx, where xxxxxxxx is the reason code.

Here's an example:

```
EQQPH35I: EQQPH35I BPX1ATX FAILED WITH RC=0157, RSN=0B1B03AC
```

To find the meaning of RSN=0B1B03AC from TSO, enter:

```
BPXMTEXT 0B1B03AC
```

You get this result:

```
BPXPRES date JRAuthCaller: The caller of this service is authorized.
Authorized callers are not permitted to load or call unauthorized programs.
Action: System key, supervisor state, or APF authorized callers cannot load
or call unauthorized programs.
```

IPCS

Interactive Program Control System (IPCS) is a powerful diagnostic tool in the MVS system that aids the diagnosis of software failures. IPCS provides formatting and analysis support for dumps and traces produced by MVS, other program products, and applications that run on MVS.

Dumps (SVC dump, stand-alone dump, SYSDUMP) and traces (system trace, GTF trace, and CTRACE) need to be formatted before analysis can begin. IPCS provides the tools to format dumps and traces in both an online and batch environment. IPCS provides you with commands that will let you interrogate specific components of the operating system and allow you to review storage locations associated with an individual task or control block. IPCS allows you to quickly review and isolate key information that will assist with your problem determination process.

Using dump and trace data sets and, in some cases, active storage as a source IPCS analyzes information and produces reports that can be viewed at a Time Sharing Option Extensions (TSO/E) terminal or can be printed.

Related information: For complete information about IPCS, see these procedures and documents:

- [“Invoking IPCS as a background job” on page 337](#)
- [z/OS MVS IPCS User's Guide](#)
- [z/OS MVS IPCS Commands](#)
- [z/OS MVS IPCS Customization](#)
- [Using IPCS to format component dump data in z/OS MVS Diagnosis: Reference.](#)

Logs

Do not overlook log data — it should be the first place to look when reviewing a problem. z/OS communicates problems through messages that it writes to logs. Six logs contain the primary sources of problem data:

SYSLOG

The SYSLOG is a SYSOUT data set provided by the job entry subsystem (either JES2 or JES3). SYSOUT data sets are output spool data sets on direct access storage devices (DASD). An installation should print the SYSLOG periodically to check for problems. The SYSLOG consists of:

- All messages issued through WTL macros
- All messages entered by LOG operator commands
- Typically, the hard-copy log
- Any messages routed to the SYSLOG from any system component or program

View SYSLOG through the Spool Display and Search Facility (SDSF) using the LOG option. A small amount of the SYSLOG is also stored in memory and is included when an address space is dumped. This is referred to as master trace (MTRACE) data and can be accessed from IPCS using the VERBX MTRACE command.

This example shows the MVS SYSLOG without time stamps.

```
STC18213 00000090 $HASP100 BPXAS ON STCINRDR
STC18213 00000090 $HASP373 BPXAS STARTED
STC18213 80000010 IEF403I BPXAS - STARTED - TIME=13.36.36 - ASID=001F - SC53
STC16316 00000291 IST663I IPS SRQ REQUEST FROM ISTAPNCP FAILED, SENSE=08570002
    111 00000291 IST664I REAL OLU=USIBMSC.S52TOS48 REAL DLU=USIBMSC.S48TO
    111 00000291 IST889I SID = ED0385CAEEAAF28
    111 00000291 IST264I REQUIRED RESOURCE S48TOS52 NOT ACTIVE
    111 00000291 IST314I END
STC16352 00000291 IST663I IPS SRQ REQUEST FROM ISTAPNCP FAILED, SENSE=087D0001
    883 00000291 IST664I REAL OLU=USIBMSC.S52TOS48 ALIAS DLU=USIBMSC.S48TO
    883 00000291 IST889I SID = ED0385CAEEAAF28
    883 00000291 IST314I END
STC28215 00000291 IST663I IPS SRQ REQUEST TO ISTAPNCP FAILED, SENSE=08570002 86
    864 00000291 IST664I REAL OLU=USIBMSC.S52TOS48 ALIAS DLU=USIBMSC.S48TO
    864 00000291 IST889I SID = ED0385CAEEAAF28
    864 00000291 IST264I REQUIRED RESOURCE S48TOS52 NOT ACTIVE
    864 00000291 IST891I USIBMSC.SC48M GENERATED FAILURE NOTIFICATION
    864 00000291 IST314I END
```

Job log

Messages sent to the job log are intended for the programmer who submitted a job. Specify the system output class for the job log in the MSGCLASS parameter of the JCL JOB statement.

OPERLOG

Operations log (OPERLOG) is an MVS system logger application that records and merges messages about programs and system functions (the hardcopy message set) from each system in a sysplex that activates OPERLOG.

In SDSF the OPERLOG panel displays the merged, sysplex-wide system message log. You can use the parameters of the LOG command to select the OPERLOG panel or the single-system SYSLOG panel. The OPERLOG panel displays the data from a log stream, a collection of log data used by the MVS System Logger to provide the merged, sysplex-wide log.

An individual product has its own log file. These log files might contain data that is valuable when diagnosing a problem. It is particularly important to look for events that precede an actual abend or failure because the problem, in many cases, will have been caused by a previous action.

This example shows the SYSOUT data sets that might be associated with a CICS® address space:

NP	DDNAME	StepName	ProcStep	DSID	Owner
	JESJCLIN			1	CICSTS
	JESMSGLG	JES2		2	CICSTS
	JESJCL	JES2		3	CICSTS
	JESYSMSG	JES2		4	CICSTS
	\$INTTEXT	JES2		5	CICSTS
	CAFF	SCSCPAA1		101	CICSTS
	CINT	SCSCPAA1		103	CICSTS
	DFHCXRF	SCSCPAA1		104	CICSTS
	COUT	SCSCPAA1		105	CICSTS
	CEEMSG	SCSCPAA1		106	CICSTS
	CEEOUT	SCSCPAA1		107	CICSTS
	PLIMSG	SCSCPAA1		108	CICSTS

CRPO	SCSCPAA1	109	CICSTS
MSGUSR	SCSCPAA1	110	CICSTS

The key SYSOUT data sets to review for problem determination data are the JESMSGLG and MSGUSR data sets. The CEEMSG and CEEOUT data sets will contain Language Environment® (LE) problem data typically associated with application problems.

The CICS JESMSGLG SYSOUT data set includes information related to CICS startup and errors related to system problems, not specifically transaction related.

Logrec Error Recording

Log recording (logrec) log stream is an MVS System Logger application that records hardware errors, selected software errors, and symptom records across the sysplex.

Use the records in the logrec data set or the logrec log stream as additional information when a dump is produced. The information in the records can point you in the right direction while supplying you with symptom data about the failure. Use the Environmental Record, Editing, and Printing program (EREP) to:

- Print reports about the system records
- Determine the history of the system
- Learn about a particular error

Logrec data is written to the SYS1.LOGREC data set and is also written to internal storage that is included in a dump. The SYS1.LOGREC data set can be interrogated using the ICFEREP1 program, or if the abend has triggered a dump, the EREP data can be reviewed using the IPCS VERBX LOGDATA command. Generally, the error log entries at the end of the display, if they have an influence on the problem being reviewed, have time stamps that relate to or immediately precede the actual abend; although there is no guarantee the error records will be written in the order they occurred. The error log entries are also written to an internal storage buffer that is included in the dump.

Using a logrec log stream rather than a logrec data set (SYS1.LOGREC, by default) for each system can streamline logrec error recording.

Console log

Console data that the installation chooses to log.

Hardcopy log

The hardcopy log is a record of the system message traffic that the installation chooses to log, such as messages to and from all consoles, commands and replies entered by the operator. In a dump, these messages are in the master trace. With JES3, the hardcopy log is always written to the SYSLOG. With JES2, the hardcopy log is typically written to the SYSLOG, but can also be written to a console printer, if your installation chooses.

Related information:

- [z/OS MVS Planning: Operations](#) contains information about OPERLOG and SYSLOG.
- [Recording logrec error records in z/OS MVS Diagnosis: Tools and Service Aids](#) contains complete information about EREP.
- [Error recording on the logrec data set in z/OS MVS Diagnosis: Reference](#) lists the incidents and the types of records that can be recorded on the logrec data set for each incident.

Traces

System trace

System trace provides an ongoing record of hardware events and software events occurring during system initialization and operation. The system activates system tracing at initialization, which runs continuously, unless your installation has changed the IBM-supplied system tracing. After system initialization, you can use the TRACE command on a console with master authority to customize system tracing. System trace is formatted in a dump using the IPCS SYSTRACE command.

For complete information, see [System trace in z/OS MVS Diagnosis: Tools and Service Aids](#).

Master trace

Master trace maintains a table of all recently issued system messages. This creates a log of external system activity; the other traces log internal system activity. Master trace is activated automatically at system initialization, but you can turn it on or off using the TRACE command. Master Trace is formatted in a dump using the VERBX MTRACE command.

For complete information, see [Master trace in *z/OS MVS Diagnosis: Tools and Service Aids*](#).

Component trace

The component trace service provides a way for z/OS components to collect problem data about events that occur in the component. Each component that uses the component trace service has set up its trace in a way that provides the unique data needed for the component. Component trace is queried and formatted using the IPCS CTRACE command. Trace data is commonly used by the IBM Support Center to:

- Diagnose problems in the component
- Check how the component is running

The IBM support center might direct you to use specific component trace options when you need to re-create a problem to gather more diagnostic data.

For complete information, see [Component trace in *z/OS MVS Diagnosis: Tools and Service Aids*](#).

Transaction trace

Transaction trace enables you to debug problems by tracing the path of a work unit running in a single system or across systems in a sysplex environment. Transaction trace provides a consolidated trace of key events for the execution path of application or transaction type work units running in a multi-system application environment.

The essential task of transaction trace is to aggregate data showing the flow of work between components in the sysplex that combine to service a transaction. Transaction trace traces events such as component entry, exit, exceptions and major events such as COMMIT, and ROLLBACK.

Restriction: Do not use transaction trace as a component tracing facility.

For complete information, see [Transaction trace in *z/OS MVS Diagnosis: Tools and Service Aids*](#).

Generalized trace facility (GTF)

GTF traces system and hardware events similar to those in system trace, but also offers the option of an external writer and to write user defined trace events. GTF trace records can be formatted in a dump or trace data set using the IPCS GTFTRACE command.

For complete information, see [The Generalized Trace Facility \(GTF\) in *z/OS MVS Diagnosis: Tools and Service Aids*](#).

GFS trace (GFS)

GFS trace is a tool that collects information about the use of the GETMAIN, FREEMAIN, or STORAGE macro. You can use GFS trace to analyze the allocation of virtual storage and identify users of large amounts of virtual storage. You must use the generalized trace facility (GTF) to get the GFS trace data output.

For complete information, see [GETMAIN, FREEMAIN, STORAGE \(GFS\) trace in *z/OS MVS Diagnosis: Tools and Service Aids*](#).

Related information:

- [z/OS DFSMSdfp Diagnosis](#)
- [z/OS Infoprint Server Messages and Diagnosis](#)
- For a comprehensive overview of tools and service aids, see the topic on [Selecting tools and service aids in *z/OS MVS Diagnosis: Tools and Service Aids*](#).
- [z/OS Communications Server: IP Configuration Reference](#)

Dumps

SVC dump

An SVC dump provides a representation of the virtual storage for the system when an error occurs. Typically, a system component requests the dump from a recovery routine when an unexpected error occurs. However, an authorized program or the operator can also request an SVC dump when diagnostic dump data is needed to solve a problem. Complete details are found in [SVC dump in z/OS MVS Diagnosis: Tools and Service Aids](#).

Transaction dump

A transaction dump provides a representation of the virtual storage for an address space when an error occurs. Typically, an application requests the dump from a recovery routine when an unexpected error occurs. Complete details are found in [Transaction dump in z/OS MVS Diagnosis: Tools and Service Aids](#)

Abend dump

An ABEND dump shows the virtual storage predominately for an unauthorized program. To produce a dump when one is requested for an error, a JCL DD statement of SYSUDUMP, SYSABEND or SYSDUMP must be included in the input job stream. See [z/OS MVS JCL Reference](#) for more information. An operator can also request an ABEND dump while ending a program, an address space, or canceling a job. There are three types of abend dumps:

- SYSDUMP – Is an unformatted dump that requires IPCS to view and format. Unformatted dumping is sometimes more efficient because only the storage requested is written to the data set, which means the application can capture diagnostic data and be brought back online faster.
- SYSABEND – The largest of the ABEND dumps, is a pre-formatted dump containing a summary dump for the failing program plus many other areas useful for analyzing processing in the failing program.
- SYSUDUMP – The smallest of the ABEND dumps, containing data and areas only about the failing program.

Complete details are found in [Abend dump in z/OS MVS Diagnosis: Tools and Service Aids](#).

SNAP dump

A SNAP dump shows virtual storage areas that a program, while running, requests the system to dump. A SNAP dump, therefore, is written while a program runs, rather than during abnormal end. The program can ask for a dump of as little as a one byte field to as much as all of the storage assigned to the current job step. The program can also ask for some system data in the dump. A SNAP dump is especially useful when testing a program. Complete details are found in [Transaction dump in z/OS MVS Diagnosis: Tools and Service Aids](#)

Stand-Alone dump

The other tools discussed in this chapter are used to collect data for individual work units on a system or a subset of components on a system. A stand-alone dump is used to collect diagnostic information about the entire system. Stand-alone dumps are not produced by z/OS but by either the IPCS SADMP dump data set utility or the AMDSADDD REXX utility. After a stand-alone dump is taken, because the system cannot resume usual processing, the IPL is of the stand-alone dump instead of z/OS.

The stand-alone dump program produces a stand-alone dump of storage that is occupied by either:

- A system that is stopped. For example, your installation has a wait state with no processing, so you must capture a stand-alone dump to diagnosis it.
- A stand-alone dump program that failed. Either the stand-alone dump program dumped itself – a self-dump –, or the operator loaded another stand-alone dump program to dump the failed stand-alone dump program.

The stand-alone dump program and the stand-alone dump together form what is known as the stand-alone dump service aid. The term stand-alone means that the dump is performed separately from usual system operations and does not require the system to be in a condition for normal operation. It

is essential to perform a store status before taking a stand-alone dump because the program gets loaded over storage that might be needed in the dump.

For more information:

- See the topics on [Chapter 3, “Best practices for large stand-alone dump,”](#) on page 23.
- See the complete details in [Stand-alone dump in z/OS MVS Diagnosis: Tools and Service Aids](#) and in [z/OS MVS IPCS User's Guide](#).

IBM Omegamon for z/OS Management Console

The OMEGAMON z/OS Management Console is a monitoring product that includes an interface for z/OS management and is designed to help eliminate, and simplify many z/OS management tasks. The OMEGAMON z/OS Management Console helps deliver real-time, check information provided by the IBM Health Checker for z/OS, and configuration status information for z/OS systems and sysplex resources.

For more information, see the [IBM OMEGAMON for z/OS Management Console library page](#) (publibz.boulder.ibm.com/cgi-bin/bookmgr_OS390/Shelves/C3169480).

Sending problem documentation to IBM

There are two tools available to send problem documentation to IBM. z/OS Problem Documentation Upload Utility (PDUU) and AMATERSE each have unique characteristics suitable for the type of documentation you must send to IBM:

Problem Documentation Upload Utility

The z/OS Problem Documentation Upload Utility (PDUU) is the primary utility for sending large volumes of documentation, such as stand-alone dumps, to the IBM FTP site. The encryption capability ensures that the transfer occurs in a secure manner. For complete details, see the topic about [AMAPDUPL: Problem Documentation Upload Utility in z/OS MVS Diagnosis: Tools and Service Aids](#).

AMATERSE

AMATERSE is useful for compressing (packing) and unpacking relatively small amounts of service data, but is incompatible with PDUU (output and input), and offers no data transfer or encryption capability. Use AMATERSE to compress and extract problem documentation you send to IBM. There are differences between AMATERSE and the former TRSMAN utility. AMATERSE is the preferred over TRSMAN because it is a supported program. For complete details, see the topic about [AMATERSE: Pack and unpack a data set in z/OS MVS Diagnosis: Tools and Service Aids](#).

IBM documentation

There are many types of documentation to aid problem determination. Here are some of the categories:

- [Chapter 26, “Diagnosis information for z/OS base elements and features,”](#) on page 333, which contains diagnosis material by element or feature name.
- [z/OS Internet library \(www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary\)](http://www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary), which contains complete, updated information about all z/OS elements and features.

Many people find it helpful to search on an individual element or feature using the z/OS elements and features search engine available in the [z/OS Internet library \(www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary\)](http://www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary).

- [z/OS Hot Topics \(www.ibm.com/systems/z/os/zos/library/hot-topics/hot-topics.html\)](http://www.ibm.com/systems/z/os/zos/library/hot-topics/hot-topics.html), written by leading z/OS experts, contains hands-on, technical information about z/OS that is not contained in the traditional product libraries.
- The Techdocs Library, which includes:

Dumps

- Flashes that alert you to significant new technical developments and provide guidance on the installation, use and, management of z/OS: [IBM Technical Support Flashes site \(www.ibm.com/support/techdocs/atmastr.nsf/Web/Flashes\)](http://www.ibm.com/support/techdocs/atmastr.nsf/Web/Flashes)
- FAQs to assist you with the installation, use, and management of z/OS: [FAQs - Frequently Asked Questions \(www.ibm.com/support/techdocs/atmastr.nsf/Web/FAQs\)](http://www.ibm.com/support/techdocs/atmastr.nsf/Web/FAQs)
- White papers, presentations, and more: [IBM Techdocs \(www.ibm.com/support/techdocs\)](http://www.ibm.com/support/techdocs).
- Technotes that includes best practices, performance evaluations, recent enhancements and helpful hints and tips: [IBM Technotes \(www.ibm.com/support/techdocs/atmastr.nsf/Web/Technotes\)](http://www.ibm.com/support/techdocs/atmastr.nsf/Web/Technotes)
- [IBM Redbooks \(www.ibm.com/redbooks\)](http://www.ibm.com/redbooks) provide positioning and value guidance, installation and implementation experiences, typical solution scenarios, and step-by-step "how-to" guidelines.

Chapter 3. Best practices for large stand-alone dump

This information describes a set of best practices for optimizing stand-alone dump (SADMP) data capture, optimizing problem analysis time, and ensuring that the stand-alone dump is successful at capturing the necessary information for use by IBM Support. In particular, the following areas:

- [“Using AutoIPL for stand-alone dumps” on page 23](#)
- [“Planning a multivolume stand-alone dump data set” on page 23](#)
- [“Creating the multivolume SADUMP” on page 24](#)
- [“Defining a dump directory for large stand-alone and SVC dumps” on page 24](#)
- [“Preparing the dump for further processing with IPCS COPYDUMP” on page 25](#)
- [“Compressing data for faster transmission and analysis” on page 26](#)
- [“Transmitting dump data to IBM” on page 26](#)
- [“Setting up remote access” on page 26](#)
- [“Testing your stand-alone dump operations” on page 26](#)
- [“Automating the SADMP process” on page 27, which includes \[“Sample JCL for post-processing” on page 27\]\(#\)](#)
- [“IBM System Test example” on page 28](#)

This information replaces existing stand-alone dump best practices information previously documented in:

- "z/OS Best Practices: Large stand-alone dump handling", found by searching for TD103286 at [IBM Technotes \(www.ibm.com/support/techdocs/atsmastr.nsf/Web/Technotes\)](http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/Technotes).
- "Every picture tells a story: Best practices for stand-alone dump", published in the February 2007 issue of z/OS Hot Topics. You can access current issues of z/OS Hot Topics Newsletter at [z/OS Hot Topics \(www.ibm.com/systems/z/os/zos/library/hot-topics/hot-topics.html\)](http://www.ibm.com/systems/z/os/zos/library/hot-topics/hot-topics.html).

Using AutoIPL for stand-alone dumps

You can enable z/OS to automatically trigger a stand alone dump using the automatic IPL (AutoIPL) function. AutoIPL is an automated function, defined in the DIAGxx parmlib member, that the system checks at wait state time. AutoIPL can re-IPL z/OS, or take a SADMP, or take a SADMP and have SADMP re-IPL z/OS when it finishes.

For details including the hard-coded table of wait state and reason codes, the wait state action table (WSAT), which triggers AutoIPL, see:

- The topic about [Using the automatic IPL function in *z/OS MVS Planning: Operations*](#).
- The setting for AutoIPL in the topic about the [DIAGxx parmlib member](#) in [z/OS MVS Initialization and Tuning Reference](#).

Planning a multivolume stand-alone dump data set

Plan a multivolume stand-alone dump data set that places each volume on a separate DASD volume on a separate control unit. You can achieve the best dump performance when the dump is taken to a multivolume DASD stand-alone dump data set. Stand-alone dump exploits multiple, independent volume paths to accelerate data recording. The dump data set is actually spread across all of the specified volumes, not each volume in succession. They should not be treated as multiple single data sets. See the topic on [“Creating the multivolume SADUMP” on page 24](#).

One of the key performance elements of stand-alone dump is the rate at which data writes to DASD. Modern DASD uses cache in the control unit to improve the performance of write operations. The placement of the multivolume stand-alone dump data set across logical subsystems (LSS) needs to avoid filling the bus or cache within the DASD with the data to be written. When the bus or cache is full of data to be written, the speed of the DASD is reduced to the (slower) speed at which the data can be written to the physical media.

There are significant performance improvements when writing the data to a multivolume stand-alone dump data set, or to specific types of DASD. For more information, review [Flash 10143: z/OS Performance: Stand-Alone Dump Performance \(www.ibm.com/support/docview.wss?uid=tss1flash10143\)](http://www.ibm.com/support/docview.wss?uid=tss1flash10143).

When defining your placement of a multivolume stand-alone dump data set, use the following guidelines:

1. Configure each volume on a separate logical subsystem (LSS) to ensure maximum parallel operation. You can achieve the best performance of stand-alone dump when the multivolume data sets have the most separation. That is, separate physical control units and separate channel paths.
2. Configure, if possible, the control units to minimize the occurrence of other activity at the time of the stand-alone dump. For example, DB2® database recovery writing to a local database volume on the same control unit as the stand-alone dump volume can result in slower dump speed and might affect the elapsed time needed to restart an alternative DB2 on an LPAR that is still running.
3. Use FICON-attached DASD volumes, when possible, to yield the best data rates. FICON® channels can deliver much better performance than ESCON channels. However, with sufficient infrastructure and I/O tuning, an ESCON configuration can still deliver high performance.
4. Dedicate more DASD volumes to SADUMP, up to the maximum of 32 volumes, for better overall performance. IPCS offers a SADMP Dump Data Set Utility, available from the IPCS Utility menu. From the data set utility panel, you can specify whether to define, clear, or reallocate your stand-alone dump data set, specify its name and the volume serial numbers for the SADMP “stripes”. This panel will then start the SADMP allocation program to define the data set that you requested. The volume names, device type, and allocated space are also confirmed. While it is not recommended by IBM, stand-alone dump can also be written to a fast tape subsystem. When directing the SADUMP to a tape drive, the dump only uses a single device and does not prompt for another device, so you cannot switch back to using a DASD device for that stand-alone dump.

Creating the multivolume SADUMP

Use the AMDSADDD utility to define a stand-alone dump data set. Specify a volume list (VOLLIST) in AMDSADDD to designate a list of VOLSERS corresponding to each DASD volume making up the data set. You can allocate a multivolume data set using the specified list of volumes. The device number of the first volume is used to specify the data set to stand-alone dump. Again, each volume should be on a different LSS to ensure parallelism when writing the dump. For a sample job that uses AMDSADDD to generate the SADMP data set, see “[IBM System Test example](#)” on page 28. Be sure to catalog your SADMP data set to prevent the possibility of accessing the wrong version of the data set when using *IPCS COPYDUMP* later.

For additional details, see the topic on [Using the AMDSADDD Utility in z/OS MVS Diagnosis: Tools and Service Aids](#).

Defining a dump directory for large stand-alone and SVC dumps

Choosing the right attributes is the key to facilitating post-processing of large stand-alone dumps and large SVC dumps. IPCS is used to consolidate and extract ASIDs from the dump, format, and analyze the dump. IPCS uses a dump directory to maintain information about the layout and content of the dump. The dump directory is a VSAM data set that you can tune for optimal performance.

You can improve IPCS performance by reducing the number of control interval (CI) splits during initialization and analysis of dumps. To do this, specify the RECORDSIZE parameter in BLSCDDIR (shipped in SYS1.SBLSCLIO). The RECORDSIZE parameter in BLSCDDIR is ‘RECORDSIZE (2560 3072)’ and yields

well performing CISIZES for the data portion of the data set. To allow IPCS to be more efficient in its processing, it is recommended that you delete old dump references from the directory periodically (especially stand-alone dumps).

Note:

1. When IBM System Test uses BLSCDDIR, they specify a CI size of 24,576 and a BUFSPACE of X'100000'.
2. You can tune the RECORDSIZE parameter by observing the number of CI splits using standard VSAM data set analysis techniques, such as the LISTCAT command.

Preparing the dump for further processing with IPCS COPYDUMP

Before you begin: If you are using IPCS on z/OS V1R7, apply PTF UA26080 to fix a problem that causes long IPCS initialization time.

After a stand-alone dump is taken to a multivolume data set, it needs to be post-processed before IPCS or other tools can view it. The IPCS COPYDUMP utility reads and processes the multivolume stand-alone dump faster than IEBGENER, and produces a merged dump ordered by ASID. The COPYDUMP utility processes the dump volumes in parallel, allowing the original dump to be read faster. Using COPYDUMP also helps IBM process the dump more quickly because it eliminates the need to reprocess the dump at IBM. Here is how it works:

1. Use IPCS COPYDUMP to produce a merged dump data set from the multivolume stand-alone dump, and a subset of the original stand-alone dump (ASIDs 1-20). [“Sample JCL for post-processing” on page 27](#) contains a sample batch job.
2. Ensure that the output data set specified to COPYDUMP is DFSMS-stripped with at least eight stripes.
3. Catalog the output dump data set to allow IPCS to access it properly.
4. Send the subset dump to IBM using a program like [“Sending problem documentation to IBM” on page 21](#). This is a smaller data set, which takes less time to send through the FTP program to IBM.
5. Keep the merged dump for later use by IBM Support, if necessary.

You can run COPYDUMP to produce a merged version of the entire multivolume stand-alone dump, or to extract a subset of the address spaces contained in the original dump and written to the merged output data set. IBM recommends that you use two COPYDUMP jobs in parallel to produce a full merged dump and a subset merged dump. The subset dump will contain ASIDs 1-20 with the primary system components of the operating system.

IPCS performance is improved when the dump being processed (the COPYDUMP output) is DFSMS-stripped. Placing the dump into a data set with at least eight stripes has shown marked improvement in IPCS response (when IBM is analyzing the problem).

A subset of ASIDs can be extracted from the full stand-alone dump into a separate data set and sent to IBM using COPYDUMP. This has been shown to reduce the data transferred by roughly 30% to 40% compared to transmitting a full stand-alone dump. Use the EASYCOPY parameter on COPYDUMP to automatically create a JOBLIST entry with a predefined list of system address space names. The JOBLIST includes the following job names: ALLOCAS, ANTAS000, ANTMAIN, CATALOG, CONSOLE, DEVMAN, DUMPSRV, IEFSCHAS, IOSAS, IXGLOGR, JESXCF, JES2, JES3, and OMVS.

The syntax of the IPCS COPYDUMP command is:

```
COPYDUMP ODS('OUTPUT DATASET NAME')
EASYCOPY
IDS('INPUT DATASET NAME') NOCONFIRM
```

Again, ensure that the specified output data set name supports DFSMS striping.

Compressing data for faster transmission and analysis

Compress dumps before sending the data to IBM using FTP. Beginning in z/OS V1R13, you can use the z/OS Problem Documentation Upload Utility (PDUU) to compress and send the dump to IBM. PDUU also offers encryption. In z/OS V1R9 through z/OS V1R12 or when sending documentation to independent software vendors, use AMATERSE. Otherwise, use TRSMMAIN. For AMATERSE and TRSMMAIN, you might need to encrypt the resulting data set, so that the data is secure when it arrives at one of the IBM Support staging points (TESTCASE or ECUREP).

- For PDUU, see the topic about [AMAPDUPL: Problem Documentation Upload Utility in z/OS MVS Diagnosis: Tools and Service Aids](#).
- For AMATERSE, see the topic about [AMATERSE: Pack and unpack a data set in z/OS MVS Diagnosis: Tools and Service Aids](#).
- For early releases, download TRSMMAIN from the [IBM Support Portal \(support.ibm.com\)](http://support.ibm.com).
- If you require that the data be encrypted prior to sending it to IBM, place the necessary decryption information in the PMR for IBM Support to use. For more information, see [Encryption Facility for z/OS \(www.ibm.com/systems/z/os/zos/tools/encryption_facility\)](http://www.ibm.com/systems/z/os/zos/tools/encryption_facility).

Transmitting dump data to IBM

The z/OS Problem Documentation Upload Utility (PDUU) is the primary utility for sending large volumes of documentation, such as stand-alone dumps, to the IBM FTP site. The encryption capability ensures that the transfer occurs in a secure manner. For complete details, see the topic about [AMAPDUPL: Problem Documentation Upload Utility in z/OS MVS Diagnosis: Tools and Service Aids](#).

Setting up remote access

Set up remote access through Remote Screen Viewing Support Facility (RSVSF) or Assist On-Site (AOS) to allow IBM to remotely view your dump in time-critical situations. Remote access products allow you to permit IBM Support personnel to immediately log into an IPCS session and view available documentation with no initial data transfer. Choices include:

- [Assist On-site \(www.ibm.com/support/assistonsite\)](http://www.ibm.com/support/assistonsite) available for use worldwide.
- Remote Screen Viewing Support Facility (RSVSF), contact your service representative for details.
- OnTop available for use in Europe, Northeast Europe, and Southwest Europe, contact your service representative for details.

This should always be the first option in a time-critical situation. Rapid viewing of the documentation has the advantage of allowing IBM Support to itemize or customize any additional documentation they may want to send to IBM for the given situation. If documentation is required to be sent through FTP, the analysis of the current documentation can continue while the requested documentation is in transit. In many cases, sending a subset of the stand-alone dump to IBM can prove sufficient for problem resolution as the complete stand-alone dump is not always required for diagnosis of a given problem.

Testing your stand-alone dump operations

It is critical for your Operations staff to train and practice taking a stand-alone dump so that they are familiar with the procedure, and to ensure that all data sets are set up properly before you run into a critical situation. This includes the process and set up for:

- Taking a SADUMP as part of the standard scheduled shutdown of an LPAR
- Using COPYDUMP to obtain the merged dump and the subset dump.

If the dump resides in a DASD dump data set, IBM recommends that you copy the dump to another data set for IPCS processing and clear (re-initialize) the dump data set using the AMDSADDD or IPCS SADMP dump data set utilities. For more information, see the topic on "Using the AMDSADDD utility" in [Using the IPCS Dialog in z/OS MVS IPCS User's Guide](#).

The best practice is to rehearse taking a stand-alone dump during scheduled disaster recovery drills. You can also consider practicing when migrating to a new z/OS release, or when moving to a new processor. If you have a test LPAR that you use to train your operations staff, one part of that training might be to take a stand-alone dump following your local procedures and prepare how to react to stand-alone dump messages.

Automating the SADMP process

The following sample JCL can help you automate several best practices. The results are is two "steps" that can be run as background jobs:

1. Use IPCS COPYDUMP to merge the data and produce a single data set to send to IBM. Because the JCL requires invoking IPCS in a background TSO environment, it is not possible to obtain condition code information from the COPYDUMP "step" to determine whether to invoke the preparation step. That means you must manually examine the results of the COPYDUMP step.
2. Use the "Preparation" job, which will compress the output data set produced by COPYDUMP, encrypt the compressed version, and send the final result through FTP to IBM using PUTDOC.

Tip: Beginning with z/OS V1R10, you can use the AutoIPL function to ensure z/OS takes a stand-alone dump when about to load a disabled wait state. For additional details, see the topic about [Using the automatic IPL function in z/OS MVS Planning: Operations](#).

Sample JCL for post-processing

Post-processing of a stand-alone dump needs to occur in two steps:

1. Run IPCS COPYDUMP to merge the data and produce a single data set to send to IBM. Examine the output from the run step to ensure that the COPYDUMP ran correctly. This JCL is identified as === IPCS COPYDUMP ===.
2. Run the following JCL, which will terse the resulting (striped) dump data set, encrypt the terse version, and send it through FTP to IBM using PUTDOC. This JCL is identified as === TERSE , ENCRYPT and FTP ===.

You can tailor the following JCL to process the data sets to be transmitted to the FTP server. Turn off the LINE NUMBERING in following job.

```
=== IPCS COPYDUMP ===
//IPCSOPYD JOB MSGLEVEL=(2,1),...
// CLASS=V,NOTIFY=&SYSUID.,MSGCLASS=H
//*****
//* IN DD IS USED TO POINT TO THE SOURCE OF INPUT WHICH WOULD BE
//* THE SYS1.SADMP... DATASET
//* OUT DD IS USED TO POINT TO THE OUTPUT OF THE COPYDUMP
//* WHERE P P P P P SHOULD BE THE NUMBER OF CYLINDERS FOR PRIMARY
//* S S S S SHOULD BE THE NUMBER OF CYLINDERS FOR SECONDARY
//* &DATACLAS SHOULD BE THE DATACLAS
//* &MGMTCLAS SHOULD BE THE MGMTCLAS
//* &STORCLAS SHOULD BE THE STORCLAS
//* IPCSDDIR DD DEFINING &SYUID..COPYDUMP.DDIR WITH NON-COMPRESS
//* DATACLAS
//* COPYDUMP SUBCOMMAND TO REQUEST FIRST 20 ADDRESS SPACES
//* IF JES OR CATALOG WERE NOT AMONG THE FIRST 20 ADDRESS SPACES
//* XXX AND YYY SHOULD BE USED FOR THESE TWO SUBSYSTEM ASIDS
//*****
//RUN EXEC PGM=IKJEFT01,REGION=200096K,DYNAMNBR=50
//IPCSPRNT DD SYSOUT=H
//IPCSTOC DD SYSOUT=H
//IPCSPARM DD DISP=SHR,DSN=SYS1.PARMLIB
//SYSTSPRT DD SYSOUT=H
//IN DD DISP=SHR,DSN=SYS1.SADMP.....
//OUT DD DISP=(NEW,CATLOG),DSN=OUTPUT.DATASET.NAME
```

Best practices for large dumps

```
//          SPACE=(CYL,(PPPPP,SSSS),RLSE),DATACLAS=&DATACLAS,
//          MGMTCLAS=&MGMTCLAS,STORCLAS=&STORCLAS
//SYSTSIN DD *
EX 'SYS1.SBLSCLI0(BLSCDDIR)' 'DSN(&SYSUID..COPYDUMP.DDIR) +
RECORDS(900000) DATACLAS(NOCOMP) MGMTCLAS(DMGDEBUG)'
IPCS NOPARM
COPYDUMP IFILE(IN) OFILE(OUT) ASIDLIST(1:20,XXX,YYY) NOCONFIRM
END
/*
```

```
==== TERSE, ENCRYPT and FTP ====
//TRENCFTP JOB CLASS=I,.....
// NOTIFY=&SYSUID.
//JOBLIB DD DISP=SHR,DSN=PDS_WITH_TERSE_ENCRYP_PGM
//TERSE EXEC PGM=TRSMAN,PARM=PACK
//SYSPRINT DD SYSOUT=H
//INFILE DD DISP=SHR,DSN=SOURCE_OF_DUMP
//OUTFILE DD DISP=(NEW,CATLG),
//          DSN=&SYSUID..PMR....TRSD,
//          UNIT=SYSDAL,
//          DATACLAS=COMPRESS,
//          SPACE=(CYL,(PPPPP,SSSS),RLSE)
//DECRYPT EXEC PGM=FTPENCRD,PARM='PASSCODE',COND=(0,NE)
//SYSOUT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//FIN DD DISP=SHR,DSN=*.TERSE.OUTFILE
//FOUT DD DSN=&SYSUID..PMR....TRSENCRP,
//       DCB=(DSORG=PS,RECFM=FB,LRECL=1024),
//       DISP=(NEW,CATLG),UNIT=SYSDAL,
//       DATACLAS=COMPRESS,
//       SPACE=(CYL,(PPPPP,SSSS),RLSE)
//FTPSTEP EXEC PGM=FTP,REGION=5000K,
//          PARM='TESTCASE.BOULDER.IBM.COM (EXIT',COND=(00,NE)
//STEPLIB DD DISP=SHR,DSN=SYS1.TCPIP.SEZALINK
//*SYSDUMP DD SYSOUT=*
//SYSPRINT DD SYSOUT=H
//OUTPUT DD SYSOUT=H
//INPUT DD *
ANONYMOUS
YOUR_EMAIL@
cd mvs/toibm
bin
PUT PMR.....TRSENCRP PMR.....TRS.ENCRP64
quit
/*
```

IBM System Test example

In a recent set of tests performed by the IBM System z® Product Evaluation Test team, a 12-volume configuration was set up to support a stand-alone dump of a 152 GB real memory system:

- Three Enterprise Storage Server® (ESS) subsystems were used:
 - ESS 2105 F20 - 2 FICON 2 GB CHP, 8 GB cache
 - ESS 2105 mod. 800 - 8 FICON 2 GB CHP, 8 GB cache
 - DS6000™ 1750 mod. 511 - 6 FICON 2GB CHP, 1.3 GB cache.
- Four volumes per CP were defined:
 - Each volume on a unique LSS
 - Each volume as 14902 cylinders
 - The DSTYPE=LARGE attribute was used.

Here is an example of the AMDSADDD JCL that used the DASD configuration:

```
//STEP1 EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
EXEC 'SYS1.SBLSCLI0(AMDSADDD)' 'DEFINE (SAD041,SAD042,SAD043,SAD044,SA+
D045,SAD046,SAD047,SAD048,SAD049,SAD050,SAD051,SAD052) (PETDUMP.J80L12.+
```

SADUMP.DSS0.STRIPE) 3390 14902 YES LARGE '

Chapter 4. Using OPERLOG

The operations log (OPERLOG) is a log stream that uses system logger to record and merge communications about programs and system functions from each system in a sysplex. Use OPERLOG as your hardcopy medium when you need a permanent log about operating conditions and maintenance for all systems in a sysplex.

Restriction: OPERLOG is designed to run in a Parallel Sysplex where the log stream is shared among the systems in a CF structure. If you have a basic sysplex, do not attempt to configure OPERLOG. Even if you set up a DASD log stream for OPERLOG on each system, only one system can connect to its OPERLOG log stream at any time because the name SYSPLEX.OPERLOG is hardcoded in the Couple Data Set.

Determining hardcopy medium settings

Use the following step to determine how your installation's hardcopy medium is set up:

- Enter the `DISPLAY CONSOLES ,HARDCOPY` command. The output displays the following information:
 - The hardcopy medium as SYSLOG, OPERLOG, or both
 - The criteria defined by your installation for selecting messages for the hardcopy message set
 - The number of messages waiting to be placed on the hardcopy medium

If you have already defined the log stream as SYSPLEX.OPERLOG in either the data administrative utility or in the IXGINVNT macro, use the `V OPERLOG ,HARDCPY` to assign OPERLOG as the hardcopy medium. You can assign both OPERLOG and SYSLOG by issuing the command separately.

To define a log stream using the system logger services, see [“Setting up OPERLOG” on page 31](#).

Setting up OPERLOG

The following instructions for setting up OPERLOG are a summary of the details found in the [Systems Programmer's Guide to: z/OS System Logger \(www.redbooks.ibm.com/abstracts/sg246898.html\)](#) available from IBM Redbooks ([www.ibm.com/redbooks](#)). For complete details about defining the log stream, see [Setting up IBM Health Checker for z/OS in z/OS MVS Setting Up a Sysplex](#).

Steps for setting up OPERLOG

Before you begin: You must define the logger subsystem.

Perform the following steps to set up OPERLOG.

1. Define the hardcopy device as OPERLOG in the HARDCOPY statement of the CONSOLxx parmlib member. You can change this setting using the `V OPERLOG ,HARDCPY` command. Specify `V OPERLOG ,HARDCOPY ,OFF` to turn off OPERLOG.
2. Define that the correspondent coupling facility structure in the CFRM policy. For example,

```
//OPERLOG JOB CLASS=A,MSGCLASS=A
//POLICY EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
          DATA TYPE(CFRM)
          STRUCTURE NAME(OPERLOG)
                                SIZE(40448)
                                INITSIZE(40448)
                                PREFLIST(FACIL01 ,FACIL02)
```

3. Activate the CFRM policy using the `START ,POLICY ,TYPE=CFRM ,POLNAME=polname` command, or the COUPLExx parmlib member.

4. Define the log stream to the LOGR policy. The following example is for illustrative purposes only; you must follow the recommendations in [z/OS MVS Setting Up a Sysplex](#) and [z/OS MVS Programming: Assembler Services Guide](#):

```
//OPERLOG JOB CLASS=A,MSGCLASS=A
//POLICY EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
DATA TYPE(LOGR)
DEFINE STRUCTURE NAME(OPERLOG)
  LOGSNUM(1)
  MAXBUFSIZE(4092)
  AVGBUFSIZE(512)
DEFINE LOGSTREAM NAME(SYSPLEX.OPERLOG)
  STRUCTNAME(OPERLOG)
  LS_DATACLAS(LOGR4K)
  HLQ(IXGLOGR)
  LS_SIZE(1024)
  LOWOFFLOAD(0)
  HIGHOFFLOAD(80)
  STG_DUPLEX(NO)
  RETPD(30)
  AUTODELETE(No)
```

5. Create the security definitions for RACF® (or for the equivalent security product). In the following example, the SYSPLEX.OPERLOG of the LOGSTRM resource CLASS is given READ permission, which allows everyone to browse the operations log and *userid1* has UPDATE access level, which allows *userid1* to delete records from the log stream. That is, the user ID associated with the job running the IEAMDBLG program. For example:

```
RDEFINE LOGSTRM SYSPLEX.OPERLOG UACC(READ)
PERMIT SYSPLEX.OPERLOG CLASS(LOGSTRM) ID(userid1) ACCESS(UPDATE)
SETROPTS CLASSACT(LOGSTRM)
```

Note: This example is for illustrative purposes only. Follow the guidelines for your installation.

6. After you activate OPERLOG, you must manage the way in which records are handled.

SYS1.SAMPLIB contains a sample program, IEAMDBLG, to read log blocks from the OPERLOG log stream and convert them to SYSLOG format. The program is an example of how to use the services of the MVS system logger to retrieve and delete records from the OPERLOG stream. It reads the records created in a given time span, converts them from Message Data Block (MDB) format to hardcopy log format (HCL or JES2 SYSLOG), and writes the SYSLOG-format records to a file. It also has an option to delete from the stream all the records created prior to a given date. When you use the delete option, a suggestion is to first copy the records on alternate media, and then conditionally delete them on a separate JCL step to ensure that you have a copy of the data before deleting. If you do not run them on two separate conditional steps, deletion occurs simultaneously with copy without any guarantee that the copy process was successful.

For additional details about handling log data, see the topic on [Managing log data: How much? For how long?](#) in [z/OS MVS Setting Up a Sysplex](#).

You know you are done when you enter the DISPLAY CONSOLES ,HARDCOPY command and see OPERLOG.

Part 2. Runtime Diagnostics

Runtime Diagnostics (component name HZR) can perform many of the same tasks you might typically perform when looking for a failure, but it can do the tasks more quickly and without the need for a storage dump.

Chapter 5. Runtime Diagnostics

Runtime Diagnostics is a base component (HZR) of z/OS designed to help you analyze a system with a potential problem or soft failure. Soft failures in particular are often difficult or impossible to detect and can slowly lead to the degradation of the solution that is using z/OS. To understand soft failures, see the definition in [Chapter 7, “Predictive Failure Analysis overview and installation,”](#) on page 73.

Runtime Diagnostics does many of the same tasks you might typically do when you are looking for a failure, such as:

- Reviewing critical messages in the OPERLOG
- Reviewing outstanding write-to-operator-with-reply messages
- Examining address spaces with high processor usage
- Looking for an address space that might be in a loop
- Evaluating local lock conditions
- Analyzing various types of contention that include ENQ, GRS latch contention, and z/OS UNIX file system latch contention
- Looking for deadlocked ENQs and GRS latches, including deadlocks involving both ENQs and GRS latches and deadlocks involving ENQs and GRS latches across multiple systems when Runtime Diagnostics follows contention blockers
- Finding JES2 health exception events
- Finding server address space health exceptions

In many cases, when Runtime Diagnostics finds a critical message, it performs additional analysis based on the job name or other information in the message text. For example, if Runtime Diagnostics identifies an XCF stalled connector message, it processes additional analysis of the identified address space to help narrow down the problem. A key feature of Runtime Diagnostics is its ability to summarize internal processing errors and return the results to you in a message response.

Predictive Failure Analysis (PFA) can also return Runtime Diagnostics report information when activity is absent or unusually low for:

- [“PFA_MESSAGE_ARRIVAL_RATE”](#) on page 133
- [“PFA_ENQUEUE_REQUEST_RATE”](#) on page 108
- [“PFA_SMF_ARRIVAL_RATE”](#) on page 163.

For details, see [“How PFA invokes Runtime Diagnostics”](#) on page 75.

How Runtime Diagnostics works

Runtime Diagnostics is a diagnostic tool to run in your environment when your system experiences events. For example, run Runtime Diagnostics when you experience system degradation to find potential problems. Do **not** run it when your system is operating normally. Do not run it on a schedule or as an attempt to use it as a monitor because the types of events it outputs can be normal operations and not indicative of a problem unless there is some type of system issue. IBM Service might also request that you run Runtime Diagnostics and report its results.

By default, Runtime Diagnostics starts automatically, but you can also manually start it using the following command:

```
S HZR,SUB=MSTR
```

When the Runtime Diagnostics address space is up, it consumes no resources until a modify command is issued to it or Predictive Failure Analysis invokes it internally. Therefore, after the address space is

started, leave it running for these purposes. Once Runtime Diagnostics is up and running, enter the following command to analyze the home system:

```
MODIFY HZR,ANALYZE
```

or

```
F HZR,ANALYZE
```

Runtime Diagnostics searches for certain critical messages and message combinations in the operations log (OPERLOG) stream and attempts to identify other system events with minimal dependencies on other system services. For a list of the messages that Runtime Diagnostics analyzes, see [Chapter 6, “Messages that Runtime Diagnostics analyzes,”](#) on page 57.

By default, Runtime Diagnostics analyzes the home system, which is the system where you issue the F HZR, ANALYZE command. However, Runtime Diagnostics might also invoke itself internally to analyze contention blockers on other systems. For more information about how Runtime Diagnostics follows blockers to other systems, see the SCOPE parameter in [“Syntax and parameters for the MODIFY HZR command”](#) on page 53.

If you want Runtime Diagnostic to analyze a system other than your home system, specify the system name on the SYSNAME parameter of the F, HZR command:

```
F HZR,ANALYZE,SYSNAME=sysname
```

When you analyze a system other than the home system, Runtime Diagnostics includes the following subset of events:

- Critical message analysis
- Outstanding write-to-operator-with-reply message analysis
- ENQ contention checking for SYSTEMS and SYSPLEX-scoped enqueues. When ENQ contention is found, Runtime Diagnostics follows the contention blockers to other systems, even when *SYSNAME=sysname* is specified.
- Deadlock detection for GRS latch and enqueue data available on the home system and GRS latch and enqueue data gathered from other systems in the sysplex where Runtime Diagnostics invokes itself internally to follow contention blockers.

The amount of time Runtime Diagnostics runs when doing analysis depends on the size of the OPERLOG, the amount of data analyzed, and the number of systems analyzed, especially if you specify or default to SCOPE=FOLLOW (see [“Syntax and parameters for the MODIFY HZR command”](#) on page 53) to specify that Runtime Diagnostics invoke itself internally to follow contention blockers on other systems.

When Runtime Diagnostics finds a problem, it displays a multi-line write-to-operator (WTO) message that lists system error events. The message contains a problem description and a suggested next action for your analysis. See [“Specific Runtime Diagnostics event reports”](#) on page 41. Runtime Diagnostics analysis creates a maximum of 999 events, or as many as will fit in the amount of storage reserved for events.

Setting up Runtime Diagnostics

Runtime Diagnostics (HZR) ships in the SYS1.PROCLIB data set and starts automatically at IPL, running as an address space under the master subsystem. The Runtime Diagnostics address space (HZR) remains active until you decide to stop it using the STOP command.

Requirements

The following are the requirements for using Runtime Diagnostics - see [“How Runtime Diagnostics works”](#) on page 35 for further details:

1. You must use OPERLOG and set read permissions for the SYSPLEX.OPERLOG in order to use Runtime Diagnostics critical message analysis. Runtime Diagnostics performs all other types of analysis when OPERLOG is not being used. For details about setting up OPERLOG for your installation, see [Chapter 4, “Using OPERLOG,”](#) on page 31.
2. By default, the system uses the job name HZR for Runtime Diagnostics. You can change the name and specify the changed name on the RDEFINE but you must also explicitly define a job name of HZR on the START command in IEACMD00. For example, if you change the Runtime Diagnostics name from HZR to HZR780, you must also specify the JOBNAME parameter on the start command as follows:

```
S HZR780, SUB=MSTR, JOBNAME=HZR
```

The JOBNAME parameter assigns the name HZR to the started address space as opposed to HZR780.

If you change the name without explicitly specifying HZR on the JOBNAME parameter of the START command in IEACMD00, HZR initialization fails during START processing.

Setting up and starting Runtime Diagnostics:

Use the following instructions to set up and start Runtime Diagnostics:

1. Define a user ID for the HZR started task by following the steps your installation has in place for its security policy. For example, you might use:

```
ADDUSER HZR OWNER(*owner*) DFLTGRP(*dfltgrp*) NOPASSWORD
```

2. Define a profile for HZR in the RACF STARTED class with the user ID from step 1. For example, you might specify:

```
RDEFINE STARTED HZR.HZR STDATA(USER(HZR) GROUP(*dfltgrp*) TRUSTED(YES))
SETROPTS CLASSACT(STARTED) RACLIST(STARTED)
```

Note: On the RDEFINE command, you must specify HZR as the job name. If you change the name, follow the instructions in [Requirement item 2](#).

If you do not define a profile for HZR in the RACF STARTED class, the system issues system message IRR913I when Runtime Diagnostics starts at IPL-time and uses the started procedures table, ICHRIN03.

3. By default, the system places Runtime Diagnostics in the workload management SYSSTC service class. It is recommended that you either leave it in this service class or put it in another importance 1 service class with a high velocity goal. See [z/OS MVS Planning: Workload Management](#) for complete information.
4. Enter the MODIFY command to invoke Runtime Diagnostics. For example:

```
MODIFY HZR,ANALYZE
```

The system returns either message HZR0200I or HZR0201I. The output shows SUMMARY: SUCCESS as in [Figure 5](#) on page 40.

Customizing Runtime Diagnostics output

- By default, the Runtime Diagnostics WTO message output and a copy of the Runtime Diagnostics message automatically goes to DDNAME HZROUT. Specify BLKSIZE=0 for the system or choose an optimum block size for DASD. If you set the block size to a number other than zero, ensure that it is an even multiple of 121.

If you want only the WTO message output, omit the HZROUT DD statement from the started task procedure. In this case, the system issues a message warning that DDNAME HZROUT is missing and no output is written to HZROUT.

- By default the HZR started task is defined with output statement HZROUT specified as DD DUMMY, but you can customize HZR to direct HZROUT to a sequential data set.

- To view the HZROUT data set without stopping HZR, you must specify DISP=SHR. For example:

```
LRECL=121, BLKSIZE=0, RECFM=FB, DISP=SHR
```

- Runtime Diagnostics displays a maximum of 20 events as WTO messages. If more than 20 events are found, it outputs a message to indicate that only 20 events are displayed and to look for additional events in the hardcopy log. The event report heading indicates when Runtime Diagnostics finds more than 20 events:

```
HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2016/02/02 11:59:34
INTERVAL: 60 MINUTES
EVENTS FOUND: 24
PRIORITIES: HIGH:24 MED:0 LOW:0
TYPES: ENQ:21 SERVERHEALTH:3
DISPLAYING 20 EVENTS. ALL EVENTS SENT TO THE HARDCOPY LOG AND
OPTIONALLY TO THE DATA SET SPECIFIED AS HZROUT.
```

Runtime Diagnostics output

When you run Runtime Diagnostics (F HZR, ANALYZE), the output is displayed in a multi-line write-to-operator (WTO) message.

If Runtime Diagnostics does not find symptom events, it returns message HZR0200I listing summary information, as shown in [Figure 2 on page 38](#).

```
HZR0201I SUCCESS. TIME (2009/06/09 - 10:25:01). NO RUNTIME DIAGNOSTICS EVENTS WERE FOUND FOR SYSTEM: SY1
```

Figure 2. No events found message

If Runtime Diagnostics does find symptom events, it returns message HZR0200I containing summary information followed by symptom event information. The message contains a problem description and a suggested next action for your analysis, as shown in [Figure 3 on page 38](#):

```
F HZR, ANALYZE
HZR0200I RUNTIME DIAGNOSTICS RESULT 048
SUMMARY: SUCCESS
REQ: 6 TARGET SYSTEM: SY1      HOME: SY1      2016/01/28 15:21:28
INTERVAL: 60 MINUTES
EVENTS FOUND: 2
PRIORITIES: HIGH:2 MED:0 LOW:0
TYPES: ENQ:1 CORRELATED:1
CORRELATED TYPES: ENQ:3
PROCESSING BYPASSED FOR SYSTEM SY3:
ALL TYPES..... FOLLOW BLOCKERS BYPASSED. NO HZR XCF SERVER.
PROCESSING FAILURES ON SYSTEM SY1:
ENQ..... ISGQUERY QSCAN ERROR.....RC=00000004 RS=00000000
```

Figure 3. Runtime Diagnostics output - summary information

The fields shown in [Figure 3 on page 38](#) are explained below:

SUMMARY:

Runtime Diagnostics reports the success of an analysis instance (F HZR, ANALYZE) in the SUMMARY: field. See [“Was the Runtime Diagnostics run successful?” on page 40](#) for more information.

REQ:

Displays the number of times that Runtime Diagnostics was invoked (with the F HZR, ANALYZE command or internal requests) since it started. This number wraps at 99,999.

TARGET SYSTEM: nnn HOME: nnn

The TARGET SYSTEM: field displays the system for which the Runtime Diagnostics command was run.

HOME: displays the system where the Runtime Diagnostics `hzr,analyze` command was run.

[Figure 3 on page 38](#) shows an instance when the test and target systems are the same.

When the target and home systems are different, output might look as follows:

```
HZR0200I RUNTIME DIAGNOSTICS RESULT 593
SUMMARY: SUCCESS - NO EVENTS FOUND
REQ: 1 TARGET SYSTEM: SY3 HOME: SY1 2015/06/01 13:24:49
INTERVAL: 60 MINUTES
EVENTS FOUND: 0
PROCESSING BYPASSED:
OMVS.....SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM.
LATCHES.....SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM.
LOOP.....SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM.
HIGHCPU.....SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM.
LOCK.....SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM.
SERVERHEALTH SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM
JES2.....SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM.
```

Figure 4. Runtime Diagnostics output where home and target systems are different

If the target system is not the same as the home system, Runtime Diagnostics analysis is limited to the following subset of events:

- Critical message analysis
- Outstanding write-to-operator-with-reply message analysis
- ENQ contention checking for SYSTEMS and SYSPLEX-scoped enqueues. When ENQ contention is found, Runtime Diagnostics follows the contention blockers to other systems, even when `SYSNAME=sysname` is specified.
- Deadlock detection for GRS latch and enqueue data available on the home system and GRS latch and enqueue data gathered from other systems in the sysplex where Runtime Diagnostics invokes itself internally to follow contention blockers.

In this figure, the lines below the `PROCESSING BYPASSED` field shows how Runtime Diagnostics bypassed processing for certain events because the target and home systems are not the same.

INTERVAL:

Displays how far back Runtime Diagnostics looks in OPERLOG for messages. Runtime Diagnostics will look back into OPERLOG for up to 60 minutes back, regardless of how long OPERLOG has been collecting messages.

EVENTS FOUND:

Displays the number of events that Runtime Diagnostics found in this run.

PRIORITIES:

Displays the number of events found for each priority, HIGH, MED, and LOW

TYPES:

Displays the number and types of events found. In [Figure 3 on page 38](#), one ENQ event and one CORRELATED event were found. For examples of all the event types, see [“Specific Runtime Diagnostics event reports” on page 41](#)

CORRELATED TYPES:

Displays the number and types of correlated events found. Correlated events are events that are for the same address space or events that are related due to following blockers. In [Figure 3 on page 38](#), one CORRELATED event was found and it contains three ENQ events. See [“CORRELATED - Correlated events” on page 49](#) for complete information .

PROCESSING BYPASSED

Information about events that Runtime Diagnostics did not analyze (because the target and home systems are different, for example), are displayed in the `PROCESSING BYPASSED` section. Runtime Diagnostics displays a reason why the event was not analyzed for each event type. [Figure 3 on page](#)

38 shows that all event types were bypassed by Runtime Diagnostics on SY3 because there was no XCF server for address space HZR on SY3.

If events are bypassed on more than one system, Runtime Diagnostics outputs multiple PROCESSING BYPASSED sections, one for each system.

PROCESSING FAILURES

Information about the unsuccessful processing is displayed under the PROCESSING FAILURES field. The system specified in this line is the system where Runtime Diagnostics was running when the failure occurred. [Figure 3 on page 38](#) shows one processing failure on system SY1.

If failures occur on more than one system Runtime Diagnostics outputs multiple PROCESSING FAILURES sections, one for each system.

For the complete list of messages that Runtime Diagnostics issues, see the topic on [HZR messages](#) in *z/OS MVS System Messages, Vol 6 (GOS-IEA)*.

See the following topics on Runtime Diagnostics output:

- [“Was the Runtime Diagnostics run successful?” on page 40](#)
- [“Specific Runtime Diagnostics event reports” on page 41.](#)

Was the Runtime Diagnostics run successful?

Runtime Diagnostics reports the success of an analysis instance (F HZR, ANALYZE) in the SUMMARY: field and reports the number of events found in the EVENTS FOUND: section of the output. When Runtime Diagnostics finds events, the report is returned in message HZR0200I.

The following list shows examples of output:

- **SUMMARY: SUCCESS** and **EVENTS FOUND:** Runtime Diagnostics analysis ran successfully and found events:

```
f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2015/06/01 13:24:49
INTERVAL: 60 MINUTES
EVENTS FOUND: 5
PRIORITIES: HIGH:5 MED:0 LOW:0
TYPES: CF:4 HIGHCPU:1
-----
```

Figure 5. Status message that reports Runtime Diagnostics success

- **SUMMARY: QUALIFIED SUCCESS** with **EVENTS FOUND:** Runtime Diagnostics was unable to complete processing for one or more events, but reports on others:

```

F HZR,ANALYZE
HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: QUALIFIED SUCCESS - SOME PROCESSING
FAILED
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2015/06/01 13:24:49
INTERVAL: 60 MINUTES
EVENTS FOUND: 1
PRIORITIES: HIGH:1 MED:0 LOW:0
TYPES:
HIGHCPU:1

PROCESSING FAILURES ON SYSTEM SY1:
ENQ..... ISGQUERY QSCAN ERROR.....RC=00000004 RS=00000000
PROCESSING BYPASSED FOR SYSTEM SY1:
OPERLOG..... OPERLOG IS NOT ACTIVE.
-----
EVENT 1: HIGH:HIGHCPU          SYSTEM: SY1      2015/06/02 08:23:01
ASID CPU RATE: 100% ASID: 0029 JOBNAME: RTDLP
STEPNAME: STEP1 PROCSTEP:          JOBID: JOB00041
USERID: WELLIE0 JOBSTART: 2015/06/02 08:17:50
ERROR : ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MIGHT BE LOOPING.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----

```

Figure 6. Status message that reports qualified success with events

Figure 6 on page 41 shows the following:

- Information about the unsuccessful processing is displayed under the PROCESSING FAILURES field. The system specified in this line is the system where Runtime Diagnostics was running when the failure occurred. Figure 3 on page 38 shows one processing failure on system SY1. If failures occur on more than one system Runtime Diagnostics outputs multiple PROCESSING FAILURES sections, one for each system.
- Although Runtime Diagnostics was unable to connect to OPERLOG, it continues analysis for other types of events. In this example, Runtime Diagnostics found a EVENT 1: HIGH:HIGHCPU symptom event.

Note that PROCESSING BYPASSED lines displaying information that Runtime Diagnostics did not analyze can be displayed for either SUMMARY: SUCCESS or SUMMARY: QUALIFIED SUCCESS outcomes.

See “[Specific Runtime Diagnostics event reports](#)” on page 41 for more report examples.

Specific Runtime Diagnostics event reports

This section covers the types of events Runtime Diagnostics detects and displays in HZS0200I reports.

Critical WTO message analysis

Requirement: In order to do critical WTO message analysis, Runtime Diagnostics requires that OPERLOG be in use and that read permissions be set for the SYSPLEX.OPERLOG. If OPERLOG is not in use, Runtime Diagnostics performs all other types of analysis. For details about setting up OPERLOG for your installation, see [Chapter 4, “Using OPERLOG,” on page 31](#).

Runtime Diagnostics reads through the last hour of OPERLOG looking for critical WTO messages. If any are found, Runtime Diagnostics lists the critical WTO message as an error event. For a subset of critical WTO messages, Runtime Diagnostics does additional analysis based on the message identifier and the content of the message. If less than one hour of message content is available in OPERLOG, message analysis is done for the messages available.

Runtime Diagnostics searches OPERLOG for messages issued on the system it is analyzing, which means the home system if that is the system being analyzed, or the target system, if the target system is not the same as the home system. When Runtime Diagnostics is following blockers, which means when an event is issued that has a blocker that does not result in an individual ASID event (see [“CORRELATED - Correlated events” on page 49](#)), Runtime Diagnostics searches the OPERLOG for critical WTO messages issued on the system where the blocker is running as well as the system Runtime Diagnostics is analyzing. Events are created for some messages regardless of which system issued the message if the message is considered critical if found on any system in the sysplex.

Some critical WTO messages for which additional analysis is done include:

IXC101I, IXC105I, IXC418I

Runtime Diagnostics compares messages IXC105I and IXC418I to determine whether they were issued for the same system as the IXC101I. If so, Runtime Diagnostics lists the activity in the IXC101I event. For example:

```
HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2015/06/02 08:28:40
INTERVAL: 60 MINUTES
EVENTS FOUND: 2
PRIORITIES: HIGH:2 MED:0 LOW:0
TYPES: XCF:2
-----
EVENT 1: HIGH:XCF              SYSTEM: SY1      2015/06/02 08:28:40
IXC101I SYSPLEX PARTITIONING IN PROGRESS FOR SY3
ERROR : MESSAGE IXC105I WAS ALSO ISSUED FOR SAME SYSNAME
ACTION: LOOK FOR AND CORRECT ANY PROBLEMS WITH THE ETR CLOCK,
ACTION: SIGNALING PATHS, OR COUPLE DATA SET.
-----
EVENT 2: HIGH:XCF              SYSTEM: SY1      2015/06/02 08:28:40
IXC105I SYSPLEX PARTITIONING HAS COMPLETED FOR SY3
ERROR : XCF REMOVED A SYSTEM FROM THE SYSPLEX.
ACTION: LOOK FOR AND CORRECT ANY PROBLEMS WITH THE ETR CLOCK,
ACTION: SIGNALING PATHS, OR COUPLE DATA SET.
-----
```

Figure 7. Runtime Diagnostics critical WTO message analysis

IXL013I

Runtime Diagnostics compares the structure name, job name, ASID, and connector name for multiple IXL013I messages. Runtime Diagnostics lists only the last message that contains a match of all four fields.

IXL030I

Runtime Diagnostics compares the structure name and the connector name for multiple IXL030I messages. Runtime Diagnostics lists only the last message that contains a match of both fields.

IXC431I

Runtime Diagnostics compares the stalled identifiers for multiple IXC431I messages. It lists the last IXC431I message that contains the stalled identifier and analyzes the job name to determine whether it is the waiter for any ENQ contention. The stalled identifiers display as ID: *s#.i#* in the message description. For example,

```
STALLED AT sdate stime ID: s#.i#
```

IXC246E

Runtime Diagnostics examines IOS messages, occurring 1 minute before and 1 minute after the system issues IXC246E. It looks for IOS messages that contain the same *devnum* (device number for the data set). If Runtime Diagnostics finds additional IOS messages with the same *devnum*, it lists the messages with the IXC246E events.

IXC467I

Runtime Diagnostics does not list this message as an error event if the reason is system partitioning.

IXC585E

Runtime Diagnostics compares the structure name and physical structure version for multiple IXC585E messages. It lists only the last message that contains a match of both fields.

For more details, see [Chapter 6, “Messages that Runtime Diagnostics analyzes,”](#) on page 57.

Outstanding WTOR message analysis

Runtime Diagnostics utilizes consoles interfaces to find WTOR messages that are outstanding. If it finds WTOR messages that are critical, it determines if the message was issued for the system it is analyzing, which means the home system if that is the system being analyzed, or the target system, if the target system is not the same as the home system. When Runtime Diagnostics is following

blockers, which means when an event is issued that has a blocker that does not result in an individual ASID event (see “CORRELATED – Correlated events” on page 55), Runtime Diagnostics searches for WTOR messages issued on the system where the blocker is running as well as the system Runtime Diagnostics is analyzing. For some critical WTOR messages, events are created regardless of which system issued the message if the message is considered critical if found on any system in the sysplex.

After Runtime Diagnostics has found the outstanding WTORs for the right systems, it determines if the message is critical. If the message is critical, an event is created for that message individually and the reply ID is provided in the ACTION text.

```
f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2018/06/02 08:28:40
INTERVAL: 60 MINUTES
EVENTS FOUND: 1
PRIORITIES: HIGH:1 MED:0 LOW:0
TYPES: DUMP:1
-----
EVENT 1: HIGH:DUMP          SYSTEM: SY1      2018/06/02 08:28:40
IEA793A NO SVC DUMP DATA SETS AVAILABLE FOR DUMPID=dumpid FOR JOB (*MASTER*). USE THE
DUMPDS COMMAND OR REPLY D TO DELETE THE
ERROR : NO SVC DUMP DATA SETS ARE AVAILABLE TO CAPTURE A DUMP.
ACTION: DD ADD,SMS=class TO ADD SMS CLASSES OR DD ADD,VOL=volser TO
ACTION: ADD DASD VOLUMES OR REPLY 0053,D TO DELETE THE CAPTURED DUMP.
```

If the message is not critical and has been issued in the last 24 hours, Runtime Diagnostics creates one event for OTHER WTORs which lists the non-critical outstanding WTORs. A maximum of 8 WTORs are displayed in an OTHER WTORs event. Multiple OTHER WTORs events will be created until all non-critical outstanding WTORs issued in the last 24 hours are printed.

```
f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2018/04/20 11:48:18
INTERVAL: 60 MINUTES
EVENTS FOUND: 1
PRIORITIES: HIGH:0 MED:0 LOW:1
TYPES: OTHER WTORs:1
-----
EVENT 1: LOW :OTHER WTORs   SYSTEM: SY1      2018/04/20 11:48:18
REPLY MESSAGE      TIME      SYSNAME  MESSAGE TEXT
0020 XYZ1234        11:48:16 SY1      AN OTHER WTOR THAT SHOWS A PROBLEM
0019 OTHER253      10:58:01 SY1      ANOTHER WTOR THAT IS NOT CRITICAL
0018 OTHER253      08:12:22 SY1      ANOTHER WTOR THAT IS NOT CRITICAL
0015 ABC1234        02:25:18 SY1      THERE IS A PROBLEM WITH THE ABC PRO
0014 KKAOTHER3     23:08:46 SY1      A REPLY IS NEEDED FOR THIS MESSAGE
ERROR : AN OUTSTANDING MESSAGE REPLY MAY BE DELAYING SYSTEM PROCESSING
ACTION: DETERMINE WHETHER TO RESPOND TO ANY OF THE ABOVE MESSAGES.
ACTION: RESPOND USING REPLY REPLY,RESPONSE WHERE REPLY IS THE VALUE
ACTION: IN THE REPLY COLUMN.
```

HIGHCPU - Address spaces with high processor usage (CPU analysis)

Runtime Diagnostics provides a point in time check of any address space that is using more than 95% of the capacity of a single CPU. High CPU usage might indicate that the address space is in a loop. The analysis is a one second sample interval based on the capacity of a single CPU within the LPAR. It is possible for the usage to be reported greater than 100% if the address space has multiple TCBS and several are each using a high percentage of the CPU capacity.

```

f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2015/06/02 08:23:01
INTERVAL: 60 MINUTES
EVENTS FOUND: 2
PRIORITIES: HIGH:2 MED:0 LOW:0
TYPES: HIGHCPU:1 LOCK:1
-----
EVENT 1: HIGH:HIGHCPU          SYSTEM: SY1      2015/06/02 08:23:01
ASID CPU RATE: 100% ASID: 0029 JOBNAME: RTDLP
STEPNAME: STEP1 PROCSTEP:          JOBID: JOB00041
USERID: WELLIE0 JOBSTART: 2015/06/02 08:17:50
ERROR : ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MIGHT BE LOOPING.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 2: HIGH:LOCK            SYSTEM: SY1      2015/06/02 08:23:01
HIGH LOCAL LOCK SUSPENSION RATE IN ASID: 002A JOBNAME: RTDLOCK6
STEPNAME: STEP1 PROCSTEP:          JOBID: JOB00045
USERID: WELLIE0 JOBSTART: 2015/06/02 08:23:09
ERROR : ADDRESS SPACE HAS HIGH LOCAL LOCK SUSPENSION RATE.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----

```

Figure 8. Runtime Diagnostics HIGHCPU and LOCK report

More details about CPU activity are found in [z/OS RMF User's Guide](#).

LOOP - Loop detection

Runtime Diagnostics looks through all tasks in all address spaces to determine whether a task is looping. Runtime Diagnostics examines various system information for indicators of consistent repetitive activity that are typical when a task is in a loop. Runtime Diagnostics detects both TCB enabled and TCB recovery loops. When both a HIGHCPU event and a LOOP event list the same job name, the task in the job is likely in a loop. The normal corrective action is to cancel the job name listed.

```

f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2015/06/02 08:23:01
INTERVAL: 60 MINUTES
EVENTS FOUND: 2
PRIORITIES: HIGH:2 MED:0 LOW:0
TYPES: HIGHCPU:1 LOOP:1
-----
EVENT 1: HIGH:HIGHCPU          SYSTEM: SY1      2015/06/02 08:23:01
ASID CPU RATE: 100% ASID: 002A JOBNAME: RTDLOOP
STEPNAME: STEP1 PROCSTEP:          JOBID: JOB00042
USERID: WELLIE0 JOBSTART: 2015/06/02 08:17:50
ERROR : ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MIGHT BE LOOPING.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 1: HIGH:LOOP            SYSTEM: SY1      2015/06/02 08:23:07
ASID: 0029 JOBNAME: RTDLP TCB: 008FE990
STEPNAME: STEP1 PROCSTEP:          JOBID: JOB00041
USERID: WELLIE0 JOBSTART: 2015/06/02 08:17:50
ERROR : ADDRESS SPACE MIGHT BE IN A TCB-ENABLED LOOP.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----

```

Figure 9. Runtime Diagnostics LOOP and HIGHCPU report

LOCK - Local lock suspension

Runtime Diagnostics provides a point in time check of local lock suspension for all address spaces. For the local lock suspension, Runtime Diagnostics calculates the amount of time an ASID is suspended waiting for the local lock. When an ASID is suspended more than 50% of the time waiting for a local lock, Runtime Diagnostics reports an event. For an example, see [Figure 8 on page 44](#).

ENQ - ENQ contention checking

Runtime Diagnostics provides a point in time check of ENQ contention equivalent to issuing the D GRS, AN, WAITER command. It compares the list of job names that are waiters with a hardcoded list of IBM-supplied address spaces (shown in [Table 4 on page 45](#)). An IBM-supplied system address

space must be in contention for **5 seconds** or more for RTD to report it. Any other non-system address space must be in contention **5 minutes** or more for RTD to report it.

If Runtime Diagnostics finds ENQ contention, it issues an error event within message HZR0200I stating the job name of the top waiter for the ENQ resource and the top blocker of the ENQ resource. The top waiter is the waiter that has been waiting for the resource the longest. The top blocker is the blocker that has been holding the resource the longest. Other waiters for the resource are listed in the report as OTHER WAITERS.

Table 4. ENQ checking: IBM-supplied address spaces

IBM-supplied address spaces				
• *MASTER*	• ALLOCAS	• ANTMMAIN	• ANTAS000	• BPXOINT
• CEA	• CONSOLE	• DEVMAIN	• DFHSM	• DUMPSRV
• GRS	• IOSAS	• IEFSCHAF	• IXGLOGR	• JES2
• JES2XCF	• JES2AUX	• JES2MON	• LLA	• OMVS
• PCAUTH	• RACF	• RASP	• SMS	• SMSPDSE
• SMSPDSE1	• SMSVSAM	• TRACE	• VLF	• WLM
• XCFAS	• ZFS			

```
f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2015/06/02 08:23:00
INTERVAL: 60 MINUTES
EVENTS FOUND: 1
PRIORITIES: HIGH:1 MED:0 LOW:0
TYPES: ENQ:1
-----
EVENT 1: HIGH:ENQ              SYSTEM: SY1      2015/06/02 08:23:00
QNAME: SYSZMCS  SCOPE: SYSTEMS
RNAME: SYSZMCS#MCS
      ASID  JOB NAME  TCB/WEB  SYSTEM  WAIT TIME
TOP WAITER : 002A  FENQ2    008C9BE0 SY1     00:05:05
TOP BLOCKER: 002A  FENQ2    008C9E00 SY1
OTHER WAITERS FOR THIS RESOURCE:
  WAITER: 0029  FENQ3    008FE990 SY2     00:02:13
  WAITER: 003B  FENQ5    008C5540 SY1     00:02:08
ERROR : ADDRESS SPACES MIGHT BE IN ENQ CONTENTION.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE BLOCKING JOBS AND
ACTION: ASIDS.
-----
```

Figure 10. Runtime Diagnostics ENQ contention report

Runtime diagnostics can invoke itself internally to run on other systems in the SYSPLEX to find additional problems with the top blocker address spaces of ENQ contention events. Runtime Diagnostics follows the blocker, and keeps following blockers until one of the following occurs:

- Runtime Diagnostics finds an individual, final event for a top blocker.
- Runtime Diagnostics does not find an individual, final event for a top blocker, so it searches OPERLOG for critical WTO messages and analyzes outstanding WTOR messages which occurred on the system on which the last top blocker resides and the system for which Runtime Diagnostics was originally invoked (whether different or the same).

In either case, all the related events that Runtime Diagnostics finds from the first contention event to the last individual event are displayed as CORRELATED events. The OPERLOG messages and the outstanding WTOR messages, however, are not correlated with the others.

To resolve contention, do the following:

- First, look for the FINAL EVENT of the CORRELATED EVENT (see [Correlated events](#) for more information).
- Determine if the blocking job is running properly or if you must cancel it. See the [Serialization summary](#) topic in *z/OS MVS Diagnosis: Reference* for a list of the ENQ names and corresponding issuing component.

Find additional contention information in the following topics:

- To identify modules, components, and products, see the topic on [Identifying modules, components, and products](#) in *z/OS MVS Diagnosis: Reference*.
- To understand ENQ contention and additional analysis steps, see the topic on [Contention management](#) in , .
- For System Logger, see the topic on [Associating latch contention with a logger TCB or WEB](#) in *z/OS MVS Diagnosis: Reference*.

LATCH - GRS latch contention

The following example shows the Runtime Diagnostics event record that summarizes latch contention. When latch contention exists for any address space for at least five minutes, a latch contention event is issued. Follow the ACTION statement in the summary to determine the source of the contention and further actions.

If Runtime Diagnostics finds contention, it issues an error event within message HZR0200I stating the job name of the top waiter for the latch and the top blocker of the latch resource. The top waiter is the waiter that has been waiting for the resource the longest. The top blocker is the blocker that has been holding the resource the longest. Other waiters for the resource are listed in the report as OTHER WAITERS.

```

HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2015/06/02 10:49:30
INTERVAL: 60 MINUTES
EVENTS FOUND: 2
PRIORITIES: HIGH:2 MED:0 LOW:0
TYPES: LATCH:2
-----
EVENT 1: HIGH:LATCH           SYSTEM: SY1      2015/06/02 10:49:30
LATCH SET NAME: SYSTEST.LATCH_TESTSET
LATCH NUMBER: 3              CASID: 0034 CJOBNAME: TSTLATC0
ASID JOB NAME TCB/WEB SYSTEM WAIT TIME
TOP WAITER : 0034 TSTLATC0 008F8368 SY1      00:05:14
TOP BLOCKER: 0034 TSTLATC0 008F8588 SY1
ERROR : ADDRESS SPACES MIGHT BE IN LATCH CONTENTION.
ACTION: D GRS,AN,LATCH,DEP,CASID=0034,LAT=(SYSTEST.L*,3),DET
ACTION: TO ANALYZE THE LATCH DEPENDENCIES. USE YOUR SOFTWARE MONITORS
ACTION: TO INVESTIGATE BLOCKING JOBS AND ASIDS.
-----
EVENT 2: HIGH:LATCH           SYSTEM: SY1      2015/06/02 10:49:30
LATCH SET NAME: SYSTEST.LATCH_TESTSET
LATCH NUMBER: 12            CASID: 0035 CJOBNAME: TSTLATC9
ASID JOB NAME TCB/WEB SYSTEM WAIT TIME
TOP WAITER : 0035 TSTLATC9 008F8368 SY1      00:05:13
TOP BLOCKER: 0035 TSTLATC9 008F8588 SY1
OTHER WAITERS FOR THIS RESOURCE:
WAITER: 0036 TSTLATC5 008F8368 SY1      00:03:25
ERROR : ADDRESS SPACES MIGHT BE IN LATCH CONTENTION.
ACTION: D GRS,AN,LATCH,DEP,CASID=0035,LAT=(SYSTEST.L*,12),DET
ACTION: TO ANALYZE THE LATCH DEPENDENCIES. USE YOUR SOFTWARE MONITORS
ACTION: TO INVESTIGATE BLOCKING JOBS AND ASIDS.
-----

```

Figure 11. HZR event for the GRS latch contention event record

Runtime diagnostics looks for additional problems with the top blocker address spaces of LATCH contention events. Runtime Diagnostics follows that blocker, and keeps following blockers until it finds an individual, final event for a top blocker or finds an ENQ contention event where the top blocker is a waiter.

OMVS - z/OS UNIX latch contention

Runtime Diagnostics gathers information about z/OS UNIX file system contention. When z/OS UNIX latch contention or waiting threads exist for at least five minutes in z/OS UNIX, the z/OS UNIX file system latch contention and waiting threads record (as shown in [Figure 12 on page 47](#)) displays at the console.

```
HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2015/06/02 10:49:30
INTERVAL: 60 MINUTES
EVENTS FOUND: 2
PRIORITIES: HIGH:1 MED:0 LOW:0
TYPES: OMVS:1
-----
EVENT 1: HIGH:OMVS              SYSTEM: SY1      2015/06/02 10:49:30
ASID: 000E JOBNAME: OMVS
MOUNT LATCH WAITERS: 1
FILE SYSTEM LATCH WAITERS: 0
XSYS AND OTHER THREADS WAITING FOR z/OS UNIX: 1
ERROR : z/OS UNIX MIGHT HAVE FILE SYSTEM LATCH CONTENTION.
ACTION: D OMVS,W,A TO INVESTIGATE z/OS UNIX FILE SYSTEM LATCH
ACTION: CONTENTION, ACTIVITY AND WAITING THREADS. USE YOUR SOFTWARE
ACTION: MONITORS TO INVESTIGATE BLOCKING JOBS AND ASIDS.
-----
```

Figure 12. HZR event for the z/OS UNIX file system latch contention and waiting threads record

If all the counts are zero, the record does not display. Follow the instructions listed in the ACTION statement in the event record by issuing D OMVS,W,A command, which returns the ASID and job names of any latch waiters.

For more information about diagnosing and resolving z/OS UNIX latch contention, see the topics on [z/OS UNIX System Services and Diagnosing and resolving latch contention in z/OS MVS Diagnosis: Reference](#). For zFS file system contention, see [z/OS File System Administration](#).

JES2 - JES2 health exception event

Runtime Diagnostics gathers information about the JES2 subsystem. Health events returned from the JES2 Subsystem interface (SSI) are available here. An example of a possible event appears below:

```
HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2015/06/02 08:28:53
INTERVAL: 60 MINUTES
EVENTS FOUND: 1
PRIORITIES: HIGH:1 MED:0 LOW:0
TYPES: JES2:1
-----
EVENT 1: HIGH:JES2              SYSTEM: SY1      2015/06/02 08:28:53
$HASP9159 JES2 EXECUTION PROCESSING STOPPED ($PXEQ)
ERROR : JES2 CANNOT START ANY NEW BATCH JOBS.
ACTION: $SXEQ TO ENABLE JES2 TO START NEW BATCH JOBS.
-----
```

Figure 13. HZR event for JES2

SERVERHEALTH - Server address space health exceptions

Runtime Diagnostics identifies server health issues. As part of Runtime Diagnostic's normal diagnostic reporting, it obtains the health status of all address spaces whose health value has been set using the IWM4HLTH API (see [z/OS MVS Programming: Workload Management Services](#)). Runtime Diagnostics creates a SERVERHEALTH event in message HZR0200I for each address space whose current health value is less than 100. The diagnostic information surfaced by Runtime Diagnostics will assist in debugging server issues with the server specified in JOB NAME: on the system specified in SYSTEM: of the SERVERHEALTH event. An example of a possible event appears below:

```

HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 1 TARGET SYSTEM: SY1      HOME: SY1      2015/06/02 08:28:40
INTERVAL: 60 MINUTES
EVENTS FOUND: 1
PRIORITIES: HIGH:1 MED:0 LOW:0
TYPES: SERVERHEALTH:1
-----
EVENT 1: HIGH:SERVERHEALTH      SYSTEM: SY1      2015/06/02 08:28:40
JOB NAME: LDAP      ASID: 0037      CURRENT HEALTH VALUE: 0
CURRENT LOWEST HEALTH VALUES:
SUBSYSTEM HEALTH      REPORTED
SUBSYSTEM NAME      SETTING      REASON      DATE AND TIME
XCF      MYXCF1      80      0000000000000001      2015/06/02 08:28:38
XES      MYXCF2      0      000000000000000C      2015/06/02 08:28:38
ERROR : ADDRESS SPACE SERVER CURRENT HEALTH VALUE LESS THAN 100.
ERROR : THIS VALUE MAY IMPACT YOUR SYSTEM OR SYSPLEX TRANSACTION
ERROR : PROCESSING.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID AND TO
ACTION: DETERMINE THE IMPACT OF THE HEALTH OF THE ADDRESS SPACE TO
ACTION: OVERALL TRANSACTION PROCESSING.
-----

```

Figure 14. HZR event for server address spaces

The report contains the following fields:

JOB NAME

ASID

The name and ASID of the address space with this health value less than 100.

CURRENT HEALTH VALUE

The current health indicator for the address space.

CURRENT LOWEST HEALTH VALUES

The list of the subsystem and subsystem names reporting the health values for the address space since the last reset occurred.

REASON

A 16-byte diagnostic reason representing the problem identified by the reporter. This value should be included if reporting the problem to IBM service.

DEADLOCK - Deadlock events

Runtime Diagnostics gathers information about enqueues and latches on the home system, which is the system where it is running. It also gathers enqueue and latch information on each system where Runtime Diagnostics invokes itself internally to follow contention blockers. Once the information is gathered, Runtime Diagnostics analyzes the data to detect deadlocks. If any deadlocks are found, Runtime Diagnostics creates a DEADLOCK event and lists all of the resources along with their WAITERS and BLOCKERS involved in the deadlock.

Runtime Diagnostics does not invoke itself internally on every system in the sysplex to retrieve this information, nor does it invoke itself on the target system when the home system is not equal to the target system. Rather, the information is gathered on the home system and on every system in the sysplex where Runtime Diagnostics follows blockers. Therefore, deadlocks can be found on the home system and across systems in the sysplex to which blockers are followed and can be found in a combination of GRS latches and enqueues. Note that the waiters do not need to be waiting any specified length of time for a deadlock event to be created nor do the TOP waiter or the TOP blocker need to be involved in the deadlock. Examples of possible events follow:

Example 1:

This example shows a deadlock between 3 tasks in an address space across 3 ENQs:

```

EVENT 1: HIGH:DEADLOCK          SYSTEM: SY1          2015/06/02 08:49:10
RESOURCE: ENQ
QNAME: SYSZDLK  SCOPE: SYSPLEX
RNAME: DLKEN32
      ASID  JOB NAME  TCB/WEB  SYSTEM  WAIT TIME
TOP WAITER : 0029  DLKSTART 008C5A60  SY1     00:02:23
BLOCKER    : 0029  DLKSTART 008C5CF0  SY1
OTHER WAITERS FOR THIS RESOURCE:
  WAITER: 0029  DLKSTART 008FE990  SY1     00:02:13
  WAITER: 0029  DLKSTART 008C5540  SY1     00:02:08
  WAITER: 0029  DLKSTART 008C57D0  SY1     00:02:08
-----
RESOURCE: ENQ
QNAME: SYSZDLK  SCOPE: SYSPLEX
RNAME: DLKEN31
      ASID  JOB NAME  TCB/WEB  SYSTEM  WAIT TIME
TOP WAITER : 0029  DLKSTART 008C5CF0  SY1     00:02:23
BLOCKER    : 003A  DLKSTART 008C5E88  SY2
-----
RESOURCE: ENQ
QNAME: SYSZDLK  SCOPE: SYSPLEX
RNAME: DLKEN33
      ASID  JOB NAME  TCB/WEB  SYSTEM  WAIT TIME
TOP WAITER : 003A  DLKSTART 008C5E88  SY2     00:02:23
TOP BLOCKER: 0029  DLKSTART 008C5A60  SY1
ERROR : ADDRESS SPACES WERE DEADLOCKED AT THE TIME OF THE ANALYZE
ERROR : REQUEST.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE BLOCKERS TO
ACTION: DETERMINE IF THE DEADLOCK STILL EXISTS AND TAKE APPROPRIATE
ACTION: ACTION TO ELIMINATE THE DEADLOCK.
-----

```

Figure 15. Deadlock event between 3 tasks in an address space across 3 ENQs

Example 2:

This example shows a deadlock between two tasks in an address space between an ENQ and a LATCH:

```

EVENT 2: HIGH:DEADLOCK          SYSTEM: SY1          2015/06/02 09:01:40
RESOURCE: ENQ
QNAME: SYSZDLK  SCOPE: STEP
RNAME: DLKCMB1
      ASID  JOB NAME  TCB/WEB  SYSTEM  WAIT TIME
TOP WAITER : 0029  DLKSTART 008C5E88  SY1     00:00:49
BLOCKER    : 0029  DLKSTART 008C5CF0  SY1
-----
RESOURCE: LATCH
LATCH SET NAME: SYSRTD.DLKCMB1TEST.LATCHSET
LATCH NUMBER: 0          CASID: 0029  CJOBNAME: DLKSTART
      ASID  JOB NAME  TCB/WEB  SYSTEM  WAIT TIME
TOP WAITER : 0029  DLKSTART 008C5CF0  SY1     00:00:49
TOP BLOCKER: 0029  DLKSTART 008C5E88  SY1
OTHER WAITERS FOR THIS RESOURCE:
  WAITER: 0029  DLKSTART 008FE990  SY1     00:00:33
ERROR : ADDRESS SPACES WERE DEADLOCKED AT THE TIME OF THE ANALYZE
ERROR : REQUEST.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE BLOCKERS TO
ACTION: DETERMINE IF THE DEADLOCK STILL EXISTS AND TAKE APPROPRIATE
ACTION: ACTION TO ELIMINATE THE DEADLOCK.
-----

```

Figure 16. Deadlock event between 2 tasks in an address space between an ENQ and a LATCH

CORRELATED - Correlated events

As Runtime Diagnostics detects events, it analyzes the detected events to determine if any of them are for the same address spaces or are related due to following blockers. When it finds events for the same address space, Runtime Diagnostics correlates the events into one event, a CORRELATED event.

The following are some of the types of events that are correlated:

- Events where the blocker job of a contention event (such as a GRS latch or ENQ event) is the job specified in another event that has only one job specified (such as a LOOP event).

- Events where the blocker job of a contention event (such as a GRS latch or ENQ event) is the waiter job of another contention event.
- Events where the job specified in the event is the same job specified for another event (such as a HIGHCPU and a LOOP event).

The following are some events that are not correlated:

- JES2 events are not correlated because each is a standalone event that requires its own action even though they are all for the JES2 address space.
- OPERLOG messages are not correlated because they are not for a specific address space.
- Outstanding WTOR messages are not correlated because they are not for a specific address space.

The header for correlated events reports looks like the header shown in [Figure 3 on page 38](#).

The layout of the body of the correlated report shown in [Figure 17 on page 51](#) is as follows:

- The priority for a correlated event is the highest priority of all the final and related events.
- The FINAL EVENT is the last event found for the correlated event while following the blockers or correlating individual events. The FINAL EVENT is always displayed first and has both ERROR : and ACTION: lines. Take action on the FINAL EVENT first to resolve the problem.
 - If the final event is a contention event, the blocker might be experiencing a problem that Runtime Diagnostics does not identify.
 - If the final event is an event with only one job name (for example, LOOP, LOCK, or SERVERHEALTH events), that event is probably the root problem, at least of the problems that Runtime Diagnostics can identify.
- The RELATED EVENTS are displayed in the order found - the first related event shown is the first event found and all related events after it are a result of following the blocker of each preceding event. The RELATED EVENTS list the ERROR but do not list an ACTION: field.
- The system name in the FINAL EVENT and the RELATED EVENT event lines is where the analysis was done by Runtime Diagnostics when the event was found. Note that when Runtime Diagnostics follows blockers to other systems, events for those blockers can be created for the other systems.
- When one or more correlated events exist, the number of correlated types are found in the TYPES: line. The types within the correlated events are described in following line CORRELATED TYPES.

To follow correlated events, start with the first RELATED EVENT after the FINAL EVENT and refer to the numbers in [Figure 17 on page 51](#).

1 and **2** - The TOP BLOCKER of each RELATED EVENT is a WAITER for the next related event down.

3 - In the last RELATED EVENT, the TOP BLOCKER always points back to the WAITER for the FINAL EVENT if the final event is a contention event. If the final event is an event with a single address space identified, the TOP BLOCKER of the last RELATED EVENT is the address space identified in the final event.

```

EVENT 2: HIGH:CORRELATED          SYSTEM: SY1      2016/01/29 16:20:13
FINAL EVENT FOR EVENT 2: HIGH:LATCH          SYSTEM: SY1
LATCH SET NAME: VAR078#SET1
LATCH NUMBER: 0          CASID: 0041  CJOBNAME: SVRASID3
          ASID  JOB NAME  TCB/WEB  SYSTEM  WAIT TIME
TOP WAITER : 0046  SVRASID2  004F8460  SY1    00:20:14
TOP BLOCKER: 0041  SVRASID3  004F8460  SY1
ERROR : ADDRESS SPACES MIGHT BE IN LATCH CONTENTION.
ACTION: D GRS,AN,LATCH,DEP,CASID=0041,LAT=(VAR078#SE*,0),DET
ACTION: TO ANALYZE THE LATCH DEPENDENCIES. USE YOUR SOFTWARE MONITORS
ACTION: TO INVESTIGATE BLOCKING JOBS AND ASIDS.
-----
RELATED EVENT FOR EVENT 2: HIGH:ENQ          SYSTEM: SY1
QNAME: SYSZFCT0  SCOPE: SYSPLEX
RNAME: RES1
          ASID  JOB NAME  TCB/WEB  SYSTEM  WAIT TIME
TOP WAITER : 0047  SVRASID1  004F8460  SY1    00:25:20
TOP BLOCKER: 0045  SVRASID1  004F8460  SY21
ERROR : ADDRESS SPACES MIGHT BE IN ENQ CONTENTION.
-----
RELATED EVENT FOR EVENT 2: HIGH:LATCH          SYSTEM: SY2
LATCH SET NAME: VAR078#SET1
LATCH NUMBER: 0          CASID: 0041  CJOBNAME: SVRASID2
          ASID  JOB NAME  TCB/WEB  SYSTEM  WAIT TIME
TOP WAITER : 0045  SVRASID1  004F8460  SY2    00:25:17
TOP BLOCKER: 0041  SVRASID2  004F8460  SY22
ERROR : ADDRESS SPACES MIGHT BE IN LATCH CONTENTION.
-----
RELATED EVENT FOR EVENT 2: HIGH:ENQ          SYSTEM: SY2
QNAME: SYSZFCT0  SCOPE: SYSPLEX
RNAME: RES2
          ASID  JOB NAME  TCB/WEB  SYSTEM  WAIT TIME
TOP WAITER : 0041  SVRASID2  004F8460  SY2    00:25:17
TOP BLOCKER: 0046  SVRASID2  004F8460  SY13
ERROR : ADDRESS SPACES MIGHT BE IN ENQ CONTENTION.

```

Figure 17. Reading correlated events

Example 1:

Figure 18 on page 52 shows ENQ contention found initially between jobs W1 and B1. Runtime Diagnostics then discovers that job B1 was waiting for an ENQ held by job B2 on another system. Runtime Diagnostics invokes itself internally on the other system and discovers that job B2 is looping. The FINAL EVENT found is the LOOP event for job B2. The first RELATED EVENT shows the first ENQ event that led to the analysis of job B1 which resulted in the second ENQ event and the analysis of job B2.

```

HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 8 TARGET SYSTEM: SY1      HOME: SY1      2014/11/18 16:09:42
INTERVAL: 60 MINUTES
EVENTS FOUND: 1
PRIORITIES: HIGH:1 MED:0 LOW:0
TYPES: CORRELATED:1
CORRELATED TYPES: ENQ:2 LOOP:1
PROCESSING BYPASSED FOR SYSTEM SY1:
OPERLOG..... OPERLOG IS NOT ACTIVE.

```

```

-----
EVENT 1: HIGH:CORRELATED          SYSTEM: SYS1      2014/11/18 16:09:54
FINAL EVENT FOR EVENT 1: HIGH:LOOP SYSTEM: SYS3
ASID: 0365 JOBNAME: B2           TCB: 004FF1C0
STEPNAME: STEP1   PROCSTEP:          JOBID: JOB00045
USERID: IBMUSER  JOBSTART: 2014/11/18 15:22:34
ERROR : ADDRESS SPACE MIGHT BE IN A LOOP ON SYS3.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.

```

```

-----
RELATED EVENT FOR EVENT 1: HIGH:ENQ          SYSTEM: SYS1
QNAME: MYQNAME  SCOPE: SYSPLEX
RNAME: MYRNAME

```

	ASID	JOB NAME	TCB/WEB	SYSTEM	WAIT TIME
TOP WAITER :	002C	W1	004E2A80	SYS1	00:05:53
TOP BLOCKER:	005A	B1	004E2A70	SYS2	

ERROR : ADDRESS SPACES MIGHT BE IN ENQ CONTENTION ON SYS1.

```

-----
RELATED EVENT FOR EVENT 1: HIGH:ENQ          SYSTEM: SYS1
QNAME: MYQNAME2 SCOPE: SYSPLEX
RNAME: MYRNAME2

```

	ASID	JOB NAME	TCB/WEB	SYSTEM	WAIT TIME
WAITER :	005A	B1	004E2A70	SYS2	00:06:25
TOP BLOCKER:	0365	B2	004FF1C0	SYS3	

OTHER WAITERS FOR THIS RESOURCE:
TOP WAITER: 0253 W2 004FF2C0 SYS3 00:09:25
ERROR : ADDRESS SPACES MIGHT BE IN ENQ CONTENTION ON SYS2.

Figure 18. Correlated ENQ contention found between three jobs across three systems

Example 2:

Figure 19 on page 53 shows two correlated events:

- Address space RTDLP is using a high amount of CPU and is looping.
- Address space RTDLOCK6 has a high local lock suspension rate and is using a high amount of CPU.

Notice that the CORRELATED TYPES section shows 2 HIGHCPU events and one each of LOCK and LOOP.

```

HZR0200I RUNTIME DIAGNOSTICS RESULT
SUMMARY: SUCCESS
REQ: 8 TARGET SYSTEM: SY1      HOME: SY1      2015/07/28 08:46:24
INTERVAL: 60 MINUTES
EVENTS FOUND: 2
PRIORITIES: HIGH:2 MED:0 LOW:0
TYPES: CORRELATED:2
CORRELATED TYPES: HIGHCPU:2 LOOP:1 LOCK:1
PROCESSING BYPASSED FOR SYSTEM SY1:
OPERLOG..... OPERLOG IS NOT ACTIVE.
-----
EVENT 1: HIGH:CORRELATED      SYSTEM: SY1      2015/07/28 08:46:12
FINAL EVENT FOR EVENT 1: HIGH:LOOP      SYSTEM: SY1
ASID: 0029 JOBNAME: RTDLP      TCB: 008FE990
STEPNAME: STEP1      PROCSTEP:      JOBID: JOB00075
USERID: WELLIE0      JOBSTART: 2015/07/28 08:44:46
ERROR : THE ADDRESS SPACE MIGHT BE IN A TCB-ENABLED LOOP.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
RELATED EVENT FOR EVENT 1: HIGH:HIGHCPU      SYSTEM: SY1
ASID CPU RATE: 99%      ASID: 0029 JOBNAME: RTDLP
STEPNAME: STEP1      PROCSTEP:      JOBID: JOB00075
USERID: WELLIE0      JOBSTART: 2015/07/28 08:44:46
ERROR : ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MIGHT BE LOOPING.
-----
EVENT 2: HIGH:CORRELATED      SYSTEM: SY1      2015/07/28 08:46:25
FINAL EVENT FOR EVENT 2: HIGH:LOCK      SYSTEM: SY1
HIGH LOCAL LOCK SUSPENSION RATE IN ASID: 002A JOBNAME: RTDLOCK6
STEPNAME: TCSTRT      PROCSTEP:      JOBID: JOB00076
USERID: WELLIE0      JOBSTART: 2015/07/28 08:44:46
ERROR : ADDRESS SPACE HAS HIGH LOCAL LOCK SUSPENSION RATE.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
RELATED EVENT FOR EVENT 2: HIGH:HIGHCPU      SYSTEM: SY1
ASID CPU RATE: 100%      ASID: 002A JOBNAME: RTDLOCK6
STEPNAME: TCSTRT      PROCSTEP:      JOBID: JOB00076
USERID: WELLIE0      JOBSTART: 2015/07/28 08:44:46
ERROR : ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MIGHT BE LOOPING.
-----

```

Figure 19. 2 Correlated events across two address spaces

Syntax and parameters for the MODIFY HZR command

Purpose

Restriction: Do not enable the debug option unless under the direction of IBM support.

If IBM Service determines they need more Runtime Diagnostics data, they might request that you specify the debug options. The `MODIFY HZR, ANALYZE, SYSNAME=sysname, DEBUG=(option, option)` command initiates an SVC dump to collect additional problem data.

Note: If the `DEBUG` option is requested when there is no event detected, only the Runtime Diagnostics address space is dumped. If the `DEBUG` option is requested when an event is detected, both the HZR address space and the specific address spaces that are identified in the events are dumped, up to a maximum of 15 total address spaces.. If any data spaces exist for the address spaces dumped, they are also included in the dump.

Format

```

MODIFY HZR, ANALYZE[, SYSNAME=sysname] ([ (, DEBUG=ALL) | (LOOP|NOLOOP), (ENQ|NOENQ),
(LOCK|NOLOCK), (HIGHCPU|NOHIGHCPU), (MSGS|NOMSGS), (CRITICALWTO|NOCRITICALWTO), (CRITICALWTOR|
NOCRITICALNWTOR), (OTHERWTOR|NOOTHERWTOR), (OMVS|NOOMVS), (LATCH|NOLATCH),
(SERVERHEALTH|NOSERVERHEALTH), (JES2|NOJES2), (DEADLOCK|NODEADLOCK) ])
([, SCOPE=(FOLLOW|LOCAL)])

```

Options

SYSNAME

The target system that Runtime Diagnostics analyzes. By default, all analysis is conducted for the home system.

Note: When the target system is not the home system, Runtime Diagnostics processes only the following debug options:

- ALL
- ENQ
- NOENQ
- MSGS
- NOMSGS
- CRITICALWTO
- NOCRITICALNOWTO
- CRITICALWTOR
- NOCRITICALNOWTOR
- OTHERWTOR
- NOOTHERWTOR
- DEADLOCK
- NODEADLOCK

DEBUG

Indicates conditions under which Runtime Diagnostics initiates an SVC dump of the HZR address space and the address spaces that are associated with events displayed in message HZR0200I. Up to 15 address spaces can be reported in the dump, where one of those 15 is always the HZR address space. The debug option controls only when the system dump occurs; the debug option does not affect the types of analysis performed by Runtime Diagnostics.

Tip: Before using this option, ensure that the system is able to take an SVC dump. To determine which dump options are currently in effect, use the DISPLAY DUMP command.

When entering two or more debug options, use parenthesis except with ALL, which is mutually exclusive. For example:

```
F HZR,ANALYZE,DEBUG=ALL
F HZR,ANALYZE,DEBUG=(LOOP,NOENQ)
```

Remember: Use the debug option only when under the direction of IBM Service.

ALL

Runtime Diagnostics initiates an SVC dump of the HZR address space and up to 14 address spaces associated with each event that is displayed in message HZR0200I.

Note: The keyword ALL supersedes all other keywords. For Example, if you specify both All and MSG on the ANALYZE command, F HZR,ANALYZE,DEBUG=(ALL,MSG), MSGS is ignored.

LOOP

Runtime Diagnostics initiates an SVC dump of the HZR address space, along with a maximum of 14 address spaces that are associated with each LOOP event that displays in the HZR0200I message. If any data spaces exist for the address spaces dumped, they are also included in the dump.

NOLOOP

When no LOOP events are found during the ANALYZE request, Runtime Diagnostics dumps the HZR address space.

ENQ

Runtime Diagnostics initiates an SVC dump of the HZR address space along with a maximum of 14 address spaces that are associated with each ENQ event that displays in the HZR0200I message. If any data spaces exist for the address spaces dumped, they are also included in the dump.

NOENQ

When no ENQ events are found during the ANALYZE request, Runtime Diagnostics dumps the HZR address space.

LOCK

Runtime Diagnostics initiates an SVC dump of the HZR address space along with a maximum of 14 address spaces that are associated with each LOCK event that displays in the HZR0200I message. If any data spaces exist for the address spaces dumped, they are also included in the dump.

NOLOCK

When no LOCK events are found during the ANALYZE request,

HIGHCPU

Runtime Diagnostics initiates an SVC dump of the HZR address space along with a maximum of 14 address spaces that are associated with each HIGHCPU event that displays in the HZR0200I message. If any data spaces exist for the address spaces dumped, they are also included in the dump.

NOHIGHCPU

When no HIGHCPU events are found during the ANALYZE request, Runtime Diagnostics dumps the HZR address space.

MSGS

Runtime Diagnostics initiates an SVC dump of the HZR address space when it finds CRITICALWTO, CRITICALWTOR, or OTHERWTOR messages during the ANALYZE request.

NOMSGS

When no CRITICALWTO, no CRITICALWTOR, and no OTHERWTOR events are found during the ANALYZE request, Runtime Diagnostics dumps the HZR address space.

CRITICALWTO

Runtime Diagnostics initiates an SVC dump of the HZR address space when any CRITICALWTO events are found during the ANALYZE request.

NOCRITICALWTO

When no CRITICALWTO events are found during the ANALYZE request, Runtime Diagnostics dumps the HZR address space.

CRITICALWTOR

Runtime Diagnostics initiates an SVC dump of the HZR address space when any CRITICALWTOR events are found during the ANALYZE request.

NOCRITICALWTOR

When no CRITICALWTOR events are found during the ANALYZE request, Runtime Diagnostics dumps the HZR address space.

OTHERWTOR

Runtime Diagnostics initiates an SVC dump of the HZR address space when any OTHERWTOR events are found during the ANALYZE request.

NOOTHERWTOR

When no OTHERWTOR events are found during the ANALYZE request, Runtime Diagnostics dumps the HZR address space.

OMVS

Runtime Diagnostics initiates an SVC dump of the HZR address space, along with a maximum of 14 address spaces that are associated with each OMVS event that displays in the HZR0200I message. If any data spaces exist for the address spaces dumped, they are also included in the dump.

NOOMVS

When no OMVS events are found during the ANALYZE request, Runtime Diagnostics dumps the HZR address space.

LATCH

Runtime Diagnostics initiates an SVC dump of the HZR address space along with a maximum of 14 address spaces that are associated with each LATCH event that displays in the HZR0200I message. If any data spaces exist for the address spaces dumped, they are also included in the dump.

NOLATCH

When no LATCH events are found, Runtime Diagnostics dumps the HZR address space .

SERVERHEALTH

Runtime Diagnostics initiates an SVC dump of the HZR address space along with a maximum of 14 address spaces (JOBNAME server address spaces) that are associated with each SERVERHEALTH event that displays in the HZR0200I message. If any data spaces exist for the address spaces dumped, they are also included in the dump.

NOSERVERHEALTH

When no SERVERHEALTH events are found during the ANALYZE request, Runtime Diagnostics dumps the HZR address space.

JES2

Runtime Diagnostics initiates an SVC dump of the HZR address space along with an SVC dump of the JES2 address space. If any data spaces exist for the address spaces dumped, they are also included in the dump.

NOJES2

When no JES2 events are found, Runtime Diagnostics dumps the HZR address space.

DEADLOCK

Runtime Diagnostics initiates an SVC dump of the HZR address space along with a maximum of 14 address spaces associated with each DEADLOCK event displayed in message HZR0200I. If any data spaces exist for the address spaces dumped, they are also included in the dump.

NODEADLOCK

When no DEADLOCK events are found, Runtime Diagnostics dumps the HZR address space.

SCOPE=FOLLOW|LOCAL**FOLLOW**

When you specify or default to SCOPE=FOLLOW, Runtime Diagnostics invokes itself internally on other systems if it finds contention blockers on other systems and no cause can be found on the original system being analyzed. If after following the blockers there are still unresolved blockers, Runtime Diagnostics searches the OPERLOG and does outstanding WTOR analysis for the system for which Runtime Diagnostics was invoked as well as the systems on which the unresolved blockers reside.

When Runtime Diagnostics follows blockers to another system and events are found that were not supported on the release of the system from which the command was invoked, the event type will be displayed as UNSUPPORTED.

LOCAL

When you specify SCOPE=LOCAL, Runtime Diagnostics does not invoke itself on other systems even if it finds blockers on another system or systems.

Examples**Example:**

Runtime Diagnostics shows a job consuming a high amount of CPU, but your monitors are not showing the job that is consuming high CPU. If instructed by IBM Service, add the following debug option to the MODIFY command the next time Runtime Diagnostics runs:

```
F HZR , ANALYZE , DEBUG=HIGHCPU
```

If Runtime Diagnostics highlights the job with a *HIGHCPU* event, it initiates an SVC dump of relevant storage areas to assist IBM Service in diagnosing analysis processing.

Chapter 6. Messages that Runtime Diagnostics analyzes

- Runtime Diagnostics analyzes the following messages:
 - “BPX message analysis” on page 57
 - “IEA message analysis” on page 59
 - “IEE message analysis” on page 60
 - “IOS message analysis” on page 61
 - “IRA message analysis” on page 61
 - “IXC message analysis” on page 62
 - “IXG message analysis” on page 65
 - “IXL message analysis” on page 68
- Runtime Diagnostics does not analyze a certain group of test messages. See “Test messages ignored by Runtime Diagnostics” on page 70.

BPX message analysis

This topic covers the messages Runtime Diagnostics uses for its analysis.

For more information on BPX messages for z/OS UNIX System Services, see *z/OS MVS System Messages, Vol 3 (ASB-BPX)*.

BPXF006I	A FILE SYSTEM WITH FILESYSTYPE <i>type</i> FAILED TO INITIALIZE. IT TERMINATED DURING INITIALIZATION	Explanation Correct the problem in <i>bpxprmxx</i> or the superuser can enter correct data using TS0/E mount with "/" as the mountpoint.
Explanation If prompted for restart, fix the problem and respond. If no prompt, the system will run without the physical file system.		BPXF032D THE INDICATED FILE SYSTEM HAS BEEN TERMINATED. RETURN CODE = <i>return_code</i>, REASON CODE = <i>reason_code</i>
BPXF020I	FILE SYSTEM <i>name</i> MAY BE DAMAGED. RETURN CODE = <i>return_code</i>, REASON CODE = <i>reason_code</i>	Explanation Review all related physical file system diagnostic messages. Reply <i>#####</i> , R to restart the file system. Reply <i>#####</i> , I to ignore it.
Explanation Access to files within the file system might still be possible. If an sdump was captured, provide it to IBM support.		BPXF076I FILE SYSTEM INIT DELAYED BY ACTIVITY ON ANOTHER SYSTEM.
BPXF029E	ROOT FILE SYSTEM <i>name</i> WAS NOT MOUNTED. RETURN CODE = <i>return_code</i>, REASON CODE = <i>reason_code</i>	Explanation See the subsequent BPXF041I message in OPERLOG. Issue D GRS, LATCH, C on all systems to investigate latch contention.
		BPXF083I THE INDICATED FILE SYSTEM IS UNUSABLE UNTIL IT IS UNQUIESCED <<i>name</i>>QUIESCING SYSTEM=<<i>sysname</i>> PID=<<i>pid</i>> JOB=<<i>jobname</i>> LATCH=<<i>latchnum</i>>

Explanation

An authorized user might be able to unquiesce the file system from the ISPF shell. See the message documentation in [z/OS MVS System Messages, Vol 3 \(ASB-BPX\)](#).

**BPXF215E ACCESS TO THE z/OS UNIX
COUPLE DATA SET IS NOT
AVAILABLE.**

Explanation

Review the error codes and correct the access problem. Issue the SETXCF couple command to enable the couple data set.

**BPXF216E FILE SYSTEM PARTITION
CLEANUP IS DELAYED DUE TO *text***

Explanation

See the subsequent BPXF041I message in OPERLOG. Issue D GRS , LATCH , C on all systems to investigate latch contention.

**BPXF217E FILE SYSTEM PARTITION
CLEANUP FAILED DUE TO *text***

Explanation

Issue F BPXOINIT , FILESYS=D , EXCEPTION to identify impacted file systems. Unmount and remount if file system does not recover.

**BPXI026I THE ETCINIT JOB COULD NOT BE
STARTED. *system_call* RETURN
CODE *return_code* REASON CODE
*reason_code***

Explanation

Creation for process failed for /etc/init or /usr/sbin/init.

Examine the return and reason codes. See the message documentation in [z/OS MVS System Messages, Vol 3 \(ASB-BPX\)](#).

**BPXI027I THE ETCINIT JOB ENDED IN
ERROR, EXIT STATUS *exit_status***

Explanation

Examine the exit status code. See the message documentation in [z/OS MVS System Messages, Vol 3 \(ASB-BPX\)](#).

**BPXI031E BPXOINIT FAILED TO INITIALIZE.
RETURN CODE *return_code*
REASON CODE *reason_code***

Explanation

Examine the return and reason codes. Correct the error. A re-IPL of the system is required to start z/OS UNIX.

**BPXI036E UNIX SYSTEM SERVICES ARE NOT
AVAILABLE**

Explanation

Correct the conditions that caused the failure. A re-IPL of the system is required to start z/OS UNIX.

**BPXI043E MOUNT TABLE LIMIT HAS
REACHED *limperc%* OF ITS
CURRENT CAPACITY OF *limtot***

Explanation

See the message documentation in [z/OS MVS System Messages, Vol 3 \(ASB-BPX\)](#).

**BPXI060I *jobname* RUNNING IN ADDRESS
SPACE *asid* IS BLOCKING
SHUTDOWN OF OMVS**

Explanation

Evaluate stopping the indicated job. The indicated job must be stopped to allow z/OS UNIX shutdown to complete.

**BPXI062I *jobname* RUNNING IN ADDRESS
SPACE *asid* IS PREVENTING THE
SHUTDOWN OF OMVS FROM
COMPLETING**

Explanation

The indicated job is blocking z/OS UNIX shutdown. Evaluate stopping the indicated job. The indicated job must be stopped to allow z/OS UNIX shutdown to complete.

**BPXI068I *jobname* RUNNING IN ADDRESS
SPACE *asid* IS USING *text***

Explanation

Evaluate stopping the indicated job. The indicated job must be stopped to allow z/OS UNIX shutdown to complete

**BPXI076E LATCH CONTENTION EXISTS THAT
MUST BE RESOLVED PRIOR TO
SHUTDOWN**

Explanation

Issue D GRS, C to determine the nature of the contention. Evaluate canceling or forcing the address space that is causing contention.

BPXI084E OMVS SHUTDOWN IS STALLED IN FILE SYSTEM TERMINATION

Explanation

z/OS UNIX shutdown delayed while terminating file systems. Get an SVC dump of z/OS UNIX and its associated data spaces.

BPXM048I BPXOINIT FILESYSTEM SHUTDOWN INCOMPLETE. notshutdown FILESYSTEM(S) ARE STILL OWNED BY THIS SYSTEM. mounted FILESYSTEM(S) WERE

IEA message analysis

For more information on IEA messages for the communication and console services, see [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IEA230E WTOR BUFFER SHORTAGE. 80% FULL

Explanation

Enter D R, R to display outstanding WTOR messages. Reply to them. Enter K M, RLIM=nnnn to increase the RLIM buffer limit.

IEA231A SEVERE WTOR BUFFER SHORTAGE. 100% FULL

Explanation

Enter D R, R to display outstanding WTOR messages. Reply to them. Enter K M, RLIM=nnnn to increase the RLIM buffer limit.

IEA359E BUFFER SHORTAGE FOR RETAINED ACTION MESSAGES - 80% FULL

Explanation

Enter D R, L to display action messages, respond to them, or delete them using K C, X, ID-ID (X = A, E OR CE) where ID is the message numbers.

IEA360A SEVERE BUFFER SHORTAGE FOR RETAINED ACTION MESSAGES - 100% FULL

MOUNTED DURING THE SHUTDOWN PROCESS

Explanation

Issue D OMVS, F to identify file systems that did not move or unmount. See the message documentation in [z/OS MVS System Messages, Vol 3 \(ASB-BPX\)](#).

BPXP007E STARTING PHYSICAL FILE SYSTEM pfsname IN ADDRESS SPACE spacename

Explanation

Get an SVC dump of the indicated address space, z/OS UNIX and its associated data spaces. Provide the dump to IBM support.

Explanation

Enter D R, L to display action messages, respond to them, or delete them using K C, X, ID-ID (X = A, E OR CE) where ID is the message numbers.

IEA404A SEVERE WTOR BUFFER SHORTAGE - 100% FULL

Explanation

Enter D C, B to display WTO backlog; enter K Q to clear message queue; enter K M, MLIM=nnnn to increase the buffer limit for MLIM (maximum number of WTO messages allowed in the system).

IEA405E WTO BUFFER SHORTAGE - 80% FULL

Explanation

Enter D C, B to display WTO backlog. Enter K Q to clear message queue. Enter K M, MLIM=nnnn to increase the buffer limit for MLIM (maximum number of WTO messages allowed in the system).

IEA406I WTO BUFFER SHORTAGE RELIEVED

IEA611I {COMPLETE|PARTIAL} DUMP ON dsname

Explanation

Use IPCS to diagnose the problem and determine the action for a partial dump, see the message documentation.

IEA793A NO SVC DUMP DATA SETS AVAILABLE FOR DUMPID=*dumpid* FOR JOB (*MASTER*). USE THE DUMPDS COMMAND OR REPLY D TO DELETE THE CAPTURED DUMP

Explanation

Enter DD ADD ,SMS=class to add SMS classes, enter DD ADD ,VOL=VOLSER to add DASD volumes, or reply D to delete the captured dump.

IEA799I AUTOMATIC ALLOCATION OF SVC DUMP DATA SET FAILED DUMPID=*dumpid* REQUESTED BY JOB (*jobname*) *reason-text* *reason text2*

Explanation

Enter D D to view allocation status. Enter DD ADD ,VOL=VOLSER to add dump resources.

IEE message analysis

For more information about IEE messages for MVS, see [z/OS MVS System Messages, Vol 7 \(IEB-IEE\)](#).

IEE012A NO LONGER SAVING MESSAGES FOR HARDCOPY, LOGLIM REACHED.

Explanation

Issue V SYSLOG ,HARDCPY to activate SYSLOG. K M ,LOGLIM=nnnnn to increase the LOGLIM value.

IEE601E PROCESSOR (y) IS IN AN EXCESSIVE DISABLED SPIN LOOP WAITING FOR *event* HELD BY PROCESSOR (x). ACR IS ALREADY ACTIVE. SPIN WILL CONTINUE.

Explanation

Reply U , ABEND or TERM to message IEE331A.

IEE711I [SYSTEM UNABLE TO DUMP] SYSTEM DUMP NOT TAKEN. *reason*]

Explanation

The system did not write the requested SVC dump. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IEE766E BUFFER SHORTAGE FOR SYSTEM LOG - 60% FULL

IEA911E {COMPLETE|PARTIAL} DUMP ON SYS1.DUMPnn JOB (*jobname*) FOR ASIDS(*id,id,...*) [REMOTE DUMPS REQUESTED | REMOTE DUMP FOR SYSNAME: *sysname*] INCIDENT TOKEN:*incident-token* [SDRSN = *vvvvvvvv* *wwwwwwwww* *xxxxxxxxx* *zzzzzzzz*] [*reason-text*] [ERRORID = SEQyyyyyy CPUzz ASID*asid* TIME*hh.mm.ss.f*] [TSOID = *tsoid*] [ID = *uuuuuuuuuu*DUMP*id=dumpid* REQUESTED BY

Explanation

Use IPCS to diagnose the problem and determine the action. For a partial dump, see the message documentation.

Explanation

K M ,LOGLIM=nnnnnn to increase the WTL buffers.

IEE767A SEVERE BUFFER SHORTAGE FOR SYSTEM LOG - 100% FULL

Explanation

K M ,LOGLIM=nnnnnn to increase the WTL buffers.

IEE786I THE PAGEADD COMMAND ENCOUNTERED A FAILURE

Explanation

Issue R nn ,U to message IEE787A to continue PAGEADD processing. Issue R nn ,END to message IEE787A to end PAGEADD processing.

IEE986E SMF HAS USED *nn*% OF AVAILABLE BUFFER SPACE

Explanation

Issue D SMF to check the status of the SMF data sets. Use the SMF dump program (IFASMFDP) to make a data set available for use.

IOS message analysis

For more information about IOS messages, see [z/OS MVS System Messages, Vol 9 \(IGF-IWM\)](#).

IOS078I I/O REQUEST HAS TIMED OUT

Explanation

Run EREP to dump data from SYS1.LOGREC and provide it to IBM Support.

IOS079I I/O REQUEST HAS TIMED OUT

Explanation

Run EREP to dump data from SYS1.LOGREC and provide it to IBM Support.

IOS431I THE INDICATED SYSTEM HOLDS THE RESERVE ON THE DASD DEVICE.

Explanation

Issue D GRS,DEV=dev to investigate and resolve the contention. See the message documentation in [z/OS MVS System Messages, Vol 9 \(IGF-IWM\)](#).

IRA message analysis

For more information about IRA messages for System Resource Manager, see [z/OS MVS System Messages, Vol 9 \(IGF-IWM\)](#).

IRA100E SQA SHORTAGE

Explanation

Determine largest users of sqa/csa storage. Get an SVC dump of those users for analysis. Include SQA/CSA/RGN in the dump.

IRA101E CRITICAL SQA SHORTAGE

Explanation

Get an SVC dump of common storage to analyze its growth and determine if SQA/CSA allocation is adequate.

IRA103I SQA/ESQA HAS EXPANDED INTO CSA/ECSA BY xxxxx PAGES

Explanation

Evaluate the system requirement for sqa storage. Increase the size of the SQA parameter in the IEASYSXX parmlib member.

IRA200E AUXILIARY STORAGE SHORTAGE

IOS1078I I/O REQUEST HAS TIMED OUT

Explanation

Run EREP to dump data from SYS1.LOGREC and provide it to IBM support.

IOS1079I I/O REQUEST HAS TIMED OUT

Explanation

Run EREP to dump data from SYS1.LOGREC and provide it to IBM support.

Explanation

Issue PA PAGE=DSNAME to add auxiliary storage. Evaluate canceling jobs indicated by messages IRA206I and IRA210E.

IRA201E CRITICAL AUXILIARY STORAGE SHORTAGE

Explanation

Issue PA PAGE=DSNAME to add auxiliary storage. Evaluate canceling any jobs indicated by messages ira206i and ira210e.

IRA206I uuuuuuuu ASID aaaa FRAMES ffffff SLOTS ssssssss % OF AUX nn.n

Explanation

Get an SVC dump of the indicated ASIDs for analysis. Evaluate canceling one or more of the indicated ASIDs.

IRA210E uuuuuuuu ASID aaaa SET NON DISPATCHABLE Frames+Slots vvvvvvvvv RATE rrrrrr

Explanation

Issue PA PAGE=DSNAME to add auxiliary storage.
Evaluate canceling the indicated job.

IRA211I *uuuuuuuu ASID aaaa SET NON
DISPATCHABLE Frames+Slots
vvvvvvvvv RATE rrrrrr*

Explanation

Issue PA PAGE=DSNAME to add auxiliary storage.
Evaluate canceling the indicated job.

IRA220I **CRITICAL AUXILIARY
SHORTAGE***text*

Explanation

The *text* is:

```
! ## ! USER      ! ASID ! PAGES      !  
SLOTS          !  
+-----+-----+-----+  
+-----+  
! ii ! uuuuuuuu ! aaaa ! xxxxxxxxxxxx !
```

IXC message analysis

For more information on IXC messages for Cross System Coupling Facility (XCF) , see [z/OS MVS System Messages, Vol 10 \(IXC-IZP\)](#).

IXC101I **SYSPLEX PARTITIONING IN
PROGRESS FOR *sysname*
REQUESTED BY *jobname* REASON:
*reason***

Explanation

See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXC102A **THE INDICATED SYSTEM IS
WAITING FOR SYSTEM RESET.
REASON: *reason***

Explanation:

System reset the indicated system before replying down. See the message documentation.

IXC105I **A SYSTEM HAS BEEN REMOVED
FROM THE SYSPLEX.**

Explanation

See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXC244E **XCF CANNOT USE {PRIMARY |
ALTERNATE} SYSPLEX COUPLE
DATA SET *dsname*, {ON VOLSER
volser, | VOLSER N/A,} *text***

```
zzzzzzzzzz !  
! ii ! uuuuuuuu ! aaaaS! xxxxxxxxxxxx !  
zzzzzzzzzz !  
! ii ! uuuuuuuu ! aaaaN! xxxxxxxxxxxx !  
zzzzzzzzzz !
```

Issue PA PAGE=DSNAME to add auxiliary storage.
Evaluate canceling a specific consumer by replying to message IRA221D.

IRA400E *return-code, PAGEABLE STORAGE
SHORTAGE*

Explanation

Evaluate canceling jobs identified by messages IRA403E and IRA404I.

IRA401E *return-code, CRITICAL PAGEABLE
STORAGE SHORTAGE*

Explanation

Evaluate canceling jobs identified by messages IRA403E and IRA404I.

Explanation

Check data set name and VOLSER to ensure they are correct.

IXC246E *typename* **COUPLE DATA SET
dsname ON VOLSER *volser*,
DEVN*devnum*, HAS BEEN
EXPERIENCING I/O DELAYS FOR
delaysec SECONDS.**

Explanation

See messages IOS078I and IOS1078I for the same device number and see the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXC255I **UNABLE TO USE DATA SET *dsname*
AS THE {PRIMARY|ALTERNATE}
FOR *typename*: *text*
[RELEVANT*typename* COUPLE
DATA SET FORMAT
INFORMATIONPRIMARY FORMAT
LEVEL: *fmtlevel* FORMAT
KEYWORDS: *fmtinfo* ALTERNATE
FORMAT LEVEL: *fmtlevel* FORMAT
KEYWORDS: *fmtinfo*]**

Explanation

See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXC256A **REMOVAL OF {PRIMARY|ALTERNATE} COUPLE DATA SET *dsname* FOR *typename* CANNOT FINISH THE {ACTION|COMPLETE} PHASE UNTIL THE FOLLOWING SYSTEM(S) ACKNOWLEDGE THE REMOVAL: *syslist***

Explanation

Missing response delaying the removal of a couple data set. If the condition persists, see the message documentation in *z/OS MVS System Messages, Vol 6 (GOS-IEA)*.

IXC259I **I/O ERROR ON DATA SET *dsname* FOR *typename*, VOLSER *volser*, *modname*, *post-code*, *text***

Explanation

Reply COUPLE=xx to message IXC207A. (xx = the suffix of the COUPLExx parmlib member to be used by XCF initialization.)

IXC267E **PROCESSING WITHOUT AN ALTERNATE COUPLE DATA SET FOR *typename*. ISSUE SETXCF COMMAND TO ACTIVATE A NEW ALTERNATE.**

Explanation

No alternate couple data set defined for indicated function. Issue SETXCF COUPLE, TYPE=*typename*, ACOUPLE=DSNAME to activate a new alternate couple data set.

IXC402D **THE INDICATED SYSTEM MIGHT BE INOPERATIVE.**

Explanation

Determine why the system is inoperative. Reply INTERVAL=SSSS for more time to resolve or reset the system then reply *###*, DOWN.

IXC406I **THIS SYSTEM IS CONNECTED TO ETR NET ID=*xx*. THE OTHER ACTIVE SYSTEMS IN THE SYSPLEX ARE USING ETR NET ID=*yy*. EFFECTIVE CLOCK VALUES ARE NOT CONSISTENT.**

Explanation

Correct any ETR problem and retry with COUPLExx parmlib member or correct any improperly defined ETR time offsets.

IXC409D **UNABLE TO COMMUNICATE BETWEEN INDICATED SYSTEMS.**

Explanation

Investigate signaling connectivity. Attempt to resolve. If not possible, specify retry or sysname to remove.

IXC427A **SYSTEM *sysname* HAS NOT UPDATED STATUS SINCE *hh:mm:ss* BUT IS SENDING XCF SIGNALS. XCF SYSPLEX FAILURE MANAGEMENT WILL REMOVE *sysname* IF NO SIGNALS ARE RECEIVED WITHIN A *interval* SECOND INTERVAL**

Explanation

Determine and resolve any contention issues with the sysplex couple data set. Reply to message IXC426D.

IXC430E **SYSTEM *sysname* HAS STALLED XCF GROUP MEMBERS**

Explanation

Issue D XCF,G to see groups with stalled members. See the message documentation in *z/OS MVS System Messages, Vol 6 (GOS-IEA)*.

IXC431I *text*

Explanation

Issue D XCF, G, *grpname*, *membername* for more information. For the complete message text, see IXC431I in *z/OS MVS System Messages, Vol 10 (IXC-IZP)*

IXC432I *text*

Explanation

The indicated XCF group member is no longer stalled. No action needed.

IXC440E **SYSTEM *hurtsys* IMPACTED BY STALLED XCF GROUP MEMBERS ON SYSTEM *stallsys***

Explanation

Issue D XCF, G to see groups with stalled members. See the message documentation in *z/OS MVS System Messages, Vol 6 (GOS-IEA)*.

IXC446I **SYSTEM *sysname* IS IN MONITOR-DETECTED STOP STATUS BUT IS SENDING XCF SIGNALS. SFM WILL TAKE SSUM ACTION AT**

actiontime IF SYSTEM REMAINS IN THIS STATE.

Explanation

The indicated system has not updated its system status. Investigate the XCF couple data sets for contention or poor performance.

IXC467I *command dir pathname RSN: text*

Explanation

XCF IS TRYING TO RESTORE THE INDICATED PATH TO SERVICE. SEE THE MESSAGE DOCUMENTATION.

Explanation

The indicated system is waiting for system reset. See the message documentation in *z/OS MVS System Messages, Vol 6 (GOS-IEA)*.

IXC500I **THE COUPLING FACILITY MIGHT BE OWNED BY ANOTHER SYSPLEX.**

Explanation

Reply to message IXC501A. Reply Y to use named coupling facility or reply N to not use it.

IXC512I **POLICY CHANGE IN PROGRESS FOR CFRM TO MAKE *polname* POLICY ACTIVE. *numpend* POLICY CHANGE(S) PENDING.**

Explanation

Issue D XCF,STR to show structures pending actions. See the message documentation in *z/OS MVS System Messages, Vol 6 (GOS-IEA)*.

IXC518I **SYSTEM *sysname* NOT USING COUPLING FACILITY *type.mfg.plant.sequence* PARTITION: *partition side* CPCID: *cpcid* NAMED *cfname* REASON: *text***

Explanation

The indicated coupling facility is unusable. See the message documentation in *z/OS MVS System Messages, Vol 6 (GOS-IEA)*.

IXC519E **COUPLING FACILITY DAMAGE RECOGNIZED FOR COUPLING FACILITY *type.mfg.plant.sequence* PARTITION: *partition side* CPCID: *cpcid* NAMED *cfname***

Explanation

Run EREP to dump data from SYS1.LOGREC and gather XCF/XES CTRACE records and provide them to IBM support.

IXC522I ***rebuildtype* FOR STRUCTURE *strname* IS BEING STOPPED *action* DUE TO *reason* [*codetype stopcode*]**

Explanation

See the message documentation in *z/OS MVS System Messages, Vol 6 (GOS-IEA)*.

IXC538I **DUPLEXING REBUILD OF STRUCTURE *strname* WAS NOT INITIATED BY MVS. REASON: *reason***

Explanation

Duplexing rebuild start request of structure not performed. See the message documentation in *z/OS MVS System Messages, Vol 6 (GOS-IEA)*.

IXC552I **DUPLEX REBUILD NEW STRUCTURE *strname* WAS ALLOCATED IN A COUPLING FACILITY THAT IS NOT FAILURE ISOLATED FROM THE OLD STRUCTURE.**

Explanation

Review your coupling facility configuration. When possible duplex failure isolation is strongly encouraged.

IXC553E **DUPLEXING REBUILD NEW STRUCTURE *strname* IS NOT FAILURE ISOLATED FROM THE DUPLEXING REBUILD OLD STRUCTURE.**

Explanation

Review your coupling facility configuration. When possible duplex failure isolation is strongly encouraged.

IXC573I ***phase* PROCESSING DURING A SYSTEM-MANAGED *process* FOR STRUCTURE *strname* ENCOUNTERED AN ERROR. ERROR DATA: *reason* [*reldata1 reldata2 reldata3 reldata4*] AUTO VERSION: *procid1 procid2***

Explanation

Error encountered during rebuild or duplex rebuild process. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXC585E **STRUCTURE *strname* IN COUPLING FACILITY *cfname*, PHYSICAL STRUCTURE VERSION *physver1 physver2*, IS AT OR ABOVE STRUCTURE FULL MONITORING THRESHOLD OF *thresh*%. ENTRIES: IN-USE: *nnnnnnnn*, TOTAL: *pppppppp*, *pct*% FULL [ELEMENTS: IN-USE: *nnnnnnnn*, TOTAL: *pppppppp*, *pct*% FULL] [EMCS: IN-USE: *nnnnnnnn*, TOTAL: *pppppppp*, *pct*% FULL]**

Explanation

Issue D XCF, STR, STRNAME=*strname* to get structure information. Increase structure size or take action against application.

IXC615I **GROUP *grpname* MEMBER *membername* JOB *jobname* ASID *asidtext***

Explanation

XCF terminated the group member to resolve critical problem. Restart the affected application, subsystem, or system. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXC631I **GROUP *grpname* MEMBER *membername* JOB *jobname* ASID *asid* STALLED, IMPACTING SYSTEM *sysname* {WHICH IS IN PARTITIONING }**

Explanation

Issue D XCF, G, *grpname*, *membername* to get more information. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXG message analysis

For more information on IXG messages for system logger, see [z/OS MVS System Messages, Vol 10 \(IXC-IZP\)](#).

IXG051I **SYSTEM LOGGER ADDRESS SPACE CREATE FAILED**

IXC633I ***text* GROUP *gnme* MEMBER *mnme* JOB *jnme* ASID *asid* {DEEMED | CONFIRMED} IMPAIRED AT *ipdate iptime* ID: *s#r#* LAST MSGX: *sgdate sgtime sgexit* STALLED *sgwork* PENDINGQ LAST GRPX: *grdate grtime grexit* STALLED *grwork* PENDINGQ LAST STAX: *stdate sttime stexit* STALLED**

Explanation

Indicated XCF group member is not operating normally. See message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXC635E **SYSTEM *sysname* HAS IMPAIRED XCF GROUP MEMBERS**

Explanation

The indicated system has impaired group members. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXC700E **SYSPLEX COUPLE DATA SET LIMIT REACHED, FUTURE REQUESTS MAY BE REJECTED. *text***

Explanation

Issue D XCF, COUPLE to review maximum values in the sysplex couple data set. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXC800I **ELEMENTS FROM TERMINATED SYSTEM *sysname* NOT RESTARTED. *text***

Explanation

Indicated system had jobs that could not be restarted. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXG055I **SYSTEM LOGGER WILL NO LONGER ISSUE SMF RECORDS**

Explanation

The system logger address space create (IXGLOGR) failed. Search the problem reporting databases for a fix. If no fix exists, contact IBM support.

Explanation

System logger SMF record handling unrecoverable task abend. Search the problem reporting databases for a fix. If no fix exists, provide the dump to IBM support. See "Restarting the system logger address space" in *z/OS MVS Diagnosis: Reference*.

IXG056I SYSTEM LOGGER ADDRESS SPACE HAS ENDED*text*

Explanation

The system logger address space has ended. If the system logger address space (IXGLOGR) is not running, find message IXG056I and follow the message documentation in *z/OS MVS System Messages, Vol 10 (IXC-IZP)*. Start IXGLOGRS to restart the system logger address space.

In the message text:

**IXGLOGR IS AUTOMATICALLY RESTARTING.
IXGLOGR CANNOT BE AUTOMATICALLY RESTARTED.
OPERATOR ISSUED FORCE COMMAND. MANUAL RESTART REQUIRED.**

IXG057I SYSTEM LOGGER ADDRESS SPACE INITIALIZATION FAILED

Explanation

The IXGLOGR address space has failed to initialize. Start IXGLOGRS to restart the system logger address space. If the problem persists, provide the diagnostic information in the message to IBM support.

IXG058E LOGSTREAM CONNECTIONS HAVE BEEN LOST

Explanation

Log stream exploiters lost their log stream connections. If the system logger address space (IXGLOGR) is not running, Find message IXG056I and follow the message documentation in *z/OS MVS System Messages, Vol 10 (IXC-IZP)*. Start IXGLOGRS to restart the system logger address space.

IXG114A OFFLOAD IS NOT PROGRESSING ON *sysname*

Explanation

Log stream recovery is stalled on the indicated system. Check for any conditions in the sysplex that could prevent offload from progressing. See the message documentation for this message and message IXG115A in *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

IXG121I STALLED ASYNCHRONOUS IXGWRITE REQUESTS DETECTED DURING REBUILD FOR STRUCTURE *strname*. REBUILD ALLOWED TO CONTINUE

Explanation

Stalled log stream writes delayed structure rebuild progress. Search the problem reporting databases for a fix. If no fix exists, get an SVC dump of the IXGLOGR address space, SYSLOGR CTRACE buffer, and CF structure and provide it to IBM support.

IXG235I STALLED ASYNCHRONOUS IXGWRITE REQUESTS DETECTED DURING DISCONNECT FOR LOGSTREAM *streamname*. DISCONNECT ALLOWED TO CONTINUE

Explanation

Stalled log stream writes delayed disconnect progress. Search the problem reporting databases for a fix. If no fix exists, get an SVC dump of the IXGLOGR address space, SYSLOGR CTRACE buffer, and CF structure and provide it to IBM support.

IXG257I DATA SET DIRECTORY FOR LOGSTREAM IS OVER 90% FULL

Explanation

Indicated log stream directory is near maximum capacity. Take action to cause unneeded offload data sets to be deleted or format a new LOGR CDS using the DSEXTENT item and define it as the alternate, then pswitch it as the primary LOGR CDS.

IXG261E SHORTAGE OF DIRECTORY EXTENT RECORDS

Explanation:

DSEXTENT Record shortage can impact log stream offloads. Relieve the DSEXTENT record shortage. See the topic on increasing the space available for DASD log data sets in *z/OS MVS Setting Up a Sysplex*.

IXG262A CRITICAL SHORTAGE OF DIRECTORY EXTENT RECORDS

Explanation:

DSEXTENT Record shortage will impact log stream offloads. Relieve the DSEXTENT record shortage. See the topic on increasing the space available for DASD log data sets in *z/OS MVS Setting Up a Sysplex*.

IXG263E IXGLOGR ATTEMPT TO ALLOCATE A DATA SET FAILED FOR

**LOGSTREAM log stream.
UNSUPPORTED CFSIZE | DSSIZE**

Explanation:

Unsupported data set was unable to be used by system logger. For CFSIZE, correct the SMS data class and *_DATACLAS parms for the log stream. For DSSIZE, see the topic on testing log data set modifications in *z/OS MVS Setting Up a Sysplex*.

**IXG266I LOGGER DATA SET DELETION
ACTIVITY FOR LOGSTREAM log
stream WAS QUIESCED**

Explanation:

Log stream data set deletions timed out. Processing delayed. Check to see if any remaining data sets that might still need to be deleted are eventually deleted on subsequent log stream offloads.

**IXG267I ALLOCATED DATA SET WITH
INCORRECT VSAM
SHAREOPTIONS, BUT IS
ACCEPTED FOR USE.**

Explanation:

Indicated data set allocated with wrong VSAM share options. Use IDCAMS to change the VSAM share options for the data set. See message documentation in *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

**IXG268I DATASET CAN NOT BE OPENED,
UNABLE TO USE IT**

Explanation:

Indicated error prevents logger access to indicated data set. If the VSAM shareoptions are not correct, change them using IDCAMS. See message documentation in *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

**IXG270I SHORTAGE OF DIRECTORY
EXTENT RECORDS FOR GROUP:
TEST**

Explanation:

DSEXTENT records shortage can lead to group(test) failures. Relieve the dsextent record shortage. See the topic on increasing the space available for DASD log data sets in *z/OS MVS Setting Up a Sysplex*.

**IXG271I LOGGER DATA SET REQUEST IN
group taskname SERVICE TASK
DELAYED**

Explanation:

Indicated service task timed out while processing log stream. See message documentation for this message and message IXG272E in *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

**IXG275I SYSTEM LOGGER EVENT
MONITORING STOPPED FOR**

**SERVICE TASKS ON SYSTEM
sysname**

Explanation:

Logger monitoring stopped for tasks on indicated system. Make sure that logger is still functioning normally. If task monitoring is essential, see the topic on restarting the system logger address space in *z/OS MVS Diagnosis: Reference*.

**IXG285I LOGGER OFFLOAD DATA SET
MANAGEMENT SUBTASK HAS
FAILED**

Explanation:

Logger offload subtask has failed for indicated log stream. Look for messages IXG286I OR IXG288E. See the message documentation for those messages in *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

**IXG286I LOGGER IS RUNNING IN A
DIMINISHED CAPACITY LOSS OF
OFFLOAD SUBTASKS**

Explanation:

System logger lost multiple offload tasks for indicated log stream. Restart the IXGLOGR address space as soon as possible. See the topic on restarting the system logger address space in *z/OS MVS Diagnosis: Reference*.

**IXG287E LOGGER NO LONGER DELETING
OFFLOAD DATA SETS**

Explanation:

Logger data set deletions failed for the indicated log stream. Restart the IXGLOGR address space as soon as possible. See the topic on restarting the system logger address space in *z/OS MVS Diagnosis: Reference*.

**IXG288E CRITICAL FAILURE OF LOGGER
OFFLOAD DATA SET MANAGEMENT**

Explanation:

Logger offload processing no longer being done on system. Restart the IXGLOGR address space as soon as possible. See the topic on restarting the system logger address space in *z/OS MVS Diagnosis: Reference*.

**IXG289E CRITICAL FAILURE OF LOGGER
DATA SET ALLOCATION
SUBCOMPONENT**

Explanation

Restart the IXGLOGR address space as soon as possible. See the topic on restarting the system logger address space in *z/OS MVS Diagnosis: Reference*.

**IXG310I SYSTEM LOGGER CURRENT
OFFLOAD IS NOT PROGRESSING**

Explanation

Offload is delayed for the indicated log stream. Look for related IXG311I message, and if located, take action indicated for that message.

**IXG311I SYSTEM LOGGER CURRENT
 OFFLOAD HAS NOT PROGRESSED**

Explanation

Logger offload has timed out for the indicated log stream. Issue command D R,R to display outstanding WTOR messages. Reply to message IXG312E if displayed. See the message documentation for message IXG312E in *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

**IXG313I DELAYED OFFLOAD FOR
 LOGSTREAM *log stream*
 CANCELLED BY INSTALLATION
 POLICY.**

Explanation

Offload for indicated log stream timed out and was canceled. Make sure the offload completes for the log stream. Look for message IXG304I for the indicated log stream.

**IXG314I SYSTEM LOGGER EVENT
 MONITORING STOPPED FOR
 LOGSTREAM OFFLOADS ON
 SYSTEM *sysname***

Explanation

Logger offload event monitoring stopped on indicated system. Make sure that logger is still functioning normally. If offload monitoring is essential, see the topic on restarting the system logger address space in *z/OS MVS Diagnosis: Reference*.

**IXG318E LOGSTREAM PRIMARY STORAGE
 CONSUMPTION HAS REACHED
 100% IN USE**

IXL message analysis

For more information on IXL messages for Cross System Extended Services (XES), see *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

**IXL010E NOTIFICATION RECEIVED FROM
 COUPLING FACILITY
 type.mfg.plant.sequence
 PARTITION: *partition side* CPCID:
 cpcid NAMED *cfname cfservrecord***

Explanation

Log stream resources have become full, exploiter impacted. Look for messages IXG310I, IXG311I, IXG312E, and IXG324I. See the message documentation for those messages in *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

**IXG324I DATA SET SWITCH FAILURE
 PREVENTING OFFLOAD
 PROGRESS FOR LOGSTREAM *log
stream*, IN STRUCTURE *strname***

Explanation

Offload data set failure preventing progress. Identify and correct the data set failure conditions. See message documentation in *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

**IXG701I NO TRACING IN EFFECT FOR
 SYSTEM LOGGER - SYSLOGR
 COMPONENT TRACE DEFINE
 FAILED.**

Explanation

CTRACE define ended with a failing return code. Check related component trace messages (prefix ITT) and search problem reporting databases for a fix. If no fix exists, contact IBM support.

**IXG733I LOGGER PARAMETER
 PROCESSING UNAVAILABLE**

Explanation

System logger parameter processing failure occurred. The system is using logger default parameters. If this is not acceptable, see the topic on restarting the system logger address space in *z/OS MVS Diagnosis: Reference*.

Explanation

Run EREP to dump data from SYS1.LOGREC and provide it to IBM support.

**IXL013I *requesttype* REQUEST FOR
 STRUCTURE *structure-name*
 FAILED. JOBNAME: *jobname* ASID:
 asid CONNECTOR NAME:
 connector-name IXLCONN RETURN**

CODE: *return-code*, **REASON CODE:** *reason-code* *errortype*
CONADIAGdiag: *diagvalue*

Explanation

See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

Note: Only the last occurrence of message IXL013I is displayed when structure name, jobname, ASID, and connector name match.

IXL030I **CONNECTOR STATISTICS FOR LOCK STRUCTURE** *structure-name*,
CONNECTOR *connector-name: n*

Explanation

Connector statistical information for indicated connector. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

Note: Only the last occurrence of message IXL030I is displayed when structure name and connector name match.

IXL040E **CONNECTOR NAME:** *connector-name*, **JOBNAME:** *jobname*, **ASID:** *asid* **HAS** *text*. **process** **FOR STRUCTURE** *structure-name* **CANNOT CONTINUE. MONITORING FOR RESPONSE STARTED:** *mm/dd/yyyy hh:mm:ss*. **DIAG:** *x*

Explanation

Evaluate job and gather XES CTRACE for specific connector terminate non-responsive job. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXL041E **CONNECTOR NAME:** *connector-name*, **JOBNAME:** *jobname*, **ASID:** *asid* **HAS NOT RESPONDED TO THE event** **FOR SUBJECT CONNECTION:** *subject-connector-name*. **process** **FOR STRUCTURE** *structure-name* **CANNOT CONTINUE. MONITORING FOR RESPONSE STARTED:** *mm/dd/yyyy hh:ss:mm*. **DIAG:** *x*

Explanation

Evaluate job and gather XES CTRACE for specific connector; terminate non-responsive job. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXL044I **COUPLING FACILITY** *cfname* **HAS EXPERIENCED** *IfccCount* **INTERFACE CONTROL CHECKS ON CHPID** *chpid* **DURING THE LAST interval** **SECONDS.**

Explanation

Run EREP to dump data from SYS1.LOGREC and provide it to IBM support.

IXL045E **[REBUILD] CONNECTOR NAME:** *connector-name*, **JOBNAME:** *jobname*, **ASID:** *asid* **FOR STRUCTURE** *structure-name* **MAY BE ENCOUNTERING DELAYS DUE TO LIMITED XES SRB SCHEDULING.**

Explanation

Run EREP to dump data from SYS1.LOGREC and gather XES CTRACE and system log and provide it to IBM support. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXL049E **HANG RESOLUTION ACTION FOR CONNECTOR NAME:** *conname* **TO STRUCTURE** *strname*, **JOBNAME:** *jobname*, **ASID:** *asid*: *actiontext*

Explanation

See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXL158I **PATH** *chpid* **IS NOW NOT-OPERATIONAL TO CUID:** *cuid* **COUPLING FACILITY** *type.mfg.plant.sequence* **PARTITION:** *partition* **side** **CPCID:** *cpcid*

Explanation

Indicated channel not operational. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXL159E **COUPLING SUPPORT FACILITY HARDWARE ERROR DETECTED.**

Explanation

Run EREP to dump data from SYS1.LOGREC and provide it to IBM Support.

IXL160E **CF REQUEST TIME ORDERING: REQUIRED AND NOT-ENABLED** *text* *reason*

Explanation

Coupling facility hardware configuration problem. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

Explanation

Coupling facility hardware configuration problem. See the message documentation in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#).

IXL162E

**CF REQUEST TIME ORDERING:
REQUIRED AND WILL NOT BE
ENABLED** *text reason*

Test messages ignored by Runtime Diagnostics

The Geographically Dispersed Parallel Sysplex® (GDPS®) integrates Parallel Sysplex technology and remote copy technology to enhance application availability and improve disaster recovery. Runtime Diagnostics ignores the following messages when GDPS issues them as test messages:

- IXC102A
- IXC105I
- IXC256A

Part 3. Predictive Failure Analysis

Soft failures are abnormal yet allowable behaviors that can slowly lead to the degradation of the operating system. To help eliminate soft failures, use Predictive Failure Analysis (PFA). PFA is intended to detect abnormal behavior early enough to allow you to correct the problem before it affects your business. PFA uses remote checks from IBM Health Checker for z/OS to collect data about your installation, and then uses machine learning to analyze this historical data to identify abnormal behavior. It warns you by issuing an exception message when a system trend might cause a problem. To help you correct the problem, it identifies a list of potential issues. PFA can invoke Runtime Diagnostics to analyze and report insufficient metric activity for specific checks and provide the next action that you can take to avoid a problem.

Chapter 7. Predictive Failure Analysis overview and installation

This chapter contains the following information:

- [“Avoiding soft failures” on page 73](#)
- [“Overview of Predictive Failure Analysis” on page 73](#)
 - [“How PFA works with a typical remote check ” on page 74](#)
 - [“How PFA interacts with IBM Health Checker for z/OS” on page 75](#)
 - [“How PFA invokes Runtime Diagnostics” on page 75](#)
 - [“Migration considerations for PFA” on page 76](#)
 - [“How PFA uses the ini file” on page 76](#)
- [“Installing PFA” on page 77](#)
 - [“Installing PFA in a z/OS UNIX shared file system environment” on page 80](#)
 - [“Updating the Java path” on page 82](#)

Avoiding soft failures

Unlike typical problems or hard failures that have a clear start and a clear cause, soft failures are caused by abnormal, but allowable behavior. Because the cause of the problem is dependent on a certain sequence or combination of events that are unique and infrequent, a solution is often difficult to determine. Multiple atypical, but legal actions performed by components on the z/OS image cause most soft failures. By design, most components of z/OS are stateless and are therefore unable to detect soft failures caused by atypical behavior.

A classic example is the exhaustion of common storage usage. A low priority, authorized task obtains common storage, but obtains significantly more common storage than usual. Then, a critical authorized system component fails while attempting to obtain a normal amount of common storage. Although the problem occurs in the second critical component, this second component is actually the victim. The first component caused the problem and is considered the villain. Soft failures usually occur in four generic areas:

- Exhaustion of shared resources
- Recurring or recursive failures often caused by damage to critical control structures
- Serialization problems such as classic deadlocks and priority inversions
- Unexpected state transition

z/OS has developed Predictive Failure Analysis (PFA) to help eliminate these soft failures.

Overview of Predictive Failure Analysis

Predictive Failure Analysis (PFA) is designed to predict potential problems with your systems. PFA extends availability by going beyond failure detection to predict problems before they occur. PFA provides this support using remote checks from IBM Health Checker for z/OS to collect data about your installation. Using this data, PFA constructs a model of the expected or future behavior of the z/OS images, compares the actual behavior with the expected behavior, and if the behavior is abnormal, PFA issues a health check exception. PFA uses a z/OS UNIX System Services (z/OS UNIX) file system to manage the historical and problem data that it collects.

Here is an LPAR view of the PFA components:

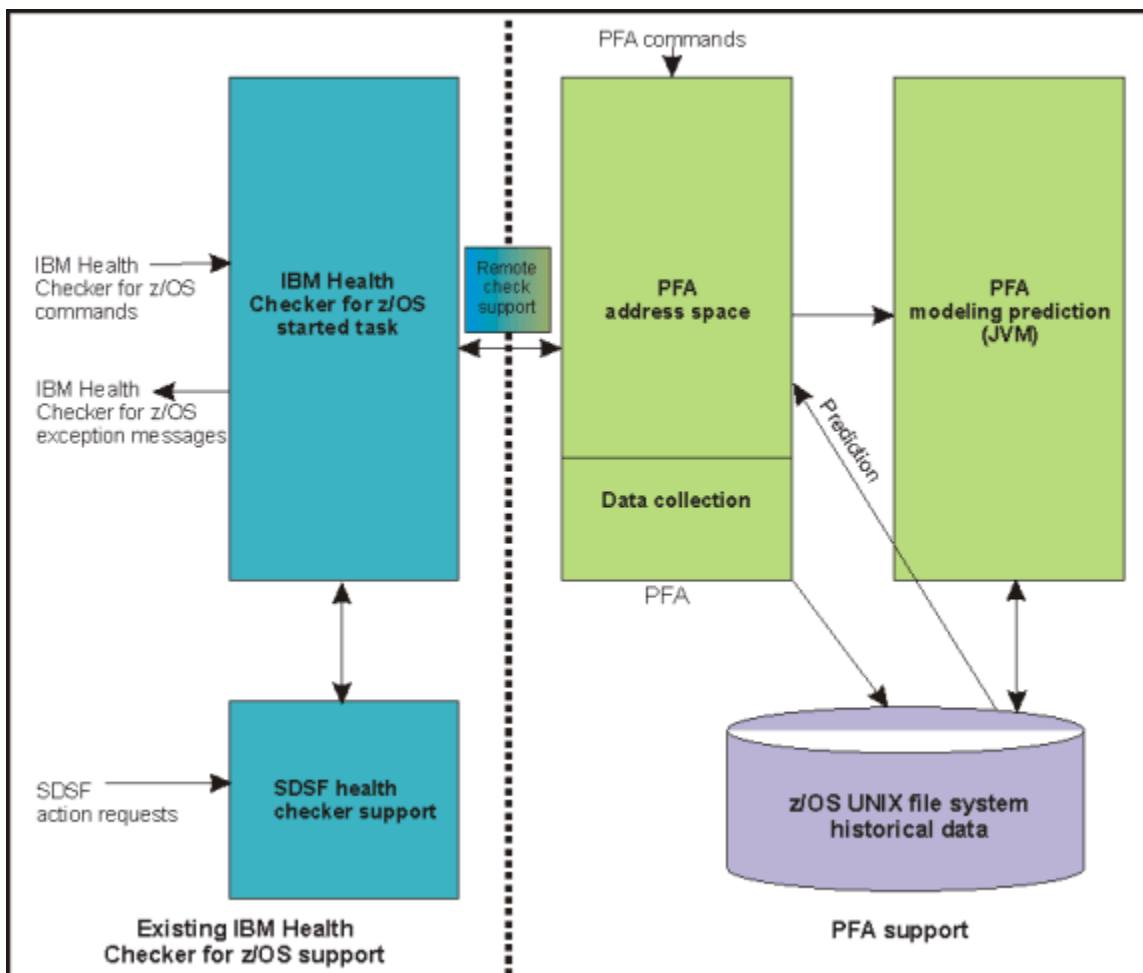


Figure 20. LPAR view of the PFA components

PFA creates report output in the following ways:

- In a z/OS UNIX file that stores the list of suspect tasks. The individual checks contain descriptions of the directory and file names.
- In an IBM Health Checker for z/OS report that is displayed by z/OS System Display and Search Facility (SDSF) and the message buffer.
- Your installation can also set up IBM Health Checker for z/OS to send output to a log stream. After you set it up, you can use the HZSPRINT utility to view PFA check output in the message buffer or in the log stream. For complete details, see [Using the HZSPRINT utility in *IBM Health Checker for z/OS User's Guide*](#).

How PFA works with a typical remote check

PFA_COMMON_STORAGE_USAGE is a remote check that evaluates the common storage use of each system. PFA, running in its own address space, periodically collects common storage area (CSA + SQA) data from the system on which the check is running. The check writes the CSA usage data, at intervals, to a z/OS UNIX file. The check identifies a list of common storage users that are abnormal and that might contribute to exhausting common storage. PFA issues an exception message to alert you if there is a potential common storage problem and provides a list of suspect tasks. You can then examine the list and stop the cause of the potential problem or move critical work off the LPAR.

How PFA interacts with IBM Health Checker for z/OS

When PFA issues an exception, the PFA check does not continue to issue exceptions to the console until the check determines a new exception must be issued or the exception resolves. For some checks, the new exception is always issued after a new model occurs. For other checks, the data must change significantly or the exception message must be different. For all checks, the check continues to run at the defined interval making the latest exception report data available using the CK panel in SDSF.

How PFA invokes Runtime Diagnostics

Runtime Diagnostics is an MVS utility (component HZR) that can perform some of the same tasks you might manually perform when looking for the cause of a hung address space as well as other tasks. See Chapter 5, “Runtime Diagnostics,” on page 35 for complete details about Runtime Diagnostics.

PFA can invoke Runtime Diagnostics to analyze and report insufficient metric activity from the PFA_ENQUEUE_REQUEST_RATE check, PFA_MESSAGE_ARRIVAL_RATE check, and PFA_SMF_ARRIVAL_RATE check. For details and examples, see:

- “PFA_ENQUEUE_REQUEST_RATE” on page 108
- “PFA_MESSAGE_ARRIVAL_RATE” on page 133
- “PFA_SMF_ARRIVAL_RATE” on page 163.

Note: PFA requires the Runtime Diagnostic address space (HZR) to be active on the system or systems running these checks for Runtime Diagnostics to detect the insufficient metric activity.

When PFA issues a check exception because metric activity is unusually low, the IBM Health Checker for z/OS report includes information from Runtime Diagnostics. The Runtime Diagnostics information in the report points to the specific job or address space and provides the next action you can take. The additional Runtime Diagnostic output can help you quickly determine your next course of action and possibly help you avoid additional problems.

The following is an example of the Runtime Diagnostics output that might appear in the message arrival rate check when PFA determined the tracked jobs had a lower than expected message arrival rate (for AIH206E):

```

.
.
.
Persistent address spaces with low rates:
                                     Predicted Message
                                     Arrival Rate
Job      Message
Name     ASID      Arrival
-----  -
JOBS4   0027      1.17      23.88      22.82      15.82
JOBS5   002D      0.30      8.34       11.11      12.11

```

Runtime Diagnostics Output:

```

Runtime Diagnostics detected a problem in ASID: 0027

EVENT 06: HIGH - HIGHCPU - SYSTEM: SY1 2011/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 0027 JOBNAME: JOBS4
STEPNAME: STEPA PROCSTEP: STEPA JOBID: STC00042 USERID: ++++++++
JOBSTART: 2011/06/12 - 13:28:35
ERROR: ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 07: HIGH - LOOP - SYSTEM: SY1 2011/06/12 - 13:28:46
ASID: 0027 JOBNAME: JOBS4 TCB: 004E6850
STEPNAME: STEPA PROCSTEP: STEPA JOBID: STC00042 USERID: ++++++++
JOBSTART: 2011/06/12 - 13:28:35
ERROR: ADDRESS SPACE APPEARS TO BE IN A LOOP.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.

```

Figure 21. Runtime Diagnostics report within the PFA message arrival rate check

Predictive Failure Analysis (PFA)

The following is an example of the Runtime Diagnostics output that might appear in the SMF arrival rate check when PFA determines the tracked jobs exception report for jobs that had a lower than expected SMF arrival rate (for AIH208E):

```
.
.
.
Persistent address spaces with low rates:
```

Job Name	ASID	SMF Arrival Rate	Predicted SMF Arrival Rate		
			1 Hour	24 Hour	7 Day
TRACKED4	0027	0.20	23.88	22.82	15.82
TRACKED5	0034	0.01	12.43	11.11	8.36

Runtime Diagnostics Output:

Runtime Diagnostics detected a problem in ASID: 0027

EVENT 06: HIGH - HIGHCPU - SYSTEM: SY1 2011/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 0027 JOBNAME: TRACKED4
STEPNAME: STEPA PROCSTEP: STEPA JOBID: STC00042 USERID: +++++++
JOBSTART: 2011/06/12 - 13:28:35
ERROR: ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.

EVENT 07: HIGH - LOOP - SYSTEM: SY1 2011/06/12 - 13:28:46
ASID: 0027 JOBNAME: TRACKED4 TCB: 004E6850
STEPNAME: STEPA PROCSTEP: STEPA JOBID: STC00042 USERID: +++++++
JOBSTART: 2011/06/12 - 13:28:35
ERROR: ADDRESS SPACE APPEARS TO BE IN A LOOP.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.

Figure 22. Runtime Diagnostics report within the SMF arrival rate check

Migration considerations for PFA

All Predictive Failure Analysis (PFA) migration actions are covered in [z/OS Upgrade Workflow](#).

This section covers the following topics:

- [“Running AIRSHREP.sh” on page 76](#)
- [“How PFA uses the ini file” on page 76](#)

For complete details about installing PFA, see [“Steps for installing PFA” on page 78](#).

For all Predictive Failure Analysis (PFA) migration actions, see [z/OS Upgrade Workflow](#).

Running AIRSHREP.sh

Previously, AIRSHREP.sh was required for installation. Starting with z/OS V2R2, AIRSHREP is not required for installation, but is provided as a tool to delete existing data, if desired. Note that deleting the data is not necessary, because PFA determines if the old data is usable when PFA starts.

When running the script AIRSHREP.sh from the home directory PFA is using or when using the sample JCL for batch provided in SYS1.SAMPLIB, provide the following parameter:

new: Use the *new* parameter to delete everything from previous releases.

If you do not append the *migrate* or *new* parameter or specify the parameter incorrectly, the script fails.

How PFA uses the ini file

PFA processing uses the ini file from /etc/PFA. The system creates /etc/PFA during the installation of z/OS. If PFA does not find the /etc/PFA directory because you deleted the /etc/PFA directory created during installation, PFA will not start. You must ensure that the /etc/PFA directory is in place before starting PFA. For more information, see the topic on [Migrate /etc and /var system control files in z/OS Upgrade Workflow](#).

If the `/etc/PFA/ini` file does not already exist, PFA attempts to create the `/etc/PFA/ini` file from an existing check or else creates it from the `/usr/lpp/bcp/samples/PFA/ini` file.

Note: When PFA starts up, it deletes the ini file from each of the existing check directories.

How to define a PFA directory path

By default, the root directory where PFA creates the check subdirectories is the home directory of the PFA user. If you want to use a different directory for PFA check subdirectories, you can define your own root directory by specifying optional parameter `PFADIR=` in the ini file, as shown in the following example:

```
PFADIR= /myPFAdirectorypath
```

Make sure that the PFA user has read / write authority to that directory. If the directory does not exist or the format of the directory is not specified correctly, PFA does not start and the system issues a message stating the reason.

If you do define your own root directory to a `PFADIR=` directory, PFA preserves the `EXCLUDED_JOBS` and `INCLUDED_JOBS` for all checks.

Note: If there is an existing PFA directory and `PFADIR=` is specified in ini file, then during the very first start of PFA, PFA will:

- Copy the existing PFA directory structure and `/config` dir files to `/mypfadirectorypath`
- Delete the existing PFA directory.

If the path to the JDK for your installation is not the same as the path in the ini file in `/etc/PFA/` or if you installed the PFA Java™ code in a location other than the default path, you must update `/etc/PFA/ini` before starting PFA. For more information, see [“Updating the Java path”](#) on page 82.

Installing PFA

Before you begin: Before you install PFA in your environment, you must initialize z/OS UNIX and install the following products on your system:

- IBM 31-bit SDK for z/OS Java Technology Edition, Version 8.0 or higher. For more information about Java, see the [z/OS Java home page \(www.ibm.com/systems/z/os/zos/tools/java\)](http://www.ibm.com/systems/z/os/zos/tools/java).

Restriction: PFA does not support IBM 64-bit SDK for z/OS Java Technology Edition.

- IBM Health Checker for z/OS. You must be familiar with the set up for IBM Health Checker for z/OS. Most of the setup for PFA involves security definitions that are similar to the setup for any other started task and remote check. You must ensure that both PFA and IBM Health Checker for z/OS have access to the necessary resources including z/OS UNIX. For IBM Health Checker for z/OS details, see [Setting up IBM Health Checker for z/OS](#) in *IBM Health Checker for z/OS User's Guide*.

Guidelines:

1. The examples this procedure show are for illustrative purposes only. Replace the example parameters with the correct specifications for your environment.
2. This documentation uses `/pfa` as the generic term for the directory PFA uses for its persistent storage.
3. The z/OS File System (zFS) is a z/OS UNIX file system that contains files and directories that can be accessed with z/OS UNIX application programming interfaces (APIs) and supports access control lists (ACLs). In this documentation, all references to the z/OS UNIX file system assume that you are using zFS. For complete zFS details, see [z/OS File System Administration](#).
4. If z/OS UNIX is shut down for any reason, restart PFA.
5. Each check has specific guidelines. For example, the `PFA_JES_SPOOL_USAGE` check works only for JES2.

Steps for installing PFA

Use the following steps to set up PFA with RACF and z/OS UNIX:

1. Define additional DASD storage for PFA. The recommended `/etc/PFA/ini` file system is zFS. Requirements for the total space for the PFA file system for each LPAR is 300 cylinders primary; 50 cylinders secondary on a 3390 device.
2. Create a PFA user ID and connect it to an existing or new RACF group.

Guideline: This documentation uses *pfouser* as the generic term for the PFA user ID. You must define authority as follows:

- If `/etc/PFA/ini` does not exist, you must define READ authority to the file you are copying (`/usr/lpp/bcp/samples/PFA/ini` or a check dir ini file) in order to copy it. If you are doing this manually, you also need READ authority. If this is happening at PFA startup, the PFA user needs READ authority.
- If `/etc/PFA/ini` does not exist, you need READ and WRITE authority to `/etc/PFA` to create and use the `/etc/PFA/ini` file. If you are doing this manually, you also need READ and WRITE authority to `/etc/PFA`. If this is happening at PFA startup, the PFA user needs READ and WRITE authority.
- If `/etc/PFA/ini` already exists, the PFA user needs READ authority to it to run PFA.
- If the `/etc/PFA/ini` file needs updating, you need WRITE authority.

Note: Certain installation tasks require UID of 0.

If you are using PFA in a sysplex that shares file systems for z/OS UNIX, use a unique directory for each LPAR. The event data that PFA writes to the file system is stored separately for each system. For details, see [“Installing PFA in a z/OS UNIX shared file system environment”](#) on page 80.

- a. Create a new user ID to own the PFA. For example, `pfouser`. The PFA user ID must be unique; do not use the same user ID that is assigned to the IBM Health Checker for z/OS.
- b. Ensure `/etc/PFA` has the same security settings as `pfouser` or `pfouser` owns `/etc/PFA`.
- c. Define the PFA started task by creating a RACF profile for the `pfouser` with the following items:
 - OMVS segment with a UID parameter (for example, `omvs(uid(7))`)
 - PROGRAM pathname of `/bin/sh` (for example, `program(/bin/sh)`)

Examples:

- This example shows how you can define and connect a new user to RACF. Ensure that you replace the parameters with the correct settings for your installation.

```
ADDUSER pfouser
          OMVS(UID(7) PROGRAM(/bin/sh)) PASSWORD(sys1)
ADDGROUP OMVSGRP OMVS(GID(46))
CONNECT pfouser GROUP(OMVSGRP)
```

- This example shows how you can change the information in a user's RACF profile:

```
altuser pfouser omvs(uid(7) shared program(/bin/sh))
```

- For information about Security Server RACF, see the [ADDUSER \(Add user profile\)](#) and [ADDGROUP \(Add group profile\)](#) topics in [z/OS Security Server RACF Command Language Reference](#).
- For information about defining z/OS UNIX users to RACF, see the topic on [Steps for defining z/OS UNIX users to RACF](#) in [z/OS UNIX System Services Planning](#).

3. Add the PFA task to the STARTED class table in RACF and refresh, if necessary. For example:

```
SETROPTS GENERIC(STARTED)
RDEFINE STARTED PFA.* STDATA(USER(pfouser) GROUP(OMVSGRP))
SETROPTS CLASSACT(STARTED)
SETROPTS RACLIST(STARTED)
```

If you have already activated RACLIST in the STARTED class, the last statement is:


```
SETROPTS RACLIST(STARTED) REFRESH
```

For more information, see the following information:

- [z/OS UNIX System Services Planning](#) and [z/OS Security Server RACF Security Administrator's Guide](#).
 - The RDEFINE (Define general resource profile) and SETROPTS (Set RACF options) sections in [z/OS Security Server RACF Command Language Reference](#).
4. Copy the sample PFA procedure, AIRPROC, from SYS1.SAMPLIB to the PFA member of SYS1.PROCLIB data set. If SMP/E does not write the executable code in the z/OS UNIX file system to PARM= ' path=(/usr/lpp/bcp) ', change the PARM value in AIRPROC to the path in which you store the executable code.
 5. Update Java configuration. If the path to the JDK for your installation is not the same as the path in the ini file in /etc/PFA/ or if you installed the PFA Java code in a location other than the default path, you must update /etc/PFA/ini before starting PFA. For more information, see [“Updating the Java path” on page 82](#).

Note: Read the topic on [“Configure supervised learning” on page 91](#) to decide whether you need to use the EXCLUDED_JOBS and INCLUDED_JOBS files.

6. Optionally override Language Environment (LE) runtime options in file /etc/PFA/ceeopts. Specify the LE runtime options as the environment variable in the following format

```
_CEE_RUNOPTS=value
```

where *value* is a null terminated character string of LE runtime options. PFA uses the LE runtime options specified each time the Java virtual machine (JVM) is launched.

For the LE runtime options, see [z/OS Language Environment Programming Reference](#)

7. Allow the appropriate people access to the PFA results in SDSF and the z/OS UNIX file system. (Both systems use standard security controls.)
8. Verify that common storage tracking (CSA Tracker) is active and the SMF exits (in SMFPRMxx) are defined.

To use the PFA_COMMON_STORAGE_USAGE check, you must ensure that your system is using the following DIAGxx parmlib member options: VSM TRACK CSA(ON) SQA(ON).

- For more information about using DIAGxx, see [z/OS MVS Initialization and Tuning Reference](#) and [z/OS MVS Initialization and Tuning Guide](#).
- For information about activating and reviewing data provided by CSA Tracker, see the topic about "Using the common storage tracking function" in [z/OS MVS Initialization and Tuning Guide](#).

To use the PFA_SMF_ARRIVAL_RATE check, you must define the SMF exits, IEFU83, IEFU84, and IEFU85.

- For information about defining SMF exits, see the topic about SMFPRMxx in [z/OS MVS Initialization and Tuning Reference](#).
9. Update the COMMNDxx parmlib member, and any system automation your installation has defined, with the PFA procedure to ensure that PFA restarts on IPL as shown here:

```
START pfa_procname
```

Make sure that you define the IBM Health Checker for z/OS procedure in COMMNDxx. See [IBM Health Checker for z/OS User's Guide](#).

Important: When updating your COMMNDxx parmlib member, remember to update any system automation your installation uses to start and restart major system address spaces.

10. Update your WLM Service Class policy for PFA to be the same priority that your installation uses for monitoring products like RMF. For more information about defining service classes, see "Defining service classes and performance goals" in [z/OS MVS Planning: Workload Management](#).

11. Increase the MAXFILEPROC settings in BPXPRMxx if your current settings are too low. The MAXFILEPROC sets the maximum number of file descriptors that a single process can have open concurrently, such as all open files, directories, sockets, and pipes. By limiting the number of open files that a process can have, you limit the amount of system resources a single process can use at one time. You can also dynamically change the MAXFILEPROC setting using the SETOMVS command. For example:

```
SETOMVS PID=123,MAXFILEPROC=value
```

References :

- See the topic on [Steps for changing the process limits for an active process in z/OS UNIX System Services Planning](#).
 - To determine if the MAXFILEPROC value is set too low, the IBM Health Checker for z/OS provides a check, USS_MAXSOCKETS_MAXFILEPROC. For details, see the check [USS_MAXSOCKETS_MAXFILEPROC in IBM Health Checker for z/OS User's Guide](#).
12. Customize your system settings for PFA:
 - a. Update your system automation to look for PFA exception messages. For complete details, see the topic about [Approaches to automation with IBM Health Checker for z/OS in IBM Health Checker for z/OS User's Guide](#).
 - b. Follow the guidelines for correcting problems by reviewing the check-specific best practice.
 - c. After the checks have been running for a while, you might decide that the default parameters are not appropriate in your environment. You can customize the PFA checks using the check-specific parameters. For details, see the check-specific parameters.

Installing PFA in a z/OS UNIX shared file system environment

In this procedure you create a z/OS UNIX file system that is shared among members of the sysplex with directories that are local to the LPAR. This procedure uses zSeries File System (zFS) because it is the strategic file system. This procedure enables you to define one started task and user ID that permits PFA to write files to a system-unique directory.

Requirement: The user ID that performs this installation must have a UID of 0. In some installations, this might mean your z/OS UNIX Administrator.

1. Define the file systems as one for each LPAR. You can use the TSO ISHELL (ISPF shell) panel to define and format the zFS file system. In this example, the Integration Test team used OMVSSPT.Z1.PFA.ZFS as the example system.




Figure 23. Defining the file system

For more information about using ISHELL, see the topic [Invoking the ISPF shell in z/OS UNIX System Services User's Guide](#).

2. After defining the file systems for each LPAR, define a symbolic link (also called a symlink) to the sysplex root. From the root directory, enter the following command (the user issuing `cd` and `ln` commands require a UID of 0):

```
cd /
ln -s \${SYSNAME}/pfa pfa
```

This results in a home directory of `/systemname/pfa`.

3. Create the PFA directory in each of the system directories by entering the following commands for each of your system names:

For example, the command for system Z1 is: `mkdir /Z1/pfa`:

```
mkdir /Z1/pfa
mkdir /Z2/pfa
mkdir /Z3/pfa
mkdir /Z4/pfa
```

4. Create the new file systems (one for each system) and mount them at the appropriate system mount point. For example, for `OMVSSPT.Z1.PFA.ZFS`, the mount point is `z1/pfa`:

```
OMVSSPT.Z1.PFA.ZFS
OMVSSPT.Z2.PFA.ZFS
OMVSSPT.Z3.PFA.ZFS
OMVSSPT.Z4.PFA.ZFS
```

Here is an example of the file system attributes:

```
File System Attributes

File system name:
OMVSSPT.Z1.PFA.ZFS
Aggregate name:
OMVSSPT.Z1.PFA.ZFS
Mount point:
/Z1/pfa

More: +
```

Predictive Failure Analysis (PFA)

```
Status . . . . . : Available
File system type . . . : ZFS
Mount mode . . . . . : R/W
Device number . . . . : 28412
Type number . . . . . : 1
DD name . . . . . :
Block size . . . . . : 1024
Total blocks . . . . . : 720000
Available blocks . . . : 561716
```

See the topic [Managing the z/OS UNIX file system in z/OS UNIX System Services Planning](#).

5. Place an entry in the SYS1.PARMLIB(BPXPRMXX) member to mount the new file systems during IPL. (If you do not want to wait until the next IPL, you can manually mount these file systems.) Use the UNMOUNT attribute on the BPXPRMXX parmlib member to unmount the file system when OMVS or the LPAR is taken down. The file system mount point is */sysname/pfa* for example:

```
SYS1.PARMLIB(BPXPRM00)
  MOUNT FILESYSTEM('OMVSSPT.&SYSNAME..PFA.ZFS') TYPE(ZFS)
  MODE(RDWR) MOUNTPOINT('/&SYSNAME./pfa') UNMOUNT
```

Updating the Java path

Starting with z/OS V2R3, PFA uses a minimum of IBM 31-bit SDK for z/OS Java Technology Edition version 8.0.

You must update the ini file in the */etc/PFA/* directory for Predictive Failure Analysis (PFA) processing to locate the Java code on your system. For more information about the ini file, see [“How PFA uses the ini file”](#) on page 76.

The **JAVAPATH=** line must be set to the SMP/E installation path for the Java code for PFA. If you installed the Java code for PFA elsewhere, you must change this line to the path where the Java code for PFA is installed. The default is */usr/lpp/bcp*. For example:

```
JAVAPATH= /usr/lpp/bcp
```

Update the **PATH=** and **LIBPATH=** statements in the ini file to point to the executable code needed for JNI calls.

The following example shows the typical Java path that PFA uses beginning with z/OS V2R3:

```
PATH= /usr/lpp/java/J8.0/lib/s390/classic:/usr/lpp/java/J8.0/lib/s390
LIBPATH=/usr/lpp/java/J8.0/lib/s390:/usr/lpp/java/J8.0/lib/s390/classic:/lib:/usr/lib:
```

Chapter 8. Managing PFA checks

You can use the MODIFY system command to display and update Predictive Failure Analysis (PFA) checks as well as the IBM Health Checker for z/OS commands. To help you understand how to manage these differences, this section contains the following topics:

- “[Understanding how to modify PFA checks](#)” on page 83 describes the unique differences PFA checks have from traditional checks.
- “[MODIFY PFA, DISPLAY](#)” on page 85 describes the MODIFY *pfa*, DISPLAY command in detail.
- “[MODIFY PFA, UPDATE](#)” on page 88 describes the MODIFY *pfa*, UPDATE command in detail.
- “[Configure supervised learning](#)” on page 91 describes how to use the EXCLUDED_JOBS and INCLUDED_JOBS files to ensure that data for your important jobs is being collected.

For the IBM Health Checker for z/OS commands, see [IBM Health Checker for z/OS User's Guide](#).

You can make installation updates to PFA checks that persist across check refreshes and restart IBM Health Checker for z/OS by activating IBM Health Checker for z/OS policies. You might do this if some check default values are not suitable for your environment or configuration. For complete details, see the topic on [Creating and maintaining IBM Health Checker for z/OS policies](#) in [IBM Health Checker for z/OS User's Guide](#).

Restriction: The IBM Health Checker for z/OS debug commands are not the same debug parameter that PFA checks use. For details, see “[Understanding how to modify PFA checks](#)” on page 83.

Understanding how to modify PFA checks

PFA checks work differently than traditional checks. Although you can modify PFA checks using the IBM Health Checker for z/OS commands and policies (as described in [IBM Health Checker for z/OS User's Guide](#)), modifications to the PFA checks are unique in the following ways:

- How to quiesce a PFA check:

Use the MODIFY command to quiesce a PFA check. If the check is to permanently quiesce until PFA restarts, use the MODIFY command to delete the check. Because PFA checks are remote health checks, you must restart PFA to add a previously deleted check. If the check is to be quiesced now, but restarted later while PFA is still running, use the MODIFY command to deactivate the check. Deactivating a check stops the comparisons from occurring, but to stop collection and modeling you must set the COLLECTINACTIVE parameter to 0. To stop collections, modeling, and comparisons for an individual check so that the check can be reactivated without restarting PFA, do the following:

1. First, stop PFA from collecting and modeling data by setting the COLLECTINACTIVE parameter to zero. For example:

```
f hzspoc,update,check(ibmpfa,pfa_logrec_arrival_rate),  
parm('collectinactive(0)')
```

2. Next, deactivate the check in IBM Health Checker for z/OS. For example:

```
f hzspoc,deactivate,check(ibmpfa,pfa_logrec_arrival_rate)
```

- Use the MODIFY *pfa* DISPLAY command to display the check-specific parameters:

You can use the MODIFY *pfa*, DISPLAY command to display the cumulative set of modified, check-specific parameters that are currently in use by a PFA check. See “[MODIFY PFA, DISPLAY](#)” on page 85 for more information.

- Use the MODIFY *pfa* UPDATE command to read specific PFA configuration files:

You can use the `MODIFY pfa, UPDATE` command to read certain configuration files that a check supports. See [“MODIFY PFA, UPDATE” on page 88](#).

- Using the `MODIFY hzsproc` command to modify individual PFA check parameters:

You can use the `MODIFY hzsproc` command to modify individual parameters of PFA checks. When specifying the `PARM` parameter on `f hzsproc, update` for PFA checks, you do not have to specify all check-specific parameters. The parameters that are not specified are not changed. If the parameters were not previously modified, the values remain the default values. If the parameters were previously modified, the previously specified values remain.

The following example of the `MODIFY hzsproc` command sets the debug parameter for PFA to ON.

```
f hzsproc,update,check(IBM PFA,PFA_COMMON_STORAGE_USAGE),parm('debug(1)')
```

The rest of the check-specific parameters retain their default values as follows:

```
CHECK SPECIFIC PARAMETERS:
COLLECTINT                : 15
MODELINT                  : 360
COLLECTINACTIVE           : 1=ON
DEBUG                     : 1=ON
EXCDIRSDAYS               : 90
FORCEMODEL                 : NO
THRESHOLD                 : 2
E_HIGH                    : 180
E_MED                     : 300
E_LOW                     : MAX
E_NONE                    : UNUSED
```

If the check is displayed using IBM Health Checker for z/OS interfaces, the user-supplied parameters are displayed as either:

```
USER SUPPLIED PARAMETERS:  debug(1)
```

or

```
CHECK PARM: debug(1)
```

Then, when the following command is entered:

```
f hzsproc,update,check(IBM PFA,PFA_COMMON_STORAGE_USAGE),parm('collectint(10) modelint(60)')
```

the values for the check change to:

```
CHECK SPECIFIC PARAMETERS:
COLLECTINT                : 10
MODELINT                  : 60
COLLECTINACTIVE           : 1=ON
DEBUG                     : 1=ON
EXCDIRSDAYS               : 90
FORCEMODEL                 : NO
THRESHOLD                 : 2
E_HIGH                    : 180
E_MED                     : 300
E_LOW                     : MAX
E_NONE                    : UNUSED

COLLECTINT                : 10
MODELINT                  : 60
COLLECTINACTIVE           : 1=ON
DEBUG                     : 1=ON
THRESHOLD                 : 5
```

If the check is displayed again using IBM Health Checker for z/OS interfaces, the user-supplied parameters that are displayed are the last ones specified on the `MODIFY` command as follows:

```
USER SUPPLIED PARAMETERS:  collectint(10) modelint(60)
```

or

```
CHECK PARM: collectint(10) modelint(60)
```

This is not the cumulative list of parameters currently in use by the check. Therefore, to display the current parameter values being used by the check, use the `f pfa,display,checks,detail` command as listed in “MODIFY PFA, DISPLAY” on page 85.

- The `debug` parameter is a PFA check-specific parameter.

The `debug` parameter is a check-specific parameter and not the same `debug` parameter as the one in IBM Health Checker for z/OS. The `debug` parameter within IBM Health Checker for z/OS applies only to the phase of performing the check when the interval is reached and not to all other phases, such as data collection and modeling that are done internally within the PFA checks. Therefore, the `debug` parameter within IBM Health Checker for z/OS has no meaning to PFA and is ignored.

To set the `debug` parameter for PFA checks, specify it as a `parm` as in the following example:

```
f hzsproc,update,check(IBM PFA,PFA_CHECK_NAME),parm('debug(1)')
```

Each PFA check contains more information about its check-specific parameter.

MODIFY PFA, DISPLAY

Purpose

MODIFY *pfa*, DISPLAY (`f pfa,display`) issues messages with information specified as different options (listed in “Parameters” on page 85).

When displaying PFA checks using the IBM Health Checker for z/OS commands, the user-supplied parameters list contains only the parameters that were specified on the last update command not the cumulative set of modified parameters. Therefore, you must use the MODIFY *pfa*, DISPLAY command to display the check-specific parameters that are currently used by a PFA check. See “Understanding how to modify PFA checks” on page 83 for more information.

Format

```
DISPLAY
  {
    [CHECKS [,filters] [,SUMMARY] | ,DETAIL]]
    |
    [filters [,SUMMARY] | ,DETAIL]]
    |
    [STATUS]
  }
```

Parameters

CHECKS

CHECKS displays information about PFA checks.

filters

Filters specify which check or checks you want to take an action against. You can specify the wildcard character `*` for filters in the last position of the filter. An asterisk (`*`) represents any string having a length of zero or more characters.

Filters must be specified in one of the following formats:

```
CHECKS,CHECK=(check_name)
or
CHECK(check_name)
```

check_name specifies the 1- through 32-character check name.

SUMMARY

PFA issues message AIR013I with summary information about the specified checks. See [“PFA DISPLAY examples” on page 87](#). For each check matching the specified filter, the following information is returned:

- Check name
- Indicator of whether the check is eligible to run at the next interval (ACTIVE(ENABLED)) in IBM Health Checker for z/OS
- The last successful collection time
- The last successful model time

SUMMARY is the default value; it does not need to be specified in the command.

DETAIL

PFA issues message AIR018I with detailed information about the specified checks. See [“PFA DISPLAY examples” on page 87](#). For each check matching the specified filter, the following information is returned:

- Check name
- Indicator of whether the check is eligible to run at the next interval (ACTIVE(ENABLED)) in IBM Health Checker for z/OS
- Total number of collections attempted
- Total number of successful collections
- The last time the collection ran
- The last successful collection time
- The next collection time
- Total number of models attempted
- Total number of successful models
- The last time the model ran
- The last successful model time
- The next model time
- The current settings of the check-specific parameters for this check:
 - The collection interval in minutes
 - The model interval in minutes
 - Indicator of whether to collect data and model data even if the check is not eligible to run (ACTIVE(ENABLED)) in IBM Health Checker for z/OS
 - Indicator if the check is generating additional diagnostic information
 - Any other parameters that are supported for this check
- If this check supports excluded jobs, the list of jobs currently being excluded for this check.
- If this check supports included jobs, the list of jobs currently being included for this check.

STATUS

PFA issues message AIR017I with general status information for PFA. The following information is returned:

- The number of checks registered to PFA
- The number of PFA checks eligible to run (ACTIVE(ENABLED)) in IBM Health Checker for z/OS
- The number of collections currently queued
- The number of models currently queued
- The number of JVM terminations that have occurred since PFA started

Flags

None.

Error conditions

None.

Version

All releases.

PFA DISPLAY examples

Example of DISPLAY STATUS message output

```
IR017I 10.31.32 PFA STATUS
NUMBER OF CHECKS REGISTERED : 5
NUMBER OF CHECKS ACTIVE : 5
COUNT OF COLLECT QUEUE ELEMENTS: 0
COUNT OF MODEL QUEUE ELEMENTS : 0
COUNT OF JVM TERMINATIONS : 0
```

The DISPLAY STATUS message output displays in response to the following commands:

- `f pfa,display`
- `f pfa,display,status`

Example of DISPLAY SUMMARY message output

```
AIR013I 12:20:53 PFA CHECK SUMMARY
                LAST SUCCESSFUL      LAST SUCCESSFUL
CHECK NAME      ACTIVE COLLECT TIME      MODEL TIME
PFA_COMMON_STORAGE_USAGE      YES 03/25/2015 12:20 03/25/2015 12:18
PFA_LOGREC_ARRIVAL_RATE       YES 03/25/2015 12:20 03/25/2015 12:17
PFA_MESSAGE_ARRIVAL_RATE      YES 03/25/2015 12:19 03/25/2015 12:18
PFA_SMF_ARRIVAL_RATE          YES 3/25/2015 12:19 03/25/2015 12:19
PFA_JES_SPOOL_USAGE           YES 03/25/2015 12:20 03/25/2015 12:17
PFA_ENQUEUE_REQUEST_RATE      YES 03/25/2015 12:20 03/25/2015 12:18
PFA_PRIVATE_STORAGE_EXHAUSTION YES 03/25/2015 12:20 03/25/2015 12:17
```

The DISPLAY SUMMARY message output displays in response to the following commands:

- `f pfa,display,checks` - results in all checks
- `f pfa,display,checks,check=(check_name)` - results in one check
- `f pfa,display,checks,check=(PFA_*)` - can result in > 1 check
- `f pfa,display,checks,check=(*) ,summary` - results in all checks
- `f pfa,display,check(check_name)` - results in one check
- `f pfa,display,check(PFA_*)` - can result in > 1 check
- `f pfa,display,check(*) ,summary` - results in all checks

Example of DISPLAY DETAIL message output

MODIFY PFA, UPDATE

```
AIR018I 02.22.54 PFA CHECK DETAIL

CHECK NAME: PFA_MESSAGE_ARRIVAL_RATE
ACTIVE : YES
TOTAL COLLECTION COUNT : 5
SUCCESSFUL COLLECTION COUNT : 5
LAST COLLECTION TIME : 04/05/2010 10.18.22
LAST SUCCESSFUL COLLECTION TIME : 04/05/2010 10.18.22
NEXT COLLECTION TIME : 04/05/2010 10.33.22
TOTAL MODEL COUNT : 1
SUCCESSFUL MODEL COUNT : 1
LAST MODEL TIME : 04/05/2010 10.18.24
LAST SUCCESSFUL MODEL TIME : 04/05/2010 10.18.24
NEXT MODEL TIME : 04/05/2010 11.18.24
CHECK SPECIFIC PARAMETERS:
COLLECTINT : 15
MODELINT : 720
COLLECTINACTIVE : 1=ON
DEBUG : 0=OFF
EXCDIRDBAYS : 90
FORCEMODEL : NO
STDDEV : 10
TRACKEDMIN : 2
EXCEPTIONMIN : 1
CHECKLOW : 1=ON
STDDEVLOW : 4
LIMITLOW : 3
EXCLUDED JOBS:
NAME SYSTEM DATE ADDED REASON ADDED
JES2 * 03/11/2010 15:15 Excluded JES* jobs on ALL.
INCLUDED JOBS:
NAME SYSTEM DATE ADDED REASON ADDED
CRIT1 SY1 03/11/2015 15:15 Include job CRIT1 on SY1.
```

The DISPLAY DETAIL message output displays in response to the following commands:

- `f pfa,display,checks,detail` - results in all checks
- `f pfa,display,checks,check=(name),detail` - results in one check
- `f pfa,display,checks,check=(check_na*),detail` - can result in > 1 check
- `f pfa,display,checks,check=(*) ,detail` - results in all checks
- `f pfa,display,check(check_name),detail` - results in one check
- `f pfa,display,check(check_na*),detail` - can result in > 1 check
- `f pfa,display,check(*) ,detail` - results in all checks

MODIFY PFA, UPDATE

Purpose

Use the MODIFY *pfa*, UPDATE (`f pfa,update`) command to read configuration files that an individual check supports and store the values for future PFA processing.

You must use the MODIFY *pfa*, update command before changes made to check-specific configuration files are used by PFA processing.

Format

```
UPDATE
{
[CHECKS [,filters] [,ALL] [,EXCLUDED_JOBS]] [,INCLUDED_JOBS]]
|
[filters [,ALL] [,EXCLUDED_JOBS] [,INCLUDED_JOBS]]
}
```

Parameters

CHECKS

CHECKS reads the configuration file requested for every check that supports that configuration file. This parameter is optional.

filters

Filters specify which check or checks you want to take an action against. This parameter is optional. You can specify the wildcard character * for filters in the last position of the filter. An asterisk (*) represents any string having a length of zero or more characters.

Filters must be specified in one of the following formats:

```
CHECKS, CHECK=(check_name)
OR
CHECK(check_name)
```

check_name specifies the 1- through 32-character check name.

ALL

PFA reads the EXCLUDED_JOBS and INCLUDED_JOBS files and updates its internal configuration to use the values in these files. See the [Restriction: under INCLUDED_JOBS](#)

You can optionally add exclusions and inclusions. For details, see [“Configure supervised learning” on page 91](#).

EXCLUDED_JOBS

PFA reads the EXCLUDED_JOBS file and updates its internal configuration to use the values in this file. You can optionally add exclusions. For details, see [“Configure supervised learning” on page 91](#).

INCLUDED_JOBS

PFA reads the INCLUDED_JOBS file and updates its internal configuration to use the values in this file. You can optionally add inclusions. For details, see [“Configure supervised learning” on page 91](#).

Restriction: The following checks do not support dynamically updating INCLUDED_JOBS. For these checks, the INCLUDED_JOBS file must contain the address spaces you want to include when you start PFA. If you want to change the INCLUDED_JOBS file after PFA is started, you must stop PFA, update the file, and then restart PFA.

- PFA_ENQUEUE_REQUEST_RATE
- PFA_MESSAGE_ARRIVAL_RATE
- PFA_SMF_ARRIVAL_RATE

Defaults

- CHECKS – all checks are updated.
- ALL – is the default value.

Processing

For each check specified that supports EXCLUDED_JOBS or INCLUDED_JOBS, the EXCLUDED_JOBS file and the INCLUDED_JOBS file (if either or both exist for the check) is read from the checks' /config directory and stored in memory for the PFA address space to use in its processing.

Usage

Use this command when PFA configuration needs to be updated, but the update is not a part of the IBM Health Checker for z/OS configuration and parameters.

Restrictions

None.

Authorization

The user running the MODIFY PFA,UPDATE command must be authorized to read the EXCLUDED_JOBS file and/or the INCLUDED_JOBS file for the check.

Flags

None.

Error conditions

None.

Messages

- AIR024I
- AIR025I
- AIR027I
- AIR028I
- AIR029I
- AIR030I
- AIR031I
- AIR059I
- AIR060I
- AIR061I
- AIR062I

For the complete message text, see the topic on [AIR messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).

Version

All releases beginning with z/OS V1R12. Updated to support INCLUDED_JOBS in z/OS V2R2.

Examples

- f pfa,update or f pfa,update,ALL
 - Reads the EXCLUDED_JOBS file (if it exists) and the INCLUDED_JOBS file (if it exists) for every check that supports it and stores the new values to use for future processing.
- f pfa,update,checks,EXCLUDED_JOBS
 - Reads the EXCLUDED_JOBS file (if it exists) for every check that supports it and stores the new values to use for future processing.
- f pfa,update,checks,INCLUDED_JOBS
 - Reads the INCLUDED_JOBS file (if it exists) for every check that supports it and stores the new values to use for future processing.
- f pfa,update,check(pfa_j*)
 - Reads the EXCLUDED_JOBS file (if it exists) and the INCLUDED_JOBS file (if it exists) for every check that starts with “pfa_m” and that supports it and stores the new values to use for future processing.
- f pfa,update,check(pfa_j*),EXCLUDED_JOBS
 - Reads the EXCLUDED_JOBS file (if it exists) for every check that starts with “pfa_m” and that supports it and stores the new values to use for future processing.
- f pfa,update,check(pfa_j*),INCLUDED_JOBS

- Reads the INCLUDED_JOBS file (if it exists) for every check that starts with “pfa_m” and that supports it and stores the new values to use for future processing.

Context

- Load module: AIRAMPVT
- Entry Point: AIRA1INI

Configure supervised learning

The supervised learning service can help you avoid false positives by excluding certain data that PFA uses when it makes predictions of future behavior. To minimize the impact to check performance, use EXCLUDED_JOBS for only the conditions that cause you the most inconvenience. Instead, use other tuning parameters for the check such as STDDEV to adjust the sensitivity of the comparisons to reduce exceptions.

The supervised learning service can also help ensure that address spaces important to your installation have data that is collected for them by PFA checks. To minimize the impact to check performance, use INCLUDED_JOBS for only address spaces that are critical. Including an address space overrides other check-specific parameters that are designed to limit the address spaces for which PFA collects. For example, the parameters **COLL%** and **COLLUPTIME** parameters on the PFA_PRIVATE_STORAGE_EXHAUSTION check are ignored for address spaces that are specified in its INCLUDED_JOBS file.

A sample EXCLUDED_JOBS file ships in the /usr/lpp/bcp/samples/PFA directory. It is named EXCLUDED_JOBS and includes an example comment line. You can modify the file using the OEDIT command, and then use the `f pfa, update` command to have PFA read in the contents of the file and start to use it in during processing. Supervised learning applies to the following checks:

- [“PFA_ENQUEUE_REQUEST_RATE” on page 108](#)
- [“PFA_LOGREC_ARRIVAL_RATE” on page 125](#)
- [“PFA_MESSAGE_ARRIVAL_RATE” on page 133](#)
- [“PFA_JES_SPOOL_USAGE” on page 118](#)
- [“PFA_SMF_ARRIVAL_RATE” on page 163](#)
- [“PFA_PRIVATE_STORAGE_EXHAUSTION” on page 147](#)

For example, lets say that the JES spool usage check shows a certain job to have erratic spool usage because the job is frequently restarted, so you exclude this job from check processing.

A sample file called INCLUDED_JOBS ships in the /usr/lpp/bcp/samples/PFA directory and includes an example comment line. You can modify the file using the OEDIT command, and then use the `f pfa, update` command to have PFA read in the contents of the file and use it during processing. Supervised learning using INCLUDED_JOBS applies to the following checks:

- [“PFA_ENQUEUE_REQUEST_RATE” on page 108](#)
- [“PFA_MESSAGE_ARRIVAL_RATE” on page 133](#)
- [“PFA_JES_SPOOL_USAGE” on page 118](#)
- [“PFA_SMF_ARRIVAL_RATE” on page 163](#)
- [“PFA_PRIVATE_STORAGE_EXHAUSTION” on page 147](#)

For example, for the critical address spaces in your installation, you will want to ensure that they are always included in the tracked jobs for the PFA_SMF_ARRIVAL_RATE check, because an inconsistent SMF arrival rate could indicate damage to a critical address space.

After PFA is installed, you can optionally use the following instructions to use the supervised learning support:

1. Create the EXCLUDED_JOBS or INCLUDED_JOBS file in the /config directory for each check for which you want jobs excluded or included and that supports supervised learning of that type. You can copy the sample from either /usr/lpp/bcp/samples/PFA/EXCLUDED_JOBS or /usr/lpp/bcp/samples/PFA/INCLUDED_JOBS.

If you did not previously define the EXCLUDED_JOBS files for the checks, PFA creates the EXCLUDED_JOBS file in corresponding check /config directory. For example, the following files are preserved if they exist:

- pfa_directory/PFA_MESSAGE_ARRIVAL_RATE/config/EXCLUDED_JOBS
- pfa_directory/PFA_ENQUEUE_REQUEST_RATE/config/EXCLUDED_JOBS

The EXCLUDED_JOBS file must exist in the local /config directory for the check on each LPAR to which it applies. The generic system name allows the file to be copied from one partition to other partitions without any changes. For additional details, see [Table 5 on page 93](#).

2. Add the jobs you want to be excluded or included.
3. If PFA has already been started, run the **f pfa,update** command to cause the EXCLUDED_JOBS and INCLUDED_JOBS files to be read for the checks that support dynamically updating supervised learning.

Restriction: The following checks do not support dynamically updating INCLUDED_JOBS. The INCLUDED_JOBS file must contain the address spaces to include when PFA is started for these checks. If the INCLUDED_JOBS needs to change after PFA has started for these checks, stop PFA, update the file, and start PFA.

- “PFA_ENQUEUE_REQUEST_RATE” on page 108
- “PFA_MESSAGE_ARRIVAL_RATE” on page 133
- “PFA_SMF_ARRIVAL_RATE” on page 163

4. If you have modified the **STDDEV** or **STDDEVLOW** parameter for a check because it was receiving too many exceptions and you are now excluding the jobs that caused the exception, consider reducing the **STDDEV** or **STDDEVLOW** parameter for those checks.

Note: There are currently excluded jobs for some checks that are hardcoded and the data for these jobs is not included in the check-specific processing, even if they are put in INCLUDED_JOBS:

- “PFA_MESSAGE_ARRIVAL_RATE” on page 133: CONSOLE
- “PFA_PRIVATE_STORAGE_EXHAUSTION” on page 147: BPXAS, *MASTER*, Initiators

The checks that support supervised learning use a z/OS UNIX file with the name EXCLUDED_JOBS or INCLUDED_JOBS in the /config directory for the check. This directory is read at check initialization (when PFA starts) and when the modify PFA,update command is issued for the check. For example, find the list of excluded jobs for the PFA_MESSAGE_ARRIVAL_RATE in the /pfa_directory/PFA_MESSAGE_ARRIVAL_RATE/config/EXCLUDED_JOBS file.

Note that if you try to use supervised learning with a check that doesn't support it, the system issues message AIR026I.

If you are using PFA in a z/OS UNIX shared file system environment, each LPAR in the sysplex has a local pfa_directory in which a directory exists for each check. For each LPAR in the sysplex that needs to exclude or include jobs, the EXCLUDED_JOBS or INCLUDED_JOBS file must exist in the /config directory of the checks for which jobs should be excluded or included. To generically specify the system on which to exclude or include the job, use a system name containing a wildcard in the system name field in the EXCLUDED_JOBS or INCLUDED_JOBS file. You can then copy the file from the /config directory for a check on one LPAR to the /config directory for the same check on another LPAR.

The PFA EXCLUDED_JOBS and INCLUDED_JOBS file format is a simple, comma-separated value format as shown in [Table 5 on page 93](#). The row ends when a new line character is reached. Each field ends when a comma is reached although not all data is stored in memory.

Table 5. PFA EXCLUDED_JOBS file format

Field	Length	Format of field
Job name	8 characters maximum	<ul style="list-style-type: none"> The name of the job to exclude or include and is required. Both the job name and the system name must match at run time in order for the job to be excluded or included on this system. Wildcard characters are allowed. <ul style="list-style-type: none"> The * character is allowed in any position and denotes one or more characters. The ? character is allowed in any position and denotes one character.
System name	8 characters maximum	<ul style="list-style-type: none"> This field identifies the system to which this excluded or included job applies. This field is required. Both the job name and the system name must match at run time in order for the job to be excluded or included on this system. If this excluded job applies to all LPARs in the sysplex, specify *. If the exclusion or inclusion applies to multiple systems, but not to all systems in the sysplex, use wildcard characters: <ul style="list-style-type: none"> The * character is allowed in any position and denotes one or more characters. The ? character is allowed in any position and denotes one character.
Date and/or time of adding	A maximum of 16 characters are stored in memory	<p>This field is a character string that provides usability so you can see when you added the exclusion or inclusion. There is no specific format and is not used by the code other than to display it when you request it.</p> <ul style="list-style-type: none"> The length is 16 characters to fit the date and time in a numerical format such as: 06/10/2009 03:45 If you do not require this field, you can skip it and specify only the comma delimiter. If the field is longer than 16 characters, only the first 16 characters are stored in memory. No error message is issued for this situation because the text is still available in the file.
Reason for adding	A maximum of 35 characters are stored in memory	<p>This field is merely a character string and provided for usability so you can see the reason for excluding or including the job or address space. There is no specific format and is not used by the code other than to display it when you request it.</p> <ul style="list-style-type: none"> If this field is not required, it can be omitted. If it is longer than 35 characters, only the first 35 characters are stored in memory. No error message is issued for this situation because the text is still available in the file. The f pfa,display output displays the first 29 characters. All 35 characters are written to the config log. You can type more characters and they remain in the file, but only the first 35 are in memory.

Example: Valid input rows for excluded jobs

```
JES2,*,06/10/2009 03:45:35,Exclude JES2 on all systems
TEST*,SYS1,Skip all of my test jobs on SYS1
DUMPSRV,*
```

```
MYJOB,SYS1,06/10/2009 03:45:35
MYJOB,SYS2,06/10/2009 03:45:35
```

Example: Input rows that are not valid:

```
JES234567,*,,This name was too long.
,,06/10/2009,Name and system are required
```

Notes:

1. When processing encounters a row that is not valid, it disregards the row and issues an error message.
2. The code checks that the maximum number of characters for the job name and the system name are 8.
3. When processing encounters a duplicate job name and system name combination from a specific file, it uses the more generic occurrence, ignores the less generic occurrence, and issues a message.

For example, if the following two rows are both found in the EXCLUDED_JOBS file, the second row is used because the first row is included in the generic name combination:

- JES2,SY1,06/10/2015,Excluded JES2 on system SY1
 - JES*,*,06/10/2015,Exclude all JES* jobs on all systems
4. PFA doesn't allow overlapping data between the EXCLUDED_JOBS and INCLUDED_JOBS files. The row used is the first one processed and the file processed first depends on whether PFA is starting or a modify PFA command is being issued.

Now, let's say that neither file existed at startup for a check that supports both files and that allows dynamic updates to both files. Later, the modify PFA command is used to process the INCLUDED_JOBS file. Then, at some point later, the modify PFA command is used to process the EXCLUDED_JOBS file, which had overlapping entries with the INCLUDED_JOBS file. In this case, the first file that is processed is the INCLUDED_JOBS file, which includes entries for PFA's processing. The second file that is processed is the EXCLUDED_JOBS file with the overlapping entries. PFA finds the EXCLUDED_JOBS file in error, issues an error message, and does not allow the overlapping entries.

5. The EXCLUDED_JOBS and INCLUDED_JOBS files supports comment lines. If a line starts with /* or #, processing ignores the line. No ending comment is necessary however, the comment lines as well as all lines must end in a new line character '\n'.

Chapter 9. Predictive Failure Analysis checks

Predictive Failure Analysis (PFA) provides the following remote checks:

- [“PFA_COMMON_STORAGE_USAGE” on page 95](#)
- [“PFA_ENQUEUE_REQUEST_RATE” on page 108](#)
- [“PFA_JES_SPOOL_USAGE” on page 118](#)
- [“PFA_LOGREC_ARRIVAL_RATE” on page 125](#)
- [“PFA_MESSAGE_ARRIVAL_RATE” on page 133](#)
- [“PFA_PRIVATE_STORAGE_EXHAUSTION” on page 147](#)
- [“PFA_SMF_ARRIVAL_RATE” on page 163](#)

PFA_COMMON_STORAGE_USAGE

Description:

The check is looking to see if there is a potential for storage to be exhausted in the upcoming predictive failure analysis (PFA) model interval. PFA analyzes the following storage locations:

- common storage area (CSA)
- system queue area (SQA)
- extended common storage area (ECSA)
- extended system queue area (ESQA)
- CSA + SQA
- ECSA + ESQA

The PFA_COMMON_STORAGE_USAGE check detects three classes of common storage exhaustion:

- Spike
- Leak
- Creep

If PFA detects that there is a potential for the exhaustion of common storage, PFA issues exception message AIRH101E and provides a list of suspect tasks in the report. During the analysis, this check writes the common storage usage data at intervals to a z/OS UNIX System Services file in comma-separated value (.csv) format. The check identifies a list of users of common storage that might contribute to exhausting common storage. If deeper analysis is necessary, PFA also provides files that contain additional diagnostic information that you can examine. See [“Best practice:” on page 95](#).

PFA also issues the following informational messages:

- AIRH102I
- AIRH103I
- AIRH132I

Reason for check:

If the system runs out of common storage, jobs and started tasks experience abends.

Best practice:

The best practice is to predict common storage problems before they occur, determine the cause of the problem, and take the appropriate action.

When IBM Health Checker for z/OS issues exception message AIRH101E, PFA has predicted that the amount of storage allocated to the common storage area is in jeopardy of being exhausted. Use the following steps to determine the appropriate action:

1. Examine the Common Storage Usage Prediction Report issued with the exception message. This report contains the total current usage and predictions for each of the six storage locations: CSA, SQA, ECSA, ESQA, CSA+SQA, and ECSA+ESQA. It also contains up to ten "users" each of CSA, SQA, ECSA, and ESQA whose usage has changed the most in the last model interval. The cause of the problem is most likely within this list of users. See ["Output:" on page 103](#) for the example report.
2. If the cause of the problem is not obvious from the common storage usage report, you can obtain additional information in the csadata and the csaAlldata files, or from other checks that list the top users of storage such as the checks owned by IBMVSM (VSM_CSA_THRESHOLD and VSM_SQA_THRESHOLD). The files are text files in comma-separated value (.csv) format and contain the historical data on the usage for each interval. You can export the files into any spreadsheet-type program.
3. Determine which type of common storage problem is occurring by examining the symptoms, and then correct the behavior:

- **Spike**: A piece of code uses more and more of the common storage area with usage growing linearly or exponentially over time. If the problem is caused by a spike, the `systemName.5day.All.data` file contains one or more users that are in the last few intervals and that consume a significant and measurable amount of common storage.

Determine if the job causing the spike can be stopped, canceled, or slowed without affecting the overall system behavior.

- **Leak**: A piece of code returns some but not all of the storage, which results in more usage of the common storage area over time. If the problem is caused by a leak, look for the contributor that is on the list multiple times, but not in every interval.

Determine if the job causing the leak can be stopped, canceled, or slowed down without affecting the overall system behavior.

- **Creep**: The common storage area usage grows slowly reflecting the overall system usage, which means there is no individual user of CSA responsible for the storage exhaustion. If there is no job or address space that appears to be using an excessive or unusual amount of common storage, the amount of work being done by the LPAR is probably causing the usage of common storage to creep.

Determine if the amount of work being sent to this LPAR can be reduced.

Note: Because of the random variation in common storage usage that typically occurs and the PFA check collects and models data at defined intervals, PFA is unable to detect all leaks, spikes, and creeps.

- PFA is sometimes unable to detect a leak or creep that is less than 750 bytes per second.
- PFA cannot detect rapid growth that occurs on a machine time frame such as within a collection interval.
- PFA cannot detect common storage exhaustion caused by fragmentation.

z/OS releases the check applies to:

z/OS V1R10 and later.

Type of check:

Remote

Restrictions

Ensure your system is using the following DIAGxx parmlib member options: VSM TRACK CSA(ON) SQA(ON)

Parameters accepted:

Yes, as follows:

Table 6. PFA_COMMON_STORAGE_USAGE check parameters

Parameter name	Default value	Minimum Value	Maximum Value	Description
collectint	15 minutes	1	360	This parameter determines the time (in minutes) to run the data collector that determines the amount of common storage being used. The default is 15 minutes (15).
modelint	720 minutes	4	1440	This parameter determines how often (in minutes) you want the system to analyze the data and construct a new common storage usage model or prediction. Note that, even when you set a value larger than 360, PFA performs the first model at 360 minutes (6 hours). By default, PFA analyzes the data and constructs a new model every 720 minutes (12 hours). The model interval must be at least four times larger than the collection interval. If necessary modeling occurs more frequently.
threshold	2 percent	1	100	The percentage of the capacity over 100% of the total capacity that defines the point of exhaustion. For example, if 2 percent is specified, the prediction must be at least 102% of the capacity for an area before comparisons will be performed. The threshold can be used to reduce false positive comparisons. Setting the threshold too high might cause exhaustion problems to be undetected. The default is 2 percent (2).
collectinactive	1 (on)	0 (off)	1 (on)	Defines whether data is collected and modeled even if the check is not eligible to run (is not ACTIVE(ENABLED)) in IBM Health Checker for z/OS.
excdirdays	90	30	365	This parameter specifies the maximum number of days that an exception directory is kept for the check. The maximum number of exception directories kept is 30 regardless of age. Directories are deleted when they are this number of days old.

Table 6. PFA_COMMON_STORAGE_USAGE check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
forcemodel	no	00	23:59	<p>This parameter specifies whether a model is forced at a specific time even if the next model time is not reached. The NO value indicates that a model is not forced. Other values must be specified in a 24-hour time format and indicate the time of day to force the model. Modeling in a stable environment then proceeds the MODELINT minutes after this model. If the data becomes unstable for this metric, modeling occurs more often as needed after this time.</p> <p>For example, specifying FORCEMODEL (05:00) MODELINT (780) causes modeling to occur at 05:00 followed by the next model 13 hours later, at 18:00, in a stable environment. When using the values in an unstable environment, modeling occurs at 05:00 and at any point when modeling is needed because of instability. The FORCEMODEL time is not in effect until PFA models one time. PFA must establish a model, without delaying the first model, until the FORCEMODEL time, and PFA must have sufficient data for the first model. Also, if modeling occurred recently, such that there are no new collections at the FORCEMODEL time, the next model time is set to MODELINT minutes from the FORCEMODEL time.</p>
debug	0 (off)	0 (off)	1 (on)	<p>This parameter (an integer of 0 or 1) is used at the direction of IBM service to generate additional diagnostic information for the IBM Support Center. This debug parameter is used in place of the IBM Health Checker for z/OS policy. The default is off (0).</p>
e_high	180	0 or UNUSED	1440 or MAX	<p>This parameter indicates when a HIGH severity exception which results in a critical eventual action WTO is issued. The value specified is the maximum number of minutes before exhaustion is detected in order for a HIGH severity exception to be issued. The default is 180 which means that if exhaustion is predicted to occur between now and 180 minutes from now, a HIGH severity exception will occur.</p> <p>A value of 0 or UNUSED indicates that this dynamic severity is not used.</p> <p>A value of MAX or 1440 indicates that this dynamic severity is used to the end of the time the current model is active. It is recommended to specify MAX or 1440 for the lowest level of severity to be used. If none of the dynamic severity parameters specify MAX or 1440, the severity used for exception messages for the time range left after the last dynamic severity specified will be SYSTEM which may resolve to a higher severity than the last value specified.</p> <p>The values for E_HIGH, E_MED, E_LOW, and E_NONE cannot overlap.</p>

Table 6. PFA_COMMON_STORAGE_USAGE check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
e_med	300	0 or UNUSED	1440 or MAX	<p>This parameter indicates when a MED severity exception which results in an eventual action WTO is issued. The value specified is the maximum number of minutes before exhaustion is detected in order for a MED severity exception to be issued. The default is 300 which means that if exhaustion is predicted to occur between the start of this dynamic severity time range and 300 minutes from now, a MED severity exception will occur.</p> <p>The start of this dynamic severity time range occurs when the time range for the dynamic severity for E_HIGH has expired if it is non-zero.</p> <p>A value of 0 or UNUSED indicates that this dynamic severity is not used.</p> <p>A value of MAX or 1440 indicates that this dynamic severity is used to the end of the time the current model is active. It is recommended to specify MAX or 1440 for the lowest level of severity to be used. If none of the dynamic severity parameters specify MAX or 1440, the severity used for exception messages for the time range left after the last dynamic severity specified will be SYSTEM which may resolve to a higher severity than the last value specified.</p> <p>The values for E_HIGH, E_MED, E_LOW, and E_NONE cannot overlap.</p>

Table 6. PFA_COMMON_STORAGE_USAGE check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
e_low	MAX	0 or UNUSED	1440 or MAX	<p>This parameter indicates when a LOW severity exception which results in an informational WTO is issued. The value specified is the maximum number of minutes before exhaustion is detected in order for a LOW severity exception to be issued. The default is 1440 which means that if exhaustion is predicted to occur between the start of this dynamic severity time range and the expiration of the prediction, a LOW severity exception will occur.</p> <p>The start of this dynamic severity time range occurs when the time range for the dynamic severity for E_HIGH and E_MED have expired if either or both are non-zero.</p> <p>A value of 0 or UNUSED indicates that this dynamic severity is not used.</p> <p>A value of MAX or 1440 indicates that this dynamic severity is used to the end of the time the current model is active. It is recommended to specify MAX or 1440 for the lowest level of severity to be used. If none of the dynamic severity parameters specify MAX or 1440, the severity used for exception messages for the time range left after the last dynamic severity specified will be SYSTEM which may resolve to a higher severity than the last value specified.</p> <p>The values for E_HIGH, E_MED, E_LOW, and E_NONE cannot overlap.</p>

Table 6. PFA_COMMON_STORAGE_USAGE check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
e_none	UNUSED	0 or UNUSED	1440 or MAX	<p>This parameter indicates when an exception with no severity is issued. The value specified is the maximum number of minutes before exhaustion is detected in order for an exception with no severity to be issued. The default is UNUSED which means that this severity level is not used by default.</p> <p>When a non-zero value is specified, the start of this dynamic severity time range occurs when the time range for the dynamic severity for E_HIGH, E_MED, and E_LOW have expired when they are non-zero.</p> <p>A value of 0 or UNUSED indicates that this dynamic severity is not used.</p> <p>A value of MAX or 1440 indicates that this dynamic severity is used to the end of the time the current model is active. It is recommended to specify MAX or 1440 for the lowest level of severity to be used. If none of the dynamic severity parameters specify MAX or 1440, the severity used for exception messages for the time range left after the last dynamic severity specified will be SYSTEM which may resolve to a higher severity than the last value specified.</p> <p>The values for E_HIGH, E_MED, E_LOW, and E_NONE cannot overlap.</p>

To determine the status of the common storage usage check, issue `f pfa,display,check(pfa_common_storage_usage),detail`. See [“PFA DISPLAY examples” on page 87](#) for the complete command example. The following is an example of the output written to message AIR018I in SDSF user log (ULOG):

```

F PFA,DISPLAY,CHECK(PFA_COMMON_STORAGE_USAGE),DETAIL
AIR018I 16:20:21 PFA CHECK DETAIL
CHECK NAME: PFA_COMMON_STORAGE_USAGE
ACTIVE : YES
TOTAL COLLECTION COUNT : 5
SUCCESSFUL COLLECTION COUNT : 5
LAST COLLECTION TIME : 09/01/2008 10:18:22
LAST SUCCESSFUL COLLECTION TIME : 09/01/2008 10:18:22
NEXT COLLECTION TIME : 09/01/2008 10:33:22
TOTAL MODEL COUNT : 1
SUCCESSFUL MODEL COUNT : 1
LAST MODEL TIME : 09/01/2008 10:18:24
LAST SUCCESSFUL MODEL TIME : 09/01/2008 10:18:24
NEXT MODEL TIME : 09/01/2008 22:18:24
CHECK SPECIFIC PARAMETERS:
COLLECTINT : 15
MODELINT : 720
COLLECTINACTIVE : 1=ON
DEBUG : 0=OFF
EXCDIRSDAYS : 90
FORCEMODEL : NO
THRESHOLD : 2
E_HIGH : 180
E_MED : 300
E_LOW : MAX
E_NONE : UNUSED

```

Dynamic Severity Settings: Note that the exception message issued will be sent with the severity based on the dynamic severity parameters, E_HIGH, E_MED, E_LOW, and E_NONE if any are non-zero

or any are a value other than UNUSED. The severity is based on the amount of time before the exhaustion is predicted to occur.

The parameters indicate the maximum number of minutes from the time the exception is issued to the predicted exhaustion. For example, the default values state that a HIGH severity message is issued if the exhaustion will occur within 180 minutes, a MED severity message is issued if the exhaustion will occur from over 180 minutes to 5 hours, a LOW severity message is issued if the exhaustion will occur from over 5 hours to the end of the model time. The severity of NONE is not used by default. If you do not specify MAX (or 1440) for one of the severity values, the system will choose the severity based on the severity in the WTOTYPE and SEVERITY parameters of the check. This system severity value may be larger than the lowest dynamic severity specified; therefore, it is recommended to specify MAX for one of the dynamic severity parameters when dynamic severity is being used.

For example, if HIGH severity is not desired and MED severity is desired for exhaustion predicted to occur from now until 6 hours from now and LOW severity is desired for anything farther than 6 hours from now to the end of the model time, use the following command or create a policy in HZSPRMxx to make the change persistent:

```
f hzsproc,update,check(ibmpfa,pfa_p*),parm('E_HIGH(UNUSED) E_MED(720)
E_LOW(MAX) E_NONE(UNUSED)')
```

For another example, if HIGH severity and MED severity defaults are desired, but you want no severity exceptions for everything beyond 5 hours rather than LOW severity exceptions, use the following command or create a policy in HZSPRMxx to make the change persistent:

```
f hzsproc,update,check(ibmpfa,pfa_p*),parm('E_HIGH(180) E_MED(300)
E_LOW(UNUSED) E_NONE(MAX)')
```

User override of IBM values:

The following example shows keywords you can use to override check values either on a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. See [Chapter 8, “Managing PFA checks,” on page 83](#). You can copy and modify this statement to override the check defaults:

```
ADDREPLACE POLICY[(polycyname)] [STATEMENT(name)]
UPDATE CHECK(IBM PFA,PFA_COMMON_STORAGE_USAGE)
ACTIVE
SEVERITY(MEDIUM)
INTERVAL(00:01)
PARMS=('COLLECTINT(15)', 'MODELINT(720)', 'THRESHOLD(2)',
'COLLECTINACTIVE(1)', 'DEBUG(0)', 'EXCDIRSDAYS(90)', 'FORCEMODEL(NO)',
'E_HIGH(180)', 'E_MED(300)', 'E_LOW(MAX) E_NONE(UNUSED)')
DATE(20140313)
REASON('The reason for the change.')
```

Verbose support:

The check does not provide additional details in verbose mode. The check uses the PFA DEBUG parameter (not the DEBUG parameter in IBM Health Checker for z/OS) for writing additional information to its log files for both exception and non-exception states.

Debug support:

The DEBUG parameter in IBM Health Checker for z/OS is ignored by this check. Rather, the debug parameter is a PFA check specific parameter. The IBM Health Checker for z/OS debug commands are not the same debug parameter that PFA checks use. For details, see [“Understanding how to modify PFA checks” on page 83](#).

Reference:

For more information about PFA, see the topic on [“Overview of Predictive Failure Analysis” on page 73](#).

Messages:

This check issues the following exception messages:

- AIRH101E

For additional message information, see the topics:

- AIRH messages in *z/OS MVS System Messages, Vol 1 (ABA-AOM)*.
- AIR messages in *z/OS MVS System Messages, Vol 1 (ABA-AOM)*.

SECLABEL recommended for MLS users:

SYSLOW

Output:

The common storage usage output report:

```

Common Storage Usage Prediction Report

Last successful model time      : 07/09/2009 11:08:44
Next model time                 : 07/09/2009 23:12:44
Model interval                  : 720
Last successful collection time  : 07/09/2009 11:10:52
Next collection time            : 07/09/2009 11:25:52
Collection interval              : 15

Storage Location      Current Usage in Kilobytes  Prediction in Kilobytes  Capacity When Predicted in Kilobytes  Percentage of Current to Capacity
-----
*CSA                  2796                      3152                      2956                      95%
SQA                   455                      455                      2460                      18%
CSA+SQA               3251                      3771                      5116                      64%
ECSA                  114922                    637703                    512700                    22%
ESQA                  8414                      9319                      13184                     64%
ECSA+ESQA            123336                    646007                    525884                    23%

Address spaces with the highest increased usage:

Job Name      Storage Location      Current Usage in Kilobytes  Predicted Usage in Kilobytes
-----
JOB3          *CSA                  1235                      1523
JOB1          *CSA                   752                      935
JOB5          *CSA                   354                      420
JOB8          *CSA                   152                      267
JOB2          *CSA                    75                      80
JOB6          *CSA                    66                      78
JOB15         *CSA                    53                      55
JOB18         *CSA                    42                      63
JOB7          *CSA                    36                      35
JOB9          *CSA                    31                      34

Owner gone storage:

Storage Location      Current Usage in Kilobytes
-----
CSA                   5
SQA                   0
CSA+SQA              5
ECSA                  0
ESQA                  0
ECSA+ESQA            0

* = Storage locations that caused the exception.

```

Figure 24. Common storage usage prediction report

Note: In accordance with the IBM Health Checker for z/OS messaging guidelines, the largest generated output length for decimal variable values up to 2147483647 (X'7FFFFFFF') is 10 bytes. When any PFA report value is greater than 2147483647, it displays using multiplier notation with a maximum of six characters. For example, if the report value is 222233333444445555, PFA displays it as 1973P (222233333444445555 ÷ 1125899906842) using the following multiplier notation:

Table 7. Multiplier notation used in values for PFA reports		
Name	Sym	Size
Kilo	K	1,024

<i>Table 7. Multiplier notation used in values for PFA reports (continued)</i>		
Name	Sym	Size
Mega	M	1,048,576
Giga	G	1,073,741,824
Tera	T	1,099,511,627,776
Peta	P	1,125,899,906,842

- **Last successful model time:** The date and time of the last successful model for this check. The predictions on this report were generated at that time.
- **Next model time:** The date and time of the next model. The next model will recalculate the predictions.
- **Model interval:** The value in the configured MODELINT parameter for this check. If PFA determines new prediction calculations are necessary, modeling can occur earlier.
- **Last successful collection time:** The date and time of the last successful data collection for this check.
- **Next collection time:** The date and time of the next collection.
- **Collection interval:** The value in the configured COLLECTINT parameter for this check.
- **Storage Location:** The storage location for the values in the row of the report. The location can be one of the following:
 - CSA
 - SQA
 - ECSA
 - ESQA
 - CSA + SQA
 - ECSA + ESQA

An asterisk (*) printed prior to the storage location indicates that location is the storage location that caused the exception.

When storage is expanded from SQA or ESQA to CSA or ECSA, an additional message prints on the report, exceptions for the original location are suppressed, and the storage is included in the CSA and ECSA current usage and predictions appropriately.

- **Current Usage in Kilobytes:** The amount of storage used in kilobytes in this storage location when the check was run.
- **Predicted Usage in Kilobytes:** The prediction of the usage in this storage location for the end of the model interval.
- **Capacity When Predicted in Kilobytes:** The total defined capacity for this storage location (both used and unused) at the time the prediction was made.
- **Percentage of Current to Capacity:** The percent of storage used in kilobytes in this storage location as compared to the capacity available.
- **Address spaces with the highest increased usage:** The address spaces whose storage usage for each individual storage location recently increased the most which was determined at the time of the last model. The report is sorted by predicted usage within each storage location. This list is only printed if the check issues an exception or the debug parameter is on. The number of jobs printed can vary. An asterisk printed prior to the storage location indicates that is the storage location that caused the exception. If debug is off, the only storage locations printed are those that caused the exception.

Note: If the SQA expands into the CSA, the CSA usage and predictions include the storage taken from the CSA as SQA and PFA no longer performs comparisons for the SQA. Similarly, if the ESQA

expands into the ECSA, the ECSA usage and predictions include the storage taken from the ECSA as ESQA and PFA no longer performs comparisons for the ESQA.

- `Owner gone storage`: Storage allocated to address spaces that have terminated.

Directories

When you install PFA_COMMON_STORAGE_USAGE, PFA creates the following directories that hold the executable program, log, error, data store, intermediate, and results files.

Note: The content and names for these files are subject to change and cannot be used as programming interfaces; these files are documented only to provide help in diagnosing problems with PFA.

pfa_directory

This directory contains all the PFA check subdirectories and is pointed to by the `PFADIR=` specified in the `/etc/PFA/ini` file, or, if you do not specify `PFADIR=`, the home directory of the PFA user.

pfa_directory/PFA_COMMON_STORAGE_USAGE/data

The directory for common storage usage that holds data and modeling results.

Results files:

- `systemName.prediction` - The predictions generated by modeling for the six storage locations. This file is used as input to the code that compares the predicted usage with the amount of current usage. The following example shows the common storage usage prediction report in .csv format, which is written to the `systemName.prediction` file:

```
A/TOTAL,22910,23484
B/TOTAL,763,763
C/CSA,316,316
E/ECSA,14832,14836
Q/ESQA,8078,8644
S/SQA,447,447
```

Storage location: The location where the storage was allocated. The possible values are:

- `A/TOTAL`: total above the line common storage (ECSA+ESQA)
- `B/TOTAL`: total below the line common storage (CSA+SQA)
- `C/CSA`: common storage area (CSA).
- `E/ECSA`: extended common storage area (ECSA).
- `Q/ESQA`: extended system queue area (ESQA).
- `S/SQA`: system queue area (SQA).
- `22910`: The current usage when predicted in kilobytes.
- `23484`: The prediction in kilobytes
- `systemName.prediction.html` - This file contains an .html report version of the data found in the `systemName.prediction` file.
- `systemName.diag` - The predictions for the address spaces whose common storage usage increased the most at the time of the last model. This file is not updated unless `debug` is on or an exception occurred. This file is used as input to the code that writes the top predicted users on the report.
- `systemName.diag.html` - The file contents for `systemName.diag` in .html report format as follows:
 - `User of Common Storage`: This is the identification of the user of common storage. It consists of the address space name, ASID, and PSW.
 - `Instance Count`: The number of records with this user that were factored into the prediction model.
 - `Current Estimated Common Storage Used`: The current amount of common storage used by this user in the last collection interval included in this model.

- Prediction Look Forward Seconds: The number of seconds the prediction should project into the future.
- Predicted Common Storage Usage: The predicted amount of common storage usage for this user.

Data store files:

- `systemName.5day.All.data` - Contains the usage for the last 5 days for each collection interval for all address spaces. The usage is categorized by each address space and the six locations of common storage tracked by this check.
- `systemName.sumAll` - Contains the summarized data for days 6 through 30 for each address space created from the `5day.All.data` file. After five days, the `5day.All.data` file has the data that is too old averaged and compressed to this file and then time-stamped with the start of that day. The data is then deleted from the `5day.All.data` file.
- `systemName.5day.data` - Contains the usage for the last 5 days for each collection interval for the six storage locations tracked by this check.
- `systemName.sum` - Contains the summarized data for days 6 through 30 for each of the six storage locations tracked by this check and is created from the `5day.data` file. After five days, the `5day.data` file has the data that is too old averaged and compressed to this file and then time-stamped with the start of that day. The data is then deleted from the `5day.data` file.

Intermediate files:

- `systemName.data` - Contains the data for the six storage locations that is used as the input to modeling.
- `systemName.mapmvs` - Convert PSW execution address to module name.
- `systemName.MAPREQF.OUT` - Contains the location of the module.
- `systemName.allData` - Contains the data for modeling the address spaces whose usage increased the most at the last model.
- `systemName.allData.mapmvs` -- Contains the getmain information for the address spaces whose usage increased the most at the last model.
- `systemName.All.mapmvs` -- Contains the individual getmains in bytes for all address spaces that were active at the last collection.
- `systemName.sumAllX` - This file is used during summarization of the `5day.All.data` file.
- `systemName.sumX` - This file is used during summarization of the `5day.data` file.

Additional information is written to these log files when `DEBUG(1)`.

- `stderr` - This file only contains data if messages are generated by the JVM.
- `stdout` - This file only contains data if messages are generated by the JVM.
- `systemName.cart.log` - The log file generated by modeling code that contains the execution details of modeling code.
- `systemName.builder.log` - The log file generated by intermediate code that builds the files that are input to modeling with details about code execution.
- `systemName.CONFIG.LOG` - The log file containing the configuration history for the last 30 days for this check.
- `systemName.COLLECT.LOG` - The log file used during data collection.
- `systemName.MODEL.LOG` - The log file used during portions of the modeling phase.
- `systemName.RUN.LOG` - The log file used when the check runs.
- `systemName.launcher.log` - The log file generated by launcher code.
- `systemName.tree` - This file is generated by the modeling code. It contains information about the model tree that is built based on collected common storage usage data.

pfa_directory/PFA_COMMON_STORAGE_USAGE/EXC_timestamp

This directory contains all the relevant data for investigating exceptions issued by this check at the timestamp provided in the directory name. PFA keeps directories only for the last 30 exceptions and only if the exception is newer than the number of days specified in the EXCDIRDAYS parameter. Therefore at each exception, if directories exist that are too old, they will be deleted. In addition, if more than 30 exception directories exist, the oldest directory is deleted so that only 30 exceptions remain after the last exception is issued.

- systemNameREPORT.LOG - The log file containing the same contents as the IBM Health Checker for z/OS report for this exception as well as other diagnostic information issued during report generation.

pfa_directory/PFA_COMMON_STORAGE_USAGE/config

This directory contains the configuration files for the check.

PFA_ENQUEUE_REQUEST_RATE

Description:

The PFA_ENQUEUE_REQUEST_RATE check detects damage to an address space or system by using the number of enqueue requests per amount of CPU used as the tracked metric. If PFA detects that the enqueue request rate is lower than expected, PFA calls Runtime Diagnostics to detect if an address space is hung. If PFA detects that the enqueue request rate is higher than expected, PFA calls Runtime Diagnostics to detect if there is a damaged address space. By detecting these conditions early, you can correct the problem before it causes the system to hang or crash.

The enqueue request rate check issues an exception for the following types of comparisons:

- tracked jobs
- total system

After the PFA_ENQUEUE_REQUEST_RATE check issues an exception, it does not run the next comparison type. To avoid skewing the enqueue request rate, PFA ignores the first hour of enqueue data after IPL and the last hour of enqueue data before shutdown. In addition, PFA attempts to track the same persistent address spaces that it tracked before IPL or PFA restart if the same persistent address spaces are still active. Read the topic about persistent jobs in [“PFA_MESSAGE_ARRIVAL_RATE” on page 133](#) to understand how the PFA_ENQUEUE_REQUEST_RATE check determines the top 20 persistent jobs.

By default, an EXCLUDED_JOBS file that contains the address spaces NETVIEW and *MASTER* on all systems is created during PFA initialization. Therefore, if you did not modify the EXCLUDED_JOBS file, these jobs are excluded. See [“Configure supervised learning” on page 91](#) for more information.

Guidelines:

- If you change the maximum number of concurrent ENQ, ISGENQ, RESERVE, GQSCAN and ISGQUERY requests or change system-wide defaults that use the SETGRS command or through GRSCNFxx parmlib, delete the files in the PFA_ENQUEUE_REQUEST_RATE/data directory to ensure that PFA is collecting relevant information.
- PFA never calls Runtime Diagnostics if PFA detects something is too high. When PFA detects that the enqueue request rate is higher than expected, PFA issues an exception that indicates that an address space or the system might be damaged.

This check supports supervised learning. See the topic on [“Configure supervised learning” on page 91](#).

Reason for check:

This check determines whether an LPAR or address space is damaged or hung by using the number of enqueues per amount of CPU used as the tracked metric.

Best practice:

Analyze the message and reports that are issued by PFA to determine what is causing the increase or decrease in the enqueue request rate.

z/OS releases the check applies to:

z/OS V1R13 and later.

Type of check:

Remote.

Parameters accepted:

Yes, as follows.

Table 8. PFA_ENQUEUE_REQUEST_RATE check parameters

Parameter name	Default value	Minimum Value	Maximum Value	Description
collectint	1 Minute	1	360	This parameter determines how often (in minutes) to run the data collector that retrieves the current enqueue request rate.
modelint	720 Minutes	60	1440	This parameter determines how often (in minutes) you want the system to analyze the data and construct a new enqueue request rate model or prediction. By default, PFA analyzes the data and constructs a new model every “default value” minutes. The model interval must be at least four times larger than the collection interval. Even when you set a value larger than 360, PFA does the first model at 360 minutes (6 hours). By default, PFA analyzes the data and constructs a new model every 720 minutes (12 hours).
stddev	10	2	100	This parameter is used to specify how much variance is allowed between the actual enqueue request rate per amount of CPU and the expected enqueue request rate. It determines whether the actual enqueue request rate increased beyond the allowable upper limit and how much variance is allowed across the time range predictions. If you set the STDDEV parameter to a smaller value, an exception issues when the actual enqueue request rate is closer to the expected enqueue request rate and the predictions across the time ranges are consistent. If you set the STDDEV parameter to a larger value, an exception issues when the actual enqueue request rate is significantly greater than the expected enqueue request rate even if the predictions across the different time ranges are inconsistent.
collectinactive	1 (on)	0 (off)	1 (on)	Defines whether data is collected and modeled even if the check is not eligible to run, not ACTIVE(ENABLED), in IBM Health Checker for z/OS.
trackedmin	3	0	1000	This parameter defines the minimum enqueue request rate that is required for a persistent job in order for it to be considered a top persistent job that should be tracked individually.
exceptionmin	1	0	1000	This parameter is used to determine if an exception must be issued for an unexpectedly high enqueue request rate. For tracked jobs, this parameter defines the minimum enqueue request rate and the minimum predicted enqueue request rate that is required to cause a too high exception. For the total system comparison, this parameter defines the minimum enqueue request rate that is required to cause a too high exception.
checklow	1	0	1	Defines whether Runtime Diagnostics is run to validate that a low enqueue request rate is caused by a problem. If this value is off, PFA does not issue exceptions for conditions in which the enqueue request rate is unexpectedly low.

Table 8. PFA_ENQUEUE_REQUEST_RATE check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
stddevlow	8	2	100	<p>This parameter is used to specify how much variance is allowed between the actual enqueue request rate per amount of CPU, and the expected enqueue request rate, when determining if the actual rate is unexpectedly low.</p> <ul style="list-style-type: none"> • If you set the STDDEVLOW parameter to a smaller value, an exception is issued when the actual enqueue request rate is closer to the expected enqueue request rate. • If you set the STDDEVLOW parameter to a larger value, an exception is issued when the actual enqueue request rate is significantly lower than the expected enqueue request rate.
limitlow	3	1	100	This parameter defines the maximum enqueue request rate that is allowed when issuing an exception for an unexpectedly low number of enqueues.
excdirdays	90	30	365	This parameter specifies the maximum number of days that an exception directory is kept for the check. The maximum number of exception directories kept is 30 regardless of age. Directories are deleted when they are this number of days old.
forcemodel	no	00	23:59	<p>This parameter specifies whether a model is forced at a specific time even if the next model time is not reached. The NO value indicates that a model is not forced. Other values must be specified in a 24-hour time format and indicate the time of day to force the model. Modeling in a stable environment then proceeds the MODELINT minutes after this model. If the data becomes unstable for this metric, modeling occurs more often as needed after this time.</p> <p>For example, specifying FORCEMODEL (05:00) MODELINT (780) causes modeling to occur at 05:00 followed by the next model 13 hours later, at 18:00, in a stable environment. When using the values in an unstable environment, modeling occurs at 05:00 and at any point when modeling is needed because of instability. The FORCEMODEL time is not in effect until PFA models one time. PFA must establish a model, without delaying the first model, until the FORCEMODEL time, and PFA must have sufficient data for the first model. Also, if modeling occurred recently, such that there are no new collections at the FORCEMODEL time, the next model time is set to MODELINT minutes from the FORCEMODEL time.</p>
debug	0 (off)	0 (off)	1 (on)	This parameter (an integer of 0 or 1) is used at the direction of IBM service to generate more diagnostic information for the IBM Support Center. This debug parameter is used in place of the IBM Health Checker for z/OS policy. The default is off (0).

To determine the status of the enqueue request rate check, issue `fpfa, display, check(PFA_ENQUEUE_REQUEST_RATE), detail`. For the command example and more

details, see “PFA DISPLAY examples” on page 87. The following example shows the output that is written to message AIR018I in SDSF:

```
AIR018I 02:22:54 PFA CHECK DETAIL

CHECK NAME:  PFA_ENQUEUE_REQUEST_RATE
ACTIVE              : YES
TOTAL COLLECTION COUNT      : 5
SUCCESSFUL COLLECTION COUNT : 5
LAST COLLECTION TIME       : 02/05/2009 10:18:22
LAST SUCCESSFUL COLLECTION TIME : 02/05/2009 10:18:22
NEXT COLLECTION TIME       : 02/05/2009 10:19:22
TOTAL MODEL COUNT          : 1
SUCCESSFUL MODEL COUNT     : 1
LAST MODEL TIME            : 02/05/2009 10:18:24
LAST SUCCESSFUL MODEL TIME : 02/05/2009 10:18:24
NEXT MODEL TIME            : 02/05/2009 22:18:24
CHECK SPECIFIC PARAMETERS:
COLLECTINT             : 1
MODELINT               : 720
COLLECTINACTIVE        : 1=ON
DEBUG                  : 0=OFF
EXCDIRSDAYS            : 90
FORCEMODEL             : NO
STDDEV                 : 10

TRACKEDMIN              : 3
EXCEPTIONMIN           : 1
CHECKLOW                : 1=ON
STDDEVLOW               : 8
LIMITLOW                : 3

EXCLUDED JOBS:
NAME    SYSTEM    DATE ADDED      REASON ADDED
*MASTER* *      2011/05/02 12:00 Exclude *MASTER* on all systems.
NETVIEW *      2011/05/02 12:00 Exclude NETVIEW on all
systems.
```

User override of IBM values:

The following shows keywords that you can use to override check values on either a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. This statement can be copied and modified to override the check defaults:

```
ADDREPLACE POLICY[(polycyname)] [STATEMENT(name)]
UPDATE CHECK(IBM PFA,PFA_ENQUEUE_REQUEST_RATE)
ACTIVE
SEVERITY(MEDIUM)
INTERVAL(ONETIME)
PARMS=('COLLECTINT(1)', 'MODELINT(720)', 'COLLECTINACTIVE(1)', 'DEBUG(0)',
       'EXCDIRSDAYS(90)', 'FORCEMODEL(NO)', 'STDDEV(10)', 'TRACKEDMIN(3)',
       'EXCEPTIONMIN(1)', 'CHECKLOW(1)', 'STDDEVLOW(8)', 'LIMITLOW(3)')
DATE(20140313)
REASON('The reason for the change.')
```

The enqueue request rate check is designed to run automatically after every data collection. Do not change the INTERVAL parameter.

Verbose support:

The check does not provide more details in verbose mode. The check uses the PFA DEBUG parameter (not the DEBUG parameter in IBM Health Checker for z/OS) for writing additional information to its log files for both exception and non-exception states.

Debug support:

The DEBUG parameter in IBM Health Checker for z/OS is ignored by this check. Rather, the debug parameter is a PFA check specific parameter. For details, see “[Understanding how to modify PFA checks](#)” on page 83.

Reference:

For more information about PFA, see the topic on “[Overview of Predictive Failure Analysis](#)” on page 73.

Messages:

The output is an enqueue request rate prediction report that corresponds to the message issued. PFA generates one of the following reports:

- AIRH190E - Enqueue request rate lower than expected exception
- AIRH192E - Enqueue request rate higher than expected exception
- AIRH210E - Total system enqueue request rate higher than expected exception
- AIRH211E - Total system enqueue request rate system lower than expected exception
- AIRH216I - Runtime Diagnostic output

For complete message information, see the topics on:

- [AIRH messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).
- [AIR messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).

SECLABEL recommended for MLS users:

SYSLOW

Output

The output is a variation of the enqueue request rate prediction report. The values that are found in the enqueue request prediction file are as follows:

Tracked jobs exception report for enqueue request rate higher than expected: PFA issues this report when any one or more tracked, persistent jobs cause an exception due to the enqueue request rate that is higher than expected. Only the tracked jobs that caused an exception are in the list of jobs on the report.

```

Last successful model time      : 01/27/2009 11:08:01
Next model time                : 01/27/2009 23:08:01
Model interval                 : 720
Last successful collection time : 01/27/2009 17:41:38
Next collection time           : 01/27/2009 17:56:38
Collection interval            : 15

Highest STDDEVLOW needed to avoid this exception:    21

Persistent address spaces with high rates:

Job Name      ASID      Enqueue Request Rate      Predicted Enqueue Request Rate
              ASID      Rate                       1 Hour      24 Hour      7 Day
-----
TRACKED1 001D      158.00                    23.88      22.82      15.82
TRACKED3 0029      11.00                     12.43      2.36       8.36
    
```

Figure 25. Prediction report for enqueue request rate higher than expected - tracked job

Tracked jobs exception report for enqueue request rate lower than expected: PFA issues this report when any one or more tracked, persistent jobs cause an exception due to the

enqueue request rate that is lower than expected. Only the tracked jobs that caused an exception are in the list of jobs on the report.

```

      Enqueue Request Rate Prediction Report

Last successful model time      : 10/10/2010 11:08:01
Next model time                : 10/10/2010 23:08:01
Model interval                 : 720
Last successful collection time : 10/10/2010 17:41:38
Next collection time           : 10/10/2010 17:56:38
Collection interval             : 15

Highest STDDEVLOW needed to avoid this exception:    21

Persistent address spaces with low rates:

      Job          Enqueue          Predicted Enqueue
      Name        ASID        Request          Request
      -----    -----    -----    -----
      Rate              Rate              Rate
      -----    -----    -----    -----
      1 Hour        24 Hour        7 Day
      -----    -----    -----
JOBS4    0027    1.17    23.88    22.82    15.82
JOBS5    002E    2.01    8.34    11.11    12.11

Runtime Diagnostics Output:
Runtime Diagnostics detected a problem in job: JOBS4
EVENT 06: HIGH - HIGHCPU - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 0027 JOBNAME: JOBS4
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 07: HIGH - LOOP - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID: 0027 JOBNAME: JOBS4 TCB: 004E6850
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE APPEARS TO BE IN A LOOP.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
Runtime Diagnostics detected a problem in job: JOBS5
EVENT 03: HIGH - HIGHCPU - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 002E JOBNAME: JOBS5
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 04: HIGH - LOOP - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID: 002E JOBNAME: JOBS5 TCB: 004E6850
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE APPEARS TO BE IN A LOOP.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----

```

Figure 26. Prediction report for enqueue request rate lower than expected - tracked jobs

No problem report and total system exception report for enqueue request rate higher than expected: The no problem report and the total system exception report (when the rate is higher than expected) show the totals at the top and the list of the tracked jobs.

```

Enqueue Request Rate Prediction Report

Last successful model time      : 01/27/2009 17:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval             : 15

Enqueue request rate
at last collection interval      : 83.52
Prediction based on 1 hour of data : 98.27
Prediction based on 24 hours of data: 85.98
Prediction based on 7 days of data : 100.22

Top persistent users:

Job Name      ASID      Enqueue Request Rate      Predicted Enqueue Request Rate
1 Hour      24 Hour      7 Day
-----
TRACKED1 001D      18.55      23.88      22.82      15.82
TRACKED2 0028      10.78      0.34      11.11      12.11
TRACKED3 0029      10.36      12.43      2.36      8.36
    
```

Figure 27. No problem report: enqueue request rate higher than expected

Total system exception report for enqueue request rate lower than expected: PFA issues the enqueue request rate exception report when there is a shortage or unusually low rate of enqueue requests. Runtime Diagnostics examines the system and PFA lists all output that it receives from Runtime Diagnostics.

```

Enqueue Request Rate Prediction Report

Last successful model time      : 01/27/2009 11:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 720
Last successful collection time : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval             : 15

Enqueue request rate
at last collection interval      : 0.28
Prediction based on 1 hour of data : 98.27
Prediction based on 24 hours of data: 85.98
Prediction based on 7 days of data : 100.22

Highest STDDEVLOW needed to avoid this exception:          TOO HIGH

Runtime Diagnostics Output:
-----
EVENT 01: HIGH - ENQ          - SYSTEM: SY1      2010/10/04 - 10:19:53
ENQ WAITER - ASID:002F - JOBNAME:IBMUSER2 - SYSTEM:SY1
ENQ BLOCKER - ASID:002E - JOBNAME:IBMUSER1 - SYSTEM:SY1
QNAME: TESTENQ
RNAME: TESTOFAVERYVERYVERYVERYL0000000000000000000000ONGRNAME1234567...
  ERROR: ADDRESS SPACES MIGHT BE IN ENQ CONTENTION.
  ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE BLOCKING JOBS AND
  ACTION: ASIDS.
-----
    
```

Figure 28. Total system exception report: low enqueue request rate

Note: In accordance with the IBM Health Checker for z/OS messaging guidelines, the largest generated output length for decimal variable values up to 2147483647 (X'7FFFFFFF') is 10 bytes. When any PFA report value is greater than 2147483647, it displays by using multiplier notation with a maximum of six characters. For example, if the report value is 2222233333444445555, PFA displays it as 1973P (2222233333444445555 ÷ 1125899906842) using the following multiplier notation:

<i>Table 9. Multiplier notation used in values for PFA reports</i>		
Name	Sym	Size
Kilo	K	1,024
Mega	M	1,048,576
Giga	G	1,073,741,824
Tera	T	1,099,511,627,776
Peta	P	1,125,899,906,842

The following fields apply to all reports:

- **Last successful model time:** The date and time of the last successful model for this check. The predictions on this report were generated at that time.
- **Next model time:** The date and time of the next model. The next model will recalculate the predictions.
- **Model interval:** The value in the configured MODELINT parameter for this check. If PFA determines new prediction calculations are necessary, modeling can occur earlier.
- **Last successful collection time:** The date and time of the last successful data collection for this check.
- **Next collection time:** The date and time of the next collection.
- **Collection interval:** The value in the configured COLLECTINT parameter for this check.
- **Enqueue request rate in last collection interval:** The actual enqueue request rate in the last collection interval where the rate is defined to be the count that is returned by the GRS ISGQUERY API normalized by the amount of CPU used.
- **Predicted rates based on...:** The enqueue request rates based on one hour, 24 hours, and seven days. If no prediction is available for a given time range, the line is not printed. For example, if the check is running for 2 days, there is not enough data for seven days of data therefore PFA does not print the "Prediction based on 7 days of data" line. If there is not enough data for a time range, INELGIBLE is printed for that time range and no comparisons are made.
- **Highest STDDEV needed to avoid this exception:** The highest STDDEV required to avoid the exception issued for all resources in exception status with the values in effect at this comparison.
- **Highest STDDEVLOW needed to avoid this exception:** The highest STDDEVLOW required to avoid the exception issued for all resources in exception status with the values in effect at this comparison.
- **TOO HIGH:** A STDDEV or STDDEVLOW needed to avoid the exception is TOO HIGH if the parameter value required is higher than the maximum allowed.
- **IRRELEVANT:** A STDDEV or STDDEVLOW needed to avoid the exception is IRRELEVANT if the parameter's value was not relevant in the comparison algorithm used for the values in the current comparison.
- **Runtime Diagnostics Output:** Runtime Diagnostics event records to assist you in diagnosing and fixing the problem. See the topic on "[Specific Runtime Diagnostics event reports](#)" on page 41 in Chapter 5, "[Runtime Diagnostics](#)," on page 35.
- **Job Name:** The name of the job that has enqueue arrivals in the last collection interval.
- **ASID:** The ASID for the job that has enqueue arrivals in the last collection interval.
- **Enqueue request rate:** The current enqueue request rate for the system.
- **Predicted enqueue request rate:** The predicted enqueue request rates based on 1 hour, 24 hours, and seven days of data. If PFA did not previously run on this system or the same jobs previously tracked are not all active, there is not enough data for two prediction times until that amount of time has passed. Gaps in the data, that are caused by stopping PFA or by an IPL, might cause the time range to not have enough data available. After the check collects enough data,

predictions are made again for the time ranges. If there is not enough data for the time range, INELIGIBLE is printed and comparisons are not made.

Directories

The content and names for these files and directories are subject to change and cannot be used as programming interfaces; these files are documented only to provide help in diagnosing problems with PFA.

pfa_directory

This directory contains all the PFA check subdirectories and is pointed to by the PFADIR= specified in the /etc/PFA/ini file, or, if you do not specify PFADIR=, the home directory of the PFA user.

pfa_directory/PFA_ENQUEUE_REQUEST_RATE/data

The directory for enqueue request rate that holds data and modeling results.

Guideline: If the use of the z/OS image is radically different after an IPL (for instance, the change from a test system to a production system) or if you modify anything that affects enqueue details, delete the files in the PFA_ENQUEUE_REQUEST_RATE/data directory to ensure that the check can collect the most accurate modeling information.

Results:

- `systemName.1hr.prediction` - This file is generated by the modeling code for the predictions that are made for one hour of historical data. It contains predictions for each of the tracked address spaces and the total system category. It also contains additional information that is required for PFA processing.
- `systemName.24hr.prediction` - This file is generated by the modeling code for the predictions that are made for 24 hours of historical data. It contains predictions for each of the tracked address spaces and the total system category. It also contains additional information that is required for PFA processing.
- `systemName.7day.prediction` - This file is generated by the modeling code for the predictions that are made for seven days of historical data. It contains predictions for each of the tracked address spaces and the total system category. It also contains additional information that is required for PFA processing.
- `systemName.1hr.prediction.html` - This file contains an .html report version of the data found in the `systemName.1hr.prediction` file.
- `systemName.24hr.prediction.html` - This file contains a .html report version of the data that is found in the `systemName.24hr.prediction` file.
- `systemName.7day.prediction.html` - This file contains a .html report version of the data that is found in the `systemName.7day.prediction` file.
- `systemName.prediction.stddev` - The file that is generated by the modeling code to list the standard deviation of the predictions across the time ranges for each job.

Data store files:

- `systemName.OUT` - The data collection file.

Intermediate files:

- `systemName.data` - The file is used as input to the modeling to track if enough data is available to model.
- `systemName.1hr.data` - The file that is used as input to modeling code. It contains one hour of historical data.
- `systemName.24hr.data` - The file that is used as input to modeling code. It contains 24 hours of historical data.
- `systemName.7day.data` - The file that is used as input to modeling code. It contains seven days of historical data.
- `systemName.1hr.holes` - The file is used to track gaps in data, caused by stopping PFA or by an IPL, for a one-hour period.

- `systemName.24hr.holes` - The file is used to track gaps in the data, caused by stopping PFA or by an IPL, for a 24-hour time period.
- `systemName.7day.holes` - The file is used to track gaps in the data, caused by stopping PFA or by an IPL, for the seven-day time period.

This directory holds the following log files. Additional information is written to these log files when `DEBUG(1)`.

- `stderr` - This file only contains data if messages are generated by the JVM.
- `stdout` - This file only contains data if messages are generated by the JVM.
- `systemName.1hr.cart.log` - The log file that is generated by modeling code with details about code execution while one hour of historical data was being modeled.
- `systemName.24hr.cart.log` - The log file that is generated by modeling code with details about code execution while 24 hours of historical data was being modeled.
- `systemName.7day.cart.log` - The log file that is generated by modeling code with details about code execution while seven days of historical data was being modeled.
- `systemName.builder.log` - The log file that is generated by intermediate code that builds the files that are input to modeling with details about code execution.
- `systemName.launcher.log` - The log file that is generated by launcher code.
- `systemName.1hr.tree` - This file is generated by the modeling code. It contains information about the model tree that was built based on the last one hour of collected data.
- `systemName.24hr.tree` - This file is generated by the modeling code. It contains information about the model tree that was built based on the last 24 hours of collected data.
- `systemName.7day.tree` - This file is generated by the modeling code. It contains information about the model tree that was built based on the last seven days of collected data.
- `systemNameCONFIG.LOG` - The log file that contains the configuration history for the last 30 days for this check.
- `systemNameCOLLECT.LOG` - The log file that is used during data collection.
- `systemNameMODEL.LOG` - The log file that is used during portions of the modeling phase.
- `systemNameRUN.LOG` - The log file that is used when the check runs.

pfa_directory/PFA_ENQUEUE_REQUEST_RATE/EXC_timestamp

This directory contains all the relevant data for investigating exceptions that are issued by this check at the timestamp that is provided in the directory name. PFA keeps directories only for the last 30 exceptions and only if the exception is newer than the number of days that are specified in the `EXCDIRDAYS` parameter. Therefore at each exception, if directories exist that are too old, they are deleted. In addition, if more than 30 exception directories exist, the oldest directory is deleted so that only 30 exceptions remain after the last exception is issued.

- `systemNameREPORT.LOG` - The log file that contains the same contents as the IBM Health Checker for z/OS report for this exception as well as other diagnostic information issued during report generation.

pfa_directory/PFA_ENQUEUE_REQUEST_RATE/config

This directory contains the configuration files for the check.

- `EXCLUDED_JOBS` - The file that contains the list of excluded jobs for this check.
- `INCLUDED_JOBS` - The file that contains the list of included jobs for this check.

PFA_JES_SPOOL_USAGE

Description:

The JES spool usage check detects abnormalities in the number of track groups used in individual address spaces. The check collects data for jobs that are INCLUDED_JOBS or that meet the COLLUPTIME configuration parameter. Data is modeled for the top 15 address spaces that have the greatest change in the number of track groups used since the last collection.

PFA uses the metric of the amount of JES spool usage to determine if an address space is behaving abnormally based on the previous behavior for the address space. If a change in JES spool usage is too high, it can indicate a damaged address space. This check detects abnormalities in the amount of JES spool usage as follows:

- Collection waits until one hour after IPL so that data from the IPL activity does not skew the results.
- Address spaces that match names in the INCLUDED_JOBS file for this check will be collected once they have been active for one full collection interval.
- Address spaces that start within one hour after IPL are collected after they have been active for one full collection interval.
- All other address spaces are collected after they have been active for the number of minutes specified in the COLLUPTIME configuration parameter and are active for one full collection interval.
- Address spaces that restart are considered new address spaces even if they retain the name and ASID of a previously active address space.
- Address spaces with duplicate names are considered different address spaces and are tracked individually.
- Modeling occurs for the 15 address spaces that had the greatest amount of change in the last collection.
- If abnormal behavior is detected based on the expected value for an address space that has been modeled, PFA issues a health check exception message.
- When an exception occurs, the check reports the address spaces whose JES spool usage was abnormal.
- When no problem exists, the check reports the address spaces that were modeled.
- This check supports supervised learning. For details, see [“Configure supervised learning” on page 91](#).

Reason for check:

The objective of the JES spool usage check is to detect address spaces that are damaged by comparing the amount of change in the size of the JES spool used by the address space to the expected value.

Best practice:

The best practice is to inspect the job log and other files written to the JES spool. If a job's spool usage is higher than expected, this can indicate a damaged job. To fix the problem, recycle the job.

z/OS releases the check applies to:

z/OS V1R13 and later.

Type of check:

Remote

Restrictions:

This check applies only to JES2.

Parameters accepted:

Yes, as follows:

Table 10. PFA_JES_SPOOL_USAGE check parameters

Parameter name	Default value	Minimum Value	Maximum Value	Description
collectint	5 Minutes	1	360	This parameter determines the time (in minutes) to run the data collector that determines the number of track groups used. The default is 5 minutes (5).
modelint	720 Minutes	4	1440	This parameter determines how often (in minutes) you want the system to analyze the data and construct a new model. Note that, even when you set a value larger than 360, PFA performs the first model at 360 minutes (6 hours). By default, PFA analyzes the data and constructs a new model every 720 minutes (12 hours).
collectinactive	1 (on)	0 (off)	1 (on)	Defines whether data will be collected and modeled even if the check is not eligible to run, not ACTIVE(ENABLED), in IBM Health Checker for z/OS.
stddev	3	2	100	The number by which to multiply the predicted JES spool used to determine if the actual used has increased beyond the allowable limit, which might indicate an address space is damaged.
exceptionmin	10	0 (off)	1000	This parameter defines the minimum JES spool usage required to cause a too high exception.
colluptime	360	30	1440	This parameter indicates the minimum number of minutes the address space must have been running (elapsed from address space start time) before collections will start. The default is 360 minutes. This parameter is ignored for jobs in the INCLUDED_JOBS file. This parameter is also ignored for jobs that start within the first hour after IPL.
excdirdays	90	30	365	This parameter specifies the maximum number of days an exception directory will be kept for a check. The maximum number of exception directories kept is 30 regardless of age. Additional directories are deleted once they are this number of days old. The default is 90 days.

Table 10. PFA_JES_SPOOL_USAGE check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
forcemodel	no	00	23:59	<p>This parameter specifies whether a model is forced at a specific time even if the next model time is not reached. The NO value indicates that a model is not forced. Other values must be specified in a 24-hour time format and indicate the time of day to force the model. Modeling in a stable environment then proceeds the MODELINT minutes after this model. If the data becomes unstable for this metric, modeling occurs more often as needed after this time.</p> <p>For example, specifying FORCEMODEL (05:00) MODELINT (780) causes modeling to occur at 05:00 followed by the next model 13 hours later, at 18:00, in a stable environment. When using the values in an unstable environment, modeling occurs at 05:00 and at any point when modeling is needed because of instability. The FORCEMODEL time is not in effect until PFA models one time. PFA must establish a model, without delaying the first model, until the FORCEMODEL time, and PFA must have sufficient data for the first model. Also, if modeling occurred recently, such that there are no new collections at the FORCEMODEL time, the next model time is set to MODELINT minutes from the FORCEMODEL time.</p>
debug	0 (off)	0 (off)	1 (on)	<p>This parameter (an integer of 0 or 1) is used at the direction of IBM service to generate additional diagnostic information for the IBM Support Center. This debug parameter is used in place of the IBM Health Checker for z/OS policy. The default is off (0).</p>

To determine the status of the JES spool usage check, issue `f pfa,display,check(PFA_JES_SPOOL_USAGE),detail`. For the command example and more details, see [“PFA DISPLAY examples” on page 87](#). The following example shows the output written to message AIR018I in SDSF:

```

AIR018I 02.22.54 PFA CHECK DETAIL
CHECK NAME: PFA_JES_SPOOL_USAGE
ACTIVE : YES
TOTAL COLLECTION COUNT : 5
SUCCESSFUL COLLECTION COUNT : 5
LAST COLLECTION TIME : 02/05/2009 10.18.22
LAST SUCCESSFUL COLLECTION TIME : 02/05/2009 10.18.22
NEXT COLLECTION TIME : 02/05/2009 10.23.22
TOTAL MODEL COUNT : 1
SUCCESSFUL MODEL COUNT : 1
LAST MODEL TIME : 02/05/2009 10.18.24
LAST SUCCESSFUL MODEL TIME : 02/05/2009 10.18.24
NEXT MODEL TIME : 02/05/2009 22.18.24
CHECK SPECIFIC PARAMETERS:
COLLECTINT : 5
MODELINT : 720
COLLECTINACTIVE : 1=YES
DEBUG : 0=NO
EXCDIRSDAYS : 90
FORCEMODEL : NO

```

```

STDDEV                : 3
EXCEPTIONMIN          : 10
COLLUPTIME            :
360

```

User override of IBM values:

The following shows keywords you can use to override check values on either a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. This statement can be copied and modified to override the check defaults:

```

ADDREPLACE POLICY[(polycyname)] [STATEMENT(name)]
UPDATE CHECK(IBM PFA,PFA_JES_SPOOL_USAGE)
ACTIVE
SEVERITY(MEDIUM)
INTERVAL(ONETIME)
PARMS=('COLLECTINT(5)', 'MODELINT(720)', 'STDDEV(3)', 'DEBUG(0)',
      'COLLECTINACTIVE(1)', 'EXCEPTIONMIN(10)', 'COLLUPTIME(360)')
DATE(20140313)
REASON('The reason for the change.')
```

Note: The JES spool usage check is designed to run automatically after every data collection. Do not change the INTERVAL parameter.

Verbose support:

The check does not provide additional details in verbose mode. The check uses the PFA DEBUG parameter (not the DEBUG parameter in IBM Health Checker for z/OS) for writing additional information to its log files for both exception and non-exception states.

Debug support:

The DEBUG parameter in IBM Health Checker for z/OS is ignored by this check. Rather, the debug parameter is a PFA check specific parameter. For details, see [“Understanding how to modify PFA checks”](#) on page 83.

Reference:

For more information about PFA, see the topic on [“Overview of Predictive Failure Analysis”](#) on page 73.

Messages:

The output is a JES spool usage prediction report that corresponds to the message issued. PFA generates one of the following reports:

- AIRH198E - JES spool usage exception report

For additional message information, see the topics on:

- [AIRH messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).
- [AIR messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).

SECLABEL recommended for MLS users:

SYSLOW

Output:

The output is a variation of the JES spool usage prediction report. The values found are as follows:

Exception Report (for AIRH198E) PFA issues this report when any one or more address spaces cause an exception. Only the jobs that caused an exception are in the list of jobs on the report.

```

JES Spool Usage Prediction Report

Last successful model time      : 01/28/2010 16:10:15
Next model time                : 01/28/2010 16:14:57
Model interval                 : 25
Last successful collection time : 01/28/2010 16:09:57
Next collection time           : 01/28/2010 16:14:57
Collection interval            : 5

Address spaces causing exception:

Job          Current Change in   Expected Change in   Current
Name ASID      Number of Track      Number of Track      Number of Track
JOB1  0019      Groups Used          Groups Used           Groups Used
JOB55 000E          252                  10                    892
                129                   3                     400
    
```

Figure 29. JES spool usage exception report

No problem report: When no exception is issued for the JES spool usage check (AIRH200I) is issued, the following report is generated:

```

JES Spool Usage Prediction Report

Last successful model time      : 01/27/2010 15:58:13
Next model time                : 01/27/2010 16:03:05
Model interval                 : 25
Last successful collection time : 01/27/2010 15:58:05
Next collection time           : 01/27/2010 16:03:05
Collection interval            : 5

Address spaces with the highest increased usage:

Job          Current Change in   Expected Change in   Current
Name ASID      Number of Track      Number of Track      Number of Track
JOB1  0019      Groups Used          Groups Used           Groups Used
JOB2  0014          52                  51                    892
JOB5  000E          30                   35                    735
JOB16 0009          29                   23                    400
                23                   16                    452
    
```

Figure 30. JES spool usage no problem report

Note: In accordance with the IBM Health Checker for z/OS messaging guidelines, the largest generated output length for decimal variable values up to 2147483647 (X'7FFFFFFF') is 10 bytes. When any PFA report value is greater than 2147483647, it displays using multiplier notation with a maximum of six characters. For example, if the report value is 2222233333444445555, PFA displays it as 1973P (2222233333444445555 ÷ 1125899906842) using the following multiplier notation:

Table 11. Multiplier notation used in values for PFA reports		
Name	Sym	Size
Kilo	K	1,024
Mega	M	1,048,576
Giga	G	1,073,741,824
Tera	T	1,099,511,627,776
Peta	P	1,125,899,906,842

The following fields apply to all four reports:

- Last successful model time: The date and time of the last successful model for this check. The predictions on this report were generated at that time.

- **Next model time:** The date and time of the next model. The next model will recalculate the predictions.
- **Model interval:** The value in the configured MODELINT parameter for this check. If PFA determines new prediction calculations are necessary, modeling can occur earlier.
- **Last successful collection time:** The date and time of the last successful data collection for this check.
- **Next collection time:** The date and time of the next collection.
- **Collection interval:** The value in the configured COLLECTINT parameter for this check.
- **Address spaces with the highest increased usage:** The address spaces with the highest recent increase in JES spool usage since the last collection.
- **Address spaces causing the exception:** The address spaces with the highest recent increase in JES spool usage whose increase is greater than expected and caused the exception.
- **Job Name:** The name of the job that has increased usage in the last collection interval.
- **ASID:** The ASID for the job that has increased usage in the last collection interval.
- **Current change in number of track groups used:** The current change in track groups (units of SPOOL space) that a job is using.
- **Expected change in number of track groups used:** The expected change in the number of track groups used by the address space job from one collection to the next.
- **Current number of track groups used:** The current number of track groups being used by the address space. This value is informational only and is not used in the algorithm to detect an abnormal increase in spool usage.

Directories

Note: The content and names for these files and directories are subject to change and cannot be used as programming interfaces; these files are documented only to provide help in diagnosing problems with PFA.

pfa_directory

This directory contains all the PFA check subdirectories and is pointed to by the PFADIR= specified in the /etc/PFA/ini file, or, if you do not specify PFADIR=, the home directory of the PFA user.

pfa_directory/PFA_JES_SPOOL_USAGE/data

The directory for JES spool usage that holds data and modeling results.

Results files

- **systemName.prediction** - This file is generated by the modeling code. It lists the jobs that have the highest recent growth in use of spool usage. For each job in the list, it records the spool usage and additional information required for PFA processing.
- **systemName.prediction.html** - This file contains an .html report version of the data found in the systemName.prediction file.

Data store files:

- **systemName.5day.data** - The data collection file for 5 days of data.
- **systemName.sum** - The summarized data for days 6-30. After 5 days, the data in the 5day.data file is averaged and compressed. The data summarized is deleted from the 5day.data file.

Intermediate files:

- **systemName.data** - The file that is used as input to modeling which contains the address spaces with the highest recent increase in track groups used in the last model.
- **systemName.sumX** - This file is used during the summarization of the 5day.data file.

This directory holds the following log files. Additional information is written to these log files when DEBUG(1).

- stderr - This file only contains data if messages are generated by the JVM.
- stdout - This file only contains data if messages are generated by the JVM.
- systemName.builder.log - The log file generated by intermediate code that builds the files that are input to modeling with details about code execution.
- systemName.launcher.log - The log file generated by launcher code.
- systemNameCONFIG.LOG - The log file containing the configuration history for the last 30 days for this check.
- systemNameCOLLECT.LOG - The log file used during data collection.
- systemNameMODEL.LOG - The log file used during portions of the modeling phase.
- systemNameRUN.LOG - The log file used when the check runs.
- systemName.cart.log - The log file generated by modeling code that contains the execution details of modeling code.
- systemName.tree - This file is generated by the modeling code. It contains information about the model tree that is built based on collected JES spool usage data.

pfa_directory/PFA_JES_SPOOL_USAGE/EXC_timestamp

This directory contains all the relevant data for investigating exceptions issued by this check at the timestamp provided in the directory name. PFA keeps directories only for the last 30 exceptions and only if the exception is newer than the number of days specified in the EXCDIRDAYS parameter. Therefore at each exception, if directories exist that are too old, they will be deleted. In addition, if more than 30 exception directories exist, the oldest directory is deleted so that only 30 exceptions remain after the last exception is issued.

- systemNameREPORT.LOG - The log file containing the same contents as the IBM Health Checker for z/OS report for this exception as well as other diagnostic information issued during report generation.

pfa_directory/PFA_JES_SPOOL_USAGE/config

This directory contains the configuration files for the check.

- EXCLUDED_JOBS - The file containing the list of excluded jobs for this check.
- INCLUDED_JOBS - The file that contains the list of included jobs for this check.

PFA_LOGREC_ARRIVAL_RATE

Description:

The check is looking at the arrival frequency of selected software logrec entries. By monitoring the arrival rate of these logrec entries over time, PFA can detect when the rate of logrec entries exceeds what is considered normal for a system. An unusually high rate of logrec entries can be indicative of recurring failures on the system. PFA can identify when these rates exceed the normal frequency and make an accurate prediction of when you need to take corrective action.

Analyzing the arrival rate by category prevents an expected, normal, but large number of logrecs in the key 8-15 category from masking an unexpected, critical, but small number of logrecs in the key 0 category.

To avoid skewing the logrec arrival rate, PFA ignores the first hour of logrec arrivals after IPL and the last hour of logrec arrivals prior to shutdown.

Tip: z/OS logrec provides two options for the logrec recording medium. Your installation either uses System Logger to produce a logrec log stream or writes logrec to a data set. When a logrec is produced, PFA is notified through an ENF listener. If your installation is set up to write to a data set and that data set fills up, PFA will stop getting notification when a logrec is produced. Therefore, for the best reliability, it is recommended that you use the log stream method with PFA_LOGREC_ARRIVAL_RATE check.

Reason for check:

By detecting recurring failures on the system early, you are able to take corrective action for the failure before a system outage results.

Best practice:

The best practice is to:

1. To determine if the number of software logrec entries is excessive, look at the Logrec Arrival Rate Prediction Report. For an example, see the LOGREC Arrival Rate Prediction Report below.
2. If the number of software logrec entries are excessive, look at the system logrec and try to identify any trend that might exist. For example, see if many logrec entries are associated with a single job (address space) or component. For IBM code, the first three letters of the module name identifies the component. For component identification, see [Identifying modules, components, and products in z/OS MVS Diagnosis: Reference](#).
 - a. If a particular job (address space) is causing the problem, look at the job logs in SDSF.
 - If the job is issuing messages, follow the directions in the message text.
 - b. If the job supports commands to evaluate its status, issue those commands.
 - If the job is responsive, if possible, schedule a recycle for the time that has the lowest business impact.
 - If the job is nonresponsive, if possible, capture diagnostic information and recycle the job.
 - c. Look in SYSLOG for messages issued by the job or about the job.
 - If the job is issuing message, follow the directions in the message text.
3. If a particular component can be identified as causing the problem, follow the standard approach to diagnosing a problem with the component. (For example, have recent changes been made that impact the component?)
 - Look for messages issued by that component in the SYSLOG. If the component is issuing message, follow the directions in the message text.
 - Some components have monitors or operator commands to further evaluate the health of the component.
4. When in a parallel sysplex, you have the option of moving work from the LPAR experiencing the problem (excessive number of logrec entries) to a different LPAR.

z/OS releases the check applies to:

z/OS V1R10 and later.

Type of check:

Remote

Parameters accepted:

Yes, as follows:

Table 12. PFA_LOGREC_ARRIVAL_RATE check parameters

Parameter name	Default value	Minimum Value	Maximum Value	Description
collectint	60 Minutes	1	360	This parameter determines the time (in minutes) to run the data collector that determines the amount of logrec entries. The default is 60 minutes (60).
modelint	720 Minutes	4	1440	This parameter determines how often (in minutes) you want the system to analyze the data and construct a new model. Note that, even when you set a value larger than 360, PFA performs the first model at 360 minutes (6 hours). By default, PFA analyzes the data and constructs a new model every 720 minutes (12 hours).
stddev	2	1	100	This parameter is used to specify how much variance is allowed between the actual logrec arrival rate and the expected logrec arrival rate. It also determines how much variance is allowed across the time range predictions. If you set the STDDEV parameter to a small value, an exception will be issued if the actual logrec arrivals are closer to the expected logrec arrivals and the predictions across the time ranges are consistent. If you set the STDDEV parameter to a larger value, an exception will be issued if the actual logrec arrivals are significantly greater than the expected logrec arrivals even if the predictions across the different time ranges are inconsistent.
collectinactive	1	0 (off)	1 (on)	Defines whether data will be collected and modeled even if the check is not eligible to run (is not ACTIVE(ENABLED)) in IBM Health Checker for z/OS.
exceptionmin	25	0	1000	This parameter defines the minimum logrec arrival rate and the minimum predicted logrec arrival rate required to cause a too high exception.
excdirdays	90	30	365	This parameter specifies the maximum number of days that an exception directory is kept for the check. The maximum number of exception directories kept is 30 regardless of age. Directories are deleted when they are this number of days old.

Table 12. PFA_LOGREC_ARRIVAL_RATE check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
forcemodel	no	00	23:59	<p>This parameter specifies whether a model is forced at a specific time even if the next model time is not reached. The NO value indicates that a model is not forced. Other values must be specified in a 24-hour time format and indicate the time of day to force the model. Modeling in a stable environment then proceeds the MODELINT minutes after this model. If the data becomes unstable for this metric, modeling occurs more often as needed after this time.</p> <p>For example, specifying FORCEMODEL (05:00) MODELINT (780) causes modeling to occur at 05:00 followed by the next model 13 hours later, at 18:00, in a stable environment. When using the values in an unstable environment, modeling occurs at 05:00 and at any point when modeling is needed because of instability. The FORCEMODEL time is not in effect until PFA models one time. PFA must establish a model, without delaying the first model, until the FORCEMODEL time, and PFA must have sufficient data for the first model. Also, if modeling occurred recently, such that there are no new collections at the FORCEMODEL time, the next model time is set to MODELINT minutes from the FORCEMODEL time.</p>
debug	0	0 (off)	1 (on)	This parameter (an integer of 0 or 1) is used at the direction of IBM service to generate additional diagnostic information for the IBM Support Center. This debug parameter is used in place of the IBM Health Checker for z/OS policy. The default is off (0).

To determine the status of the logrec arrival usage check, issue `f pfa,display,check(pfa_logrec_arrival_rate),detail`. For the command example and more details, see [“PFA DISPLAY examples” on page 87](#). The following is an example of the output written to message AIR018I in SDSF user log (ULOG):

```
AIR018I 02:22:54 PFA CHECK DETAIL
CHECK NAME: PFA_LOGREC_ARRIVAL_RATE
ACTIVE : YES
TOTAL COLLECTION COUNT : 5
SUCCESSFUL COLLECTION COUNT : 5
LAST COLLECTION TIME : 04/05/2008 10:18:22
LAST SUCCESSFUL COLLECTION TIME : 04/05/2008 10:18:22
NEXT COLLECTION TIME : 04/05/2008 11:18:22
TOTAL MODEL COUNT : 1
SUCCESSFUL MODEL COUNT : 1
LAST MODEL TIME : 04/05/2008 10:18:24
LAST SUCCESSFUL MODEL TIME : 04/05/2008 10:18:24
NEXT MODEL TIME : 04/05/2008 22:18:24
CHECK SPECIFIC PARAMETERS:
COLLECTINT : 60
MODELINT : 720
COLLECTINACTIVE : 1=ON
DEBUG : 0=OFF
EXCDIRSDAYS : 90
FORCEMODEL : NO
STDDEV : 2
EXCEPTIONMIN : 25
```

User override of IBM values:

The following shows keywords you can use to override check values on either a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. This statement can be copied and modified to override the check defaults:

```
ADDREPLACE POLICY[(policyname)] [STATEMENT(name)]
UPDATE CHECK(IBM PFA,PFA_LOGREC_ARRIVAL_RATE)
ACTIVE
SEVERITY(MEDIUM)
INTERVAL(00:15)
PARMS=( 'COLLECTINT(60)', 'MODELINT(720)', 'STDDEV(2)', 'DEBUG(0)',
        'COLLECTINACTIVE(1)', 'EXCEPTIONMIN(25)', 'EXCDIRDAYS(90)',
        'FORCEMODEL(NO)')
DATE(20140313)
REASON('The reason for the change.')
```

If you change the COLLECTINT parameter or the INTERVAL parameter, the minutes set in the larger of the two parameters must elapse before the reports will be accurate.

Verbose support:

The check does not provide additional details in verbose mode. The check uses the PFA DEBUG parameter (not the DEBUG parameter in IBM Health Checker for z/OS) for writing additional information to its log files for both exception and non-exception states.

Debug support:

The DEBUG parameter in IBM Health Checker for z/OS is ignored by this check. Rather, the debug parameter is a PFA check specific parameter. For details, see [“Understanding how to modify PFA checks”](#) on page 83.

Reference:

For more information about PFA, see the topic on [“Overview of Predictive Failure Analysis”](#) on page 73.

Messages:

This check issues the following exception messages:

- AIRH110E

For additional message information, see the topics on:

- [AIRH messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).
- [AIR messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).

SECLABEL recommended for MLS users:

SYSLOW

Output:

The LOGREC Arrival Rate Prediction Report:

LOGREC Arrival Rate Prediction Report			
Last successful model time	:	11/06/2008 17:32:44	
Next model time	:	11/06/2008 23:33:44	
Model interval	:	360	
Last successful collection time	:	11/06/2008 18:33:49	
Next collection time	:	11/06/2008 19:34:49	
Collection interval	:	60	
		Key 0	Key 1-7
		-----	-----
Arrivals in last collection interval:		1	0
Predicted rates based on...			
1 hour of data:		1	0
24 hours of data:		0	0
7 days of data:		0	0
30 days of data:		0	0
Jobs having LOGREC arrivals in last collection interval:			
Job Name	ASID	Arrivals	
LOGREC08	0029	2	
LOGREC00	0027	1	

Figure 31. LOGREC arrival rate prediction report

Note: In accordance with the IBM Health Checker for z/OS messaging guidelines, the largest generated output length for decimal variable values up to 2147483647 (X'7FFFFFFF') is 10 bytes. When any PFA report value is greater than 2147483647, it displays using multiplier notation with a maximum of six characters. For example, if the report value is 2222233333444445555, PFA displays it as 1973P (2222233333444445555 ÷ 1125899906842) using the following multiplier notation:

Table 13. Multiplier notation used in values for PFA reports		
Name	Sym	Size
Kilo	K	1,024
Mega	M	1,048,576
Giga	G	1,073,741,824
Tera	T	1,099,511,627,776
Peta	P	1,125,899,906,842

- Last successful model time: The date and time of the last successful model for this check. The predictions on this report were generated at that time.
- Next model time: The date and time of the next model. The next model will recalculate the predictions.
- Model interval: The value in the configured MODELINT parameter for this check. If PFA determines new prediction calculations are necessary, modeling can occur earlier.
- Last successful collection time: The date and time of the last successful data collection for this check.
- Next collection time: The date and time of the next collection.
- Collection interval: The value in the configured COLLECTINT parameter for this check.
- Key column headings: The program key of the issuer of the logrec entry.
- Arrivals in last collection interval: The value is one of the following:

nnnn

The actual count of the logrec entries received in the last collection interval period.

The number of arrivals in the last collection interval is not available because the system generated a very large number of logrec entries during this period. For example:

```

Arrivals in last      Key 0      Key 1-7      Key 8-15
collection interval:  -----      -----      -----
                        ****              ****              ****
    
```

If the arrival count is unavailable (****), the following message also appears on the report immediately following the predicted rates:

```

**** Arrivals in last collection interval unavailable
due to high level of system activity.
    
```

In addition, if the check was unable to report all the logrec entries found in the last collection interval, it can also display the following message before the job listings:

```

Job list may not include all arrivals in last collection
interval due to high level of system activity.
    
```

This message typically appears when the arrival count is unavailable (with ****) however, it can also appear when a logrec entry is skipped because of a high level of system activity

- **Predicted rates based on...:** The logrec entry counts based on different ranges of historical data. If the required amount of data is not available, the line in this report is not generated. For example, if PFA has been running for a week, it has not yet collected 30 days of historical data and therefore the "30 days of data" line is not generated in the report.
- **Jobs having LOGREC arrivals in last collection interval:** The jobs that contributed to the logrec arrivals in the last collection interval.
- **Job name:** The name of the job that had logrec arrivals in the last collection interval.
- **ASID:** The ASID of the job that had logrec arrivals in the last collection interval.
- **Arrivals:** The actual count of the logrec arrivals for the job.

Directories:

When you install the PFA_LOGREC_ARRIVAL_RATE check, the shell script creates the following directories:

pfa_directory

This directory contains all the PFA check subdirectories and is pointed to by the PFADIR= specified in the /etc/PFA/ini file, or, if you do not specify PFADIR=, the home directory of the PFA user.

pfa_directory/PFA_LOGREC_ARRIVAL_RATE/data/

The directory for logrec arrival rate that holds data and modeling results.

Guideline: If the use of the z/OS image is radically different after an IPL (for instance, the change from a test system to a production system), delete the files from PFA_LOGREC_ARRIVAL_RATE/data directory to enable the check to collect the most accurate modeling information. After the PFA_LOGREC_ARRIVAL_RATE check finds a problem and you correct it, delete the files in PFA_LOGREC_ARRIVAL_RATE/data to prevent that previous error from hiding any new errors.

Results files:

- **systemName.1hr.prediction** - This file is generated by the modeling code for the predictions made for one hour of historical data. It contains the list of program keys with their corresponding predictions and additional information required for PFA processing.
- **systemName.24hr.prediction** - This file is generated by the modeling code for the predictions made for 24 hours of historical data. It contains the list of program keys with their corresponding predictions and additional information required for PFA processing.

- `systemName.7day.prediction` - This file is generated by the modeling code for the predictions made for seven days of historical data. It contains the list of program keys with their corresponding predictions and additional information required for PFA processing.
- `systemName.30day.prediction` - This file is generated by the modeling code for the predictions made for 30 days of historical data. It contains the list of program keys with their corresponding predictions and additional information required for PFA processing.
- `systemName.1hr.prediction.html` - This file contains an .html report version of the data found in the `systemName.1hr.prediction` file.
- `systemName.24hr.prediction.html` - This file contains an .html report version of the data found in the `systemName.24hr.prediction` file.
- `systemName.7day.prediction.html` - This file contains an .html report version of the data found in the `systemName.7day.prediction` file.
- `systemName.30day.prediction.html` - This file contains an .html report version of the data found in the `systemName.30day.prediction` file.
- `systemName.prediction.stddev` - The file generated by the modeling code to list the standard deviation of the predictions across the time ranges for each program key.

The values for the logrec arrival prediction file in .html (prediction.html) files are as follows:

- **Program Key:** The program key when the error described by the logrec entry occurred.
- **Instance Count:** The number of records with this program key that were factored into the prediction model.
- **Current Logrec Arrivals:** The number of logrec arrivals for this program key in the last collection interval included in this model.
- **Prediction Look Forward Seconds:** The number of seconds the prediction should project into the future.
- **Predicted Logrec Arrivals:** The predicted logrec arrivals for this program key.

Intermediate files:

- `systemName.OUT` - The data collection file.
- `systemName.1hr.data` - The file used as input to modeling code. It contains one hour of historical data.
- `systemName.24hr.data` - The file used as input to modeling code. It contains 24 hours of historical data.
- `systemName.7day.data` - The file used as input to modeling code. It contains seven days of historical data.
- `systemName.30day.data` - The file is used as input to modeling code. It contains 30 days of historical data.
- `systemName.data --` The file is used as input to the modeling to track if enough data is available to model.

This directory holds the following log files. Additional information is written to these log files when `DEBUG(1)`.

- `stderr` - This file only contains data if messages are generated by the JVM.
- `stdout` - This file only contains data if messages are generated by the JVM.
- `systemName.1hr.cart.log` - The log file generated by modeling code with details about code execution while one hour of historical data was being modeled.
- `systemName.24hr.cart.log` - The log file generated by modeling code with details about code execution while 24 hours of historical data was being modeled.
- `systemName.7day.cart.log` - The log file generated by modeling code with details about code execution while seven days of historical data was being modeled.

- `systemName.30day.cart.log` - The log file generated by modeling code with details about code execution while 30 days of historical data was being modeled.
- `systemName.builder.log` - The log file generated by intermediate code that builds the files that are input to modeling with details about code execution.
- `systemName.launcher.log` - The log file generated by launcher code.
- `systemName.1hr.tree` - This file is generated by the modeling code. It contains information about the model tree which was built based on the last one hour of collected data.
- `systemName.24hr.tree` - This file is generated by the modeling code. It contains information about the model tree which was built based on the last 24 hours of collected data.
- `systemName.7day.tree` - This file is generated by the modeling code. It contains information about the model tree which was built based on the last seven days of collected data.
- `systemName.30day.tree` - This file is generated by the modeling code. It contains information about the model tree which was built based on the last 30 days of collected data.
- `systemName.1hr.holes` - The file is used to track gaps in data for a one hour period. Gaps are caused by stopping PFA or by an IPL.
- `systemName.24hr.holes` - The file is used to track gaps in the data for a 24 hour time period. Gaps are caused by stopping PFA or by an IPL.
- `systemName.7day.holes` - The file is used to track gaps in the data for the seven day time period. Gaps are caused by stopping PFA or by an IPL.
- `systemName.30day.holes` - The file is used to track gaps in the data for the 30 day time period. Gaps are caused by stopping PFA or by an IPL.
- `systemNameCONFIG.LOG` - The log file containing the configuration history for the last 30 days for this check.
- `systemNameCOLLECT.LOG` - The log file used during data collection.
- `systemNameMODEL.LOG` - The log file used during portions of the modeling phase.
- `systemNameRUN.LOG` - The log file used when the check runs.

pfa_directory/PFA_LOGREC_ARRIVAL_RATE/EXC_timestamp

This directory contains all the relevant data for investigating exceptions issued by this check at the timestamp provided in the directory name. PFA keeps directories only for the last 30 exceptions and only if the exception is newer than the number of days specified in the EXCDIRDAYS parameter. Therefore at each exception, if directories exist that are too old, they will be deleted. In addition, if more than 30 exception directories exist, the oldest directory is deleted so that only 30 exceptions remain after the last exception is issued.

- `systemNameREPORT.LOG` - The log file containing the same contents as the IBM Health Checker for z/OS report for this exception as well as other diagnostic information issued during report generation.

pfa_directory/PFA_LOGREC_ARRIVAL_RATE/config

This directory contains configuration information for the check.

- `EXCLUDED_JOBS` - The file containing the list of excluded jobs for this check.

PFA_MESSAGE_ARRIVAL_RATE

Description:

This check is determining when an LPAR is damaged by checking for an abnormal arrival rate of messages per amount of CPU used.

To avoid skewing the message arrival rate, PFA ignores the first hour of message data after IPL and the last hour of message data before shutdown. In addition, PFA attempts to track the same persistent jobs that it tracked before IPL or PFA restart if the same persistent jobs are still active. (The same persistent jobs must still be active for PFA to track.)

This check is not designed to detect performance problems that are caused by insufficient resources, faulty WLM policy, or spikes in work. However, it might help to determine if a performance problem detected by a performance monitor or WLM is caused by a damaged system.

The message arrival rate check issues an exception using four types of comparisons, which are described in more detail in the next section:

- top persistent jobs
- other persistent
- non-persistent jobs
- total system

After PFA issues an exception, the next comparison type is not performed. The CONSOLE address space and any jobs that match the job and system combinations defined in the config/EXCLUDED_JOBS file that have been read for PFA processing are not included in the processing for any of the four types of comparisons. By default, an EXCLUDED_JOBS file containing the all address spaces that match JES* on all systems is created during installation. Therefore, if you have not made any modifications to the EXCLUDED_JOBS file, these jobs will be excluded. See [“Configure supervised learning” on page 91](#) for more information.

Top persistent jobs

PFA tracks the top persistent jobs individually. Jobs are considered persistent if they start within an hour after IPL. PFA determines which jobs to track individually based on the following criteria:

- If PFA previously ran on this system and the same 10 jobs that were previously tracked are active, PFA tracks the same jobs.
- If PFA did not previously run on this system or the same jobs previously tracked are not all active, PFA collects data for a period of time to use in determining which jobs have the highest arrival rates. After this time passes, PFA individually tracks the jobs that have the highest arrival rates for that period.
 - During the first hour after IPL and during the time PFA is determining the jobs to track individually, normal data collection and modeling are suspended.
 - Changing the COLLECTINT or the MODELINT parameters during these times is allowed, but the changes are not used until after these times have passed.
 - Next collection and model times change automatically during these times to reflect the most accurate times known at each phase of the initial processing.

This top persistent jobs comparison is performed to determine if the message rate is higher than expected or lower than expected.

Other persistent jobs

The persistent jobs that PFA does not track individually are the other persistent jobs. PFA generates the predictions using the totals for this group. When determining if the message arrival rate is higher than expected, PFA performs the comparisons individually using a mathematical formula. When determining if the message rate is lower than expected, PFA performs the comparisons using the totals for the group.

Non-persistent jobs

The jobs that start over an hour after IPL are the non-persistent jobs. PFA performs the predictions and the comparisons using the totals for this group. This type of comparison is only used to determine if the message rate is higher than expected.

Total system

This group includes all jobs. PFA performs the predictions and the comparisons for the entire system to determine if the message rate is higher than expected or lower than expected.

Reason for check:

The objective of this check is to determine if an LPAR is damaged and if an address space or partition is hung by checking for an abnormal arrival rate of messages per amount of CPU used.

Best practice:

If PFA detects an unexpectedly high amount of messages, the best practice is to analyze the messages being sent by the address spaces identified on the report by examining the system log to determine what is causing this burst of message activity. Establish which messages were issued around the time of the activity and review the message details. Follow the directions provided by the message to continue to diagnose and fix the problem.

If PFA detects an unexpectedly low number of messages, examine the report in SDSF for details about why the exception was issued. Use the Runtime Diagnostics output in the report to assist you in diagnosing and fixing the problem. For more information about Runtime Diagnostics see [Chapter 5, "Runtime Diagnostics,"](#) on page 35.

z/OS releases the check applies to:

z/OS V1R11 and later.

Type of check:

Remote

Parameters accepted:

Yes, as follows:

Table 14. PFA_MESSAGE_ARRIVAL_RATE check parameters

Parameter name	Default value	Minimum Value	Maximum Value	Description
collectint	15 Minutes	15	360	This parameter determines how often (in minutes) to run the data collector that retrieves the current message arrival rate.
modelint	720 Minutes	60	1440	This parameter determines how often (in minutes) you want the system to analyze the data and construct a new message arrival rate model or prediction. By default, PFA analyzes the data and constructs a new model every "default value" minutes. The model interval must be at least four times larger than the collection interval. Note that, even when you set a value larger than 360, PFA performs the first model at 360 minutes (6 hours). By default, PFA analyzes the data and constructs a new model every 720 minutes (12 hours).

Table 14. PFA_MESSAGE_ARRIVAL_RATE check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
stddev	10	2	100	This parameter is used to specify how much variance is allowed between the actual message arrival rate per amount of CPU and the expected message arrival rate. It is used when determining if the actual message arrival rate has increased beyond the allowable upper limit. It also determines how much variance is allowed across the time range predictions. If you set the STDDEV parameter to a smaller value, an exception is issued if the actual message arrival rate is closer to the expected message arrival rate and the predictions across the time ranges are consistent. If you set the STDDEV parameter to a larger value, an exception is issued if the actual message arrival rate is significantly greater than the expected message arrival rate even if the predictions across the different time ranges are inconsistent.
collectinactive	1 (on)	0 (off)	1 (on)	Defines whether data is collected and modeled even if the check is not eligible to run, not ACTIVE(ENABLED), in IBM Health Checker for z/OS.
trackedmin	3	0	1000	This parameter defines the minimum message arrival rate required for a persistent job for it to be considered a top persistent job that must be tracked individually.
exceptionmin	1	0	1000	This parameter is used to determine if an exception is issued for an unexpectedly high message arrival rate. For tracked jobs and other persistent jobs, this parameter defines the minimum message arrival rate and the minimum predicted message arrival rate required to cause a too high exception. For non-persistent jobs and the total system comparisons, this parameter defines the minimum message arrival rate required to cause a too high exception.
checklow	1 (on)	0 (off)	1	Defines whether Runtime Diagnostics is run to validate that the absence of messages is caused by a problem. If this value is off, exceptions are not issued for conditions in which the message arrival rate is unexpectedly low.

Table 14. PFA_MESSAGE_ARRIVAL_RATE check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
stddevlow	2	2	100	<p>This parameter is used to specify how much variance is allowed between the actual message arrival rate per amount of CPU and the expected message arrival rate when determining if the actual rate is unexpectedly low.</p> <ul style="list-style-type: none"> • If you set the STDDEVLOW parameter to a smaller value, an exception is issued when the actual message arrival rate is closer to the expected message arrival rate. • If you set the STDDEVLOW parameter to a larger value, an exception is issued when the actual message arrival rate is significantly lower than the expected message arrival rate.
limitlow	3	1	100	This parameter defines the maximum message arrival rate allowed when issuing an exception for an unexpectedly low number of messages.
excdirdays	90	30	365	This parameter specifies the maximum number of days that an exception directory is kept for the check. The maximum number of exception directories kept is 30 regardless of age. Directories are deleted when they are this number of days old.
forcemodel	no	00	23:59	<p>This parameter specifies whether a model is forced at a specific time even if the next model time is not reached. The NO value indicates that a model is not forced. Other values must be specified in a 24-hour time format and indicate the time of day to force the model. Modeling in a stable environment then proceeds the MODELINT minutes after this model. If the data becomes unstable for this metric, modeling occurs more often as needed after this time.</p> <p>For example, specifying FORCEMODEL (05:00) MODELINT (780) causes modeling to occur at 05:00 followed by the next model 13 hours later, at 18:00, in a stable environment. When using the values in an unstable environment, modeling occurs at 05:00 and at any point when modeling is needed because of instability. The FORCEMODEL time is not in effect until PFA models one time. PFA must establish a model, without delaying the first model, until the FORCEMODEL time, and PFA must have sufficient data for the first model. Also, if modeling occurred recently, such that there are no new collections at the FORCEMODEL time, the next model time is set to MODELINT minutes from the FORCEMODEL time.</p>
debug	0 (off)	0 (off)	1 (on)	The debug parameter (an integer of 0 or 1) is used at the direction of IBM service to generate additional diagnostic information for the IBM Support Center. This debug parameter is used in place of the IBM Health Checker for z/OS policy. The default is off (0).

To determine the status of the message arrival rate check, issue `f pfa,display,check(pfa_message_arrival_rate),detail`. For the command example and more details, see [“PFA DISPLAY examples”](#) on page 87. The following example shows the output written to message AIR018I in SDSF:

```
AIR018I 02:22:54 PFA CHECK DETAIL
CHECK NAME: PFA_MESSAGE_ARRIVAL_RATE
ACTIVE : YES
TOTAL COLLECTION COUNT : 5
SUCCESSFUL COLLECTION COUNT : 5
LAST COLLECTION TIME : 02/05/2009 10:18:22
LAST SUCCESSFUL COLLECTION TIME: 02/05/2009 10:18:22
NEXT COLLECTION TIME : 02/05/2009 10:33:22
TOTAL MODEL COUNT : 1
SUCCESSFUL MODEL COUNT : 1
LAST MODEL TIME : 02/05/2009 10:18:24
LAST SUCCESSFUL MODEL TIME : 02/05/2009 10:18:24
NEXT MODEL TIME : 02/05/2009 22:18:24
CHECK SPECIFIC PARAMETERS:
COLLECTINT : 15
MODELINT : 720
COLLECTINACTIVE : 1=ON
DEBUG : 0=OFF
EXCDIRSDAYS : 90
FORCEMODEL : NO
STDDEV : 10
TRACKEDMIN : 3
EXCEPTIONMIN : 1
CHECKLOW : 1=ON
STDDEVLOW : 2
LIMITLOW : 3
EXCLUDED JOBS:
NAME SYSTEM DATE ADDED REASON ADDED
JES * 2010/03/31 00:00 Exclude JES* jobs on ALL.
INCLUDED JOBS:
NAME SYSTEM DATE ADDED REASON ADDED
CRIT1 SY1 03/11/2015 15:15 Include job CRIT1 on SY1.
```

User override of IBM values:

The following shows keywords you can use to override check values on either a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. This statement can be copied and modified to override the check defaults:

```
ADDREPLACE POLICY[(policyname)] [STATEMENT(name)]
UPDATE CHECK(IBM PFA,PFA_MESSAGE_ARRIVAL_RATE)
ACTIVE
SEVERITY(MEDIUM)
INTERVAL(ONETIME)
PARMS=('COLLECTINT(15)', 'MODELINT(720)', 'STDDEV(10)', 'DEBUG(0)',
'COLLECTINACTIVE(1)', 'EXCEPTIONMIN(1)', 'TRACKEDMIN(3)',
'CHECKLOW(1)', 'STDDEVLOW(2)', 'LIMITLOW(3)',
'EXCDIRSDAYS(90)', 'FORCEMODEL(NO)')
DATE(20140313)
REASON('The reason for the change.')
```

The message arrival rate check is designed to run automatically after every data collection. Do not change the INTERVAL parameter.

Verbose support:

The check does not provide additional details in verbose mode. The check uses the PFA DEBUG parameter (not the DEBUG parameter in IBM Health Checker for z/OS) for writing additional information to its log files for both exception and non-exception states.

Debug support:

The DEBUG parameter in IBM Health Checker for z/OS is ignored by this check. Rather, the debug parameter is a PFA check specific parameter. For details, see [“Understanding how to modify PFA checks”](#) on page 83.

Reference:

For more information about PFA, see the topic on [“Overview of Predictive Failure Analysis”](#) on page 73.

Messages:

The output is a message arrival rate prediction report that corresponds to the message issued. One of the following reports is generated:

- AIRH152E or AIRH153E – total system exception report
- AIRH165E or AIRH206E – tracked jobs exception report
- AIRH166E or AIRH207E – other persistent jobs exception report
- AIRH169E – other non-persistent jobs exception report

For additional message information, see the topics on:

- [AIRH messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).
- [AIR messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).

SECLABEL recommended for MLS users:

SYSLOW

Output:

The output is a variation of the message arrival rate prediction report. The values found in the message arrival prediction report are as follows:

Tracked top persistent jobs exception report:

PFA issues the message arrival rate tracked jobs exception report when any one or more tracked, persistent jobs cause an exception. The exception can be caused by a higher than expected message arrival rate or a lower than expected message arrival rate. Only the tracked jobs that caused the exception are included in the list of jobs on the report. If the report was generated due to a lower than expected message arrival rate, the report includes Runtime Diagnostics output, which can help you diagnose the behavior. The following example is the message arrival rate tracked jobs exception report for jobs that had a higher than expected message arrival rate (for AIRH165E):

```

Message Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 17:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time  : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval             : 15

Highest STDDEV needed to avoid this exception:                10

Persistent address spaces with high rates:

Job Name      ASID      Message
                Arrival
                Rate      1 Hour      24 Hour      7 Day
-----
TRACKED1 001D  ----- 158.00  ----- 23.88  ----- 22.82  ----- 15.82
TRACKED3 0029  ----- 11.00  ----- 12.43  ----- 2.36  ----- 8.36
    
```

Figure 32. Message arrival rate prediction report: tracked jobs higher than expected

The following example shows the message arrival rate prediction report for jobs that had a lower than expected message arrival rate (for AIRH206E):

```

Message Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 17:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time  : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval              : 15
Highest STDDEVLOW needed to avoid this exception:
21

Persistent address spaces with low rates:

Job Name      ASID      Enqueue Request Rate      Predicted Enqueue Request Rate
              ASID      Rate                      1 Hour      24 Hour      7 Day
-----
JOBS4      0027      1.17      23.88      22.82      15.82
JOBS5      002E      2.01      8.34      11.11      12.11

Runtime Diagnostics Output:
Runtime Diagnostics detected a problem in job: JOBS4
EVENT 06: HIGH - HIGHCPU - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 0027 JOBNAME: JOBS4
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 07: HIGH - LOOP - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID: 0027 JOBNAME: JOBS4 TCB: 004E6850
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE APPEARS TO BE IN A LOOP.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
Runtime Diagnostics detected a problem in job: JOBS5
EVENT 03: HIGH - HIGHCPU - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 0027 JOBNAME: JOBS5
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 04: HIGH - LOOP - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID: 0027 JOBNAME: JOBS5 TCB: 004E6850
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE APPEARS TO BE IN A LOOP.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----

```

Figure 33. Message arrival rate prediction report: tracked jobs lower than expected

Other persistent jobs exception report:

PFA issues the message arrival rate other persistent jobs exception report when a comparison of a persistent job (that is not being individually tracked) causes an exception when compared to the totals of the other persistent jobs. The exception can be caused by a higher than expected message arrival rate or a lower than expected message arrival rate. The predictions listed on this report are the predicted rates for the total other persistent jobs group. The list of jobs is only those persistent jobs (not tracked individually) that have a problem and is only generated for a higher than expected message arrival rate. No predictions are given for these jobs because PFA does not model individual predictions for jobs that are not tracked individually. If there is more than one job with the same name, four asterisks **** are printed for the ASID in the report. If the report was

generated due to a lower than expected message arrival rate, the report will include Runtime Diagnostics output which can help diagnose the behavior. The following example is the message arrival rate exception report for other persistent jobs that had an unexpectedly high message arrival rate (for AIRH166E):

```

Message Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 17:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval              : 15

Other persistent jobs group:
Message arrival rate
  in last collection interval   : 502.72
Prediction based on 1 hour of data : 20.27
Prediction based on 24 hours of data : 27.98
Prediction based on 7 days of data  : 31.22

Highest STDEVLOW needed to avoid this exception:          25

Persistent address spaces with high rates:

Job Name      ASID      Message
              ASID      Arrival
              ASID      Rate
-----
PERS1      001E      83.22
PERS2      0038      75.52
PERS3      0039      47.47
    
```

Figure 34. Message arrival rate prediction report: other persistent jobs with high arrival rate

The following example is the message arrival rate exception report for other persistent jobs that had an unexpectedly low message arrival rate (for AIRH207E):

```

Message Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 17:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time  : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval              : 15

Other persistent jobs group:
Message arrival rate
  in last collection interval    : 1.13
Prediction based on 1 hour of data : 20.27
Prediction based on 24 hours of data : 27.98
Prediction based on 7 days of data  : 31.22

Highest STDEVLOW needed to avoid this exception:          28

Runtime Diagnostics Output:

EVENT 01: HIGH - HIGHCPU - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 0027 JOBNAME: PERS4
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 02: HIGH - LOOP - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID: 0027 JOBNAME: PERS4 TCB: 004E6850
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE APPEARS TO BE IN A LOOP.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----

```

Figure 35. Message arrival rate prediction report: other persistent jobs with low arrival rate

Non-persistent jobs exception report:

PFA issues the message arrival rate non-persistent jobs exception report when the non-persistent jobs as a group can cause an exception. This exception is only issued for a higher than expected message arrival rate. The message arrival rate and predictions listed on this report are calculated for the total non-persistent jobs group. The list of jobs contains only three non-persistent jobs that have high arrival counts. No predictions are given for these jobs because PFA does not model

individual predictions for jobs that are not tracked individually. The following example is the message arrival rate non-persistent jobs exception report (for AIRH169E):

```

Message Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 17:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time : 01/27/2009 17:41:38
Next collection time           : 01/27/2009 17:56:38
Collection interval             : 15

Highest STDEV needed to avoid this exception:          14

Non-persistent jobs group:
Message arrival rate
  in last collection interval      : 265.49
Prediction based on 1 hour of data : 20.27
Prediction based on 24 hours of data : 27.98
Prediction based on 7 days of data  : 31.22

Address spaces with high arrivals:

Job Name      ASID      Message
              ASID      Arrival
              ASID      Counts
-----
NONPERS1 001F      83
NONPERS2 0048      52
NONPERS3 0049      47
    
```

Figure 36. Message arrival rate prediction report: non-persistent jobs higher than expected

No problem and total system exception report:

PFA issues the message arrival rate system report when no exception is issued or when the total message arrival rate exception is issued. When there is no problem or when an exception occurs due to a higher than expected message arrival rate, the list of jobs contains all of the jobs being tracked individually. The Runtime Diagnostics section is written in the report when the exception is issued due to a lower than expected message arrival rate. When the report is issued due to a lower than expected message arrival rate, the list of tracked jobs will not be printed on the report. The following example is the message arrival rate no problem total system report issued due to a higher than expected message arrival rate (AIRH152E):

```

Message Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 17:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time : 01/27/2009 17:41:38
Next collection time           : 01/27/2009 17:56:38
Collection interval             : 15

Message arrival rate
  at last collection interval    : 1083.52
Prediction based on 1 hour of data : 98.27
Prediction based on 24 hours of data : 85.98
Prediction based on 7 days of data  : 100.22

Highest STDEV needed to avoid this exception:          13

Top persistent users:

Job Name      ASID      Message
              ASID      Arrival
              ASID      Rate
              ASID      1 Hour
              ASID      24 Hour
              ASID      7 Day
-----
JOB1 001D      58.00      23.88      22.82      15.82
JOB2 0028      11.00      0.34      11.11      12.11
JOB3 0029      11.00      12.43      2.36      8.36
...
    
```

Figure 37. Message arrival rate prediction report: total system higher than expected

This example is the message arrival rate total system exception report issued due to a lower than expected message arrival rate (for AIRH153E) with Runtime Diagnostic output:

```

Message Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 17:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time  : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval              : 15

Message arrival rate
in last collection interval     : 2.02
Prediction based on 1 hour of data : 98.27
Prediction based on 24 hours of data : 85.98
Prediction based on 7 days of data  : 100.22

Highest STDEVLOW needed to avoid this exception:          50

Runtime Diagnostics Output:

EVENT 01: HIGH - CF - SYSTEM: SY1 2009/06/12 - 13:23:25
IXL013I IXLCONN REQUEST FOR STRUCTURE SYSZWLM_WORKUNIT FAILED.
JOBNAME: WLM ASID: 000A CONNECTOR NAME: #SY1
IXLCONN RETURN CODE: 0000000C, REASON CODE: 02010C05
STRUCTURE NOT DEFINED IN THE CFRM ACTIVE POLICY
CONADIAG0: 00000002
-----
EVENT 02: HIGH - XCF - SYSTEM: SY1 2009/06/12 - 13:23:55
IXC467I STOPPING PATHIN STRUCTURE IXCT_SIGNAL
RSN: START REQUEST FAILED
-----
EVENT 03: HIGH - XCF - SYSTEM: SY1 2009/06/12 - 13:23:58
IXC467I STOPPING PATHIN STRUCTURE IXCT_SIGNAL
RSN: START REQUEST FAILED
-----
EVENT 04: HIGH - CF - SYSTEM: SY1 2009/06/12 - 13:24:02
IXL013I IXLCONN REQUEST FOR STRUCTURE IXCT_SIGNAL FAILED.
JOBNAME: XCFAS ASID: 0006 CONNECTOR NAME: SIGPATH_01000008
IXLCONN RETURN CODE: 0000000C, REASON CODE: 02010C05
STRUCTURE NOT DEFINED IN THE CFRM ACTIVE POLICY
CONADIAG0: 00000002
-----
EVENT 05: HIGH - XCF - SYSTEM: SY1 2009/06/12 - 13:24:02
IXC467I STOPPING PATHIN STRUCTURE IXCT_SIGNAL
RSN: START REQUEST FAILED
-----
EVENT 06: HIGH - HIGHCPU - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 0027 JOBNAME: DAVIDZ
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 07: HIGH - LOOP - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID: 0027 JOBNAME: DAVIDZ TCB: 004E6850
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE APPEARS TO BE IN A LOOP.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----

```

Figure 38. Message arrival rate prediction report: total system lower than expected

Note: In accordance with the IBM Health Checker for z/OS messaging guidelines, the largest generated output length for decimal variable values up to 2147483647 (X'7FFFFFFF') is 10 bytes. When any PFA report value is greater than 2147483647, it displays using multiplier notation with a maximum of six characters. For example, if the report value is 222233333444445555, PFA displays it as 1973P (222233333444445555 ÷ 1125899906842) using the following multiplier notation:

<i>Table 15. Multiplier notation used in values for PFA reports</i>		
Name	Sym	Size
Kilo	K	1,024
Mega	M	1,048,576
Giga	G	1,073,741,824
Tera	T	1,099,511,627,776
Peta	P	1,125,899,906,842

The following fields apply to all reports:

- **Last successful model time:** The date and time of the last successful model for this check. The predictions on this report were generated at that time.
 - **Next model time:** The date and time of the next model. The next model will recalculate the predictions.
 - **Model interval:** The value in the configured MODELINT parameter for this check. If PFA determines new prediction calculations are necessary, modeling can occur earlier.
 - **Last successful collection time:** The date and time of the last successful data collection for this check.
 - **Next collection time:** The date and time of the next collection.
 - **Collection interval:** The value in the configured COLLECTINT parameter for this check.
 - **Message arrival rate in last collection interval:** The actual message arrival rate in the last collection interval where the rate is defined to be the number of messages divided by the amount of CPU used.
 - **Predicted rates based on...:** The message arrival rates based on one hour, 24 hours, and seven days. If no prediction is available for a given time range, the line is not printed. For example, if the check has been running for two days, seven days of data is not available and the "Prediction based on 7 days of data" line is not printed.
 - **Highest STDDEV needed to avoid this exception:** The highest STDDEV required to avoid the exception issued for all resources in exception status with the values in effect at this comparison.
 - **Highest STDDEVLOW needed to avoid this exception:** The highest STDDEVLOW required to avoid the exception issued for all resources in exception status with the values in effect at this comparison.
 - **TOO HIGH:** A STDDEV or STDDEVLOW needed to avoid the exception is TOO HIGH if the parameter value required is higher than the maximum allowed.
 - **IRRELEVANT:** A STDDEV or STDDEVLOW needed to avoid the exception is IRRELEVANT if the parameter's value was not relevant in the comparison algorithm used for the values in the current comparison.
 - **Runtime Diagnostics Output:** Runtime Diagnostics event records to assist you in diagnosing and fixing the problem. See the topic on "[Specific Runtime Diagnostics event reports](#)" on page 41 in Chapter 5, "[Runtime Diagnostics](#)," on page 35.
 - **Job Name:** The name of the job that has message arrivals in the last collection interval.
 - **ASID:** The ASID for the job that has message arrivals in the last collection interval.
 - **Message Arrival Rate:** The current message arrival rate for the persistent job.
 - **Message Arrival Counts:** The message arrival count for the non-persistent job.
- Note:** The "Message Arrival Count" field is unique to the non-persistent jobs exception report.
- **Predicted Message Arrival Rate:** The predicted message arrival rate based on one hour, 24 hours, and seven days of data. If PFA did not previously run on this system or the same jobs

previously tracked are not all active, there will not be enough data for a time range until that amount of time has passed. Also, gaps in the data caused by stopping PFA or by an IPL might cause the time range to not have enough data available. After the check collects enough data for any time range, predictions are made again for that time range. If there is not enough data for a time range, INELIGIBLE is printed and comparisons are not made for that time range.

Directories

Note: The content and names for these files and directories are subject to change and cannot be used as programming interfaces; these files are documented only to provide help in diagnosing problems with PFA.

pfa_directory

This directory contains all the PFA check subdirectories and is pointed to by the PFADIR= specified in the /etc/PFA/ini file, or, if you do not specify PFADIR=, the home directory of the PFA user.

pfa_directory/PFA_MESSAGE_ARRIVAL_RATE/data

The directory for message arrival rate that holds data and modeling results.

Guideline: If the use of the z/OS image is radically different after an IPL (for instance, the change from a test system to a production system), delete the files in the *PFA_MESSAGE_ARRIVAL_RATE/data* directory to enable the check to collect the most accurate modeling information.

Results files

- *systemName.1hr.prediction* - This file is generated by the modeling code for the predictions made for one hour of historical data. It contains predictions for each of the tracked address spaces, the other persistent category, the non-persistent category, and the total system category. It also contains additional information required for PFA processing.
- *systemName.24hr.prediction* - This file is generated by the modeling code for the predictions made for 24 hours of historical data. It contains predictions for each of the tracked address spaces, the other persistent category, the non-persistent category, and the total system category. It also contains additional information required for PFA processing.
- *systemName.7day.prediction* - This file is generated by the modeling code for the predictions made for seven days of historical data. It contains predictions for each of the tracked address spaces, the other persistent category, the non-persistent category, and the total system category. It also contains additional information required for PFA processing.
- *systemName.1hr.prediction.html* - This file lists the persistent address spaces in an .html report format for the predictions made for one hour of historical data.
- *systemName.24hr.prediction.html* - This file lists the persistent address spaces in an .html report format for the predictions made for 24 hours of historical data.
- *systemName.7day.prediction.html* - This file lists the persistent address spaces in an .html report format for the predictions made for seven days of historical data.
- *systemName.prediction.stddev* - The file generated by the modeling code to list the standard deviation of the predictions across the time ranges for each address space.

Data store files:

- *systemName.OUT* - The data collection file.

Intermediate files:

- *systemName.data* - The file is used as input to the modeling to track if enough data is available to model.
- *systemName.1hr.data* - The file used as input to modeling code. It contains one hour of historical data.
- *systemName.24hr.data* - The file used as input to modeling code. It contains 24 hours of historical data.

- `systemName.7day.data` - The file used as input to modeling code. It contains seven days of historical data.
- `systemName.1hr.holes` - The file is used to track gaps in the data for a one hour time period. Gaps are caused by stopping PFA or by an IPL.
- `systemName.24hr.holes` - The file is used to track gaps in the data for a 24 hour time period. Gaps are caused by stopping PFA or by an IPL.
- `systemName.7day.holes` - The file is used to track gaps in the data for a seven day time period. Gaps are caused by stopping PFA or by an IPL.

This directory holds the following log files. Additional information is written to these log files when `DEBUG(1)`.

- `stderr` - This file only contains data if messages are generated by the JVM.
- `stdout` - This file only contains data if messages are generated by the JVM.
- `systemName.1hr.cart.log` - The log file generated by modeling code with details about code execution while one hour of historical data was being modeled.
- `systemName.24hr.cart.log` - The log file generated by modeling code with details about code execution while 24 hours of historical data was being modeled.
- `systemName.7day.cart.log` - The log file generated by modeling code with details about code execution while seven days of historical data was being modeled.
- `systemName.1hr.tree` - This file is generated by the modeling code. It contains information about the model tree which was built based on the last one hour of collected data.
- `systemName.24hr.tree` - This file is generated by the modeling code. It contains information about the model tree which was built based on the last 24 hours of collected data.
- `systemName.7day.tree` - This file is generated by the modeling code. It contains information about the model tree which was built based on the last seven days of collected data.
- `systemName.builder.log` - The log file generated by intermediate code that builds the files that are input to modeling with details about code execution.
- `systemName.launcher.log` - The log file generated by launcher code.
- `systemNameCONFIG.LOG` - The log file containing the configuration history for the last 30 days for this check.
- `systemNameCOLLECT.LOG` - The log file used during data collection.
- `systemNameMODEL.LOG` - The log file used during portions of the modeling phase.
- `systemNameRUN.LOG` - The log file used when the check runs.

pfa_directory/PFA_MESSAGE_ARRIVAL_RATE/EXC_timestamp

This directory contains all the relevant data for investigating exceptions issued by this check at the timestamp provided in the directory name. PFA keeps directories only for the last 30 exceptions and only if the exception is newer than the number of days specified in the `EXCDIRDAYS` parameter. Therefore at each exception, if directories exist that are too old, they will be deleted. In addition, if more than 30 exception directories exist, the oldest directory is deleted so that only 30 exceptions remain after the last exception is issued.

- `systemNameREPORT.LOG` - The log file containing the same contents as the IBM Health Checker for z/OS report for this exception as well as other diagnostic information issued during report generation.

pfa_directory/PFA_MESSAGE_ARRIVAL_RATE/config

This directory contains the configuration files for the check.

- `EXCLUDED_JOBS` - The file containing the list of excluded jobs for this check.
- `INCLUDED_JOBS` - the file containing the list of included jobs for this check.

PFA_PRIVATE_STORAGE_EXHAUSTION

Description:

The check is looking to see if there is a potential for private virtual storage below the 2GB bar to be exhausted by an address space in the future. PFA analyzes the following storage locations for individual address spaces:

- USER: The user region in the private area.
- EUSER: The user region in the extended private area.
- AUTH: The LSQA/SWA/229/230 area in the private area.
- EAUTH: The LSQA/SWA/229/230 area in the extended private area.
- ABOVE: The extended user private area (above 16M). This is also EUSER+EAUTH.
- BELOW: The user private area (below 16M). This is also USER+AUTH.

The PFA_PRIVATE_STORAGE_EXHAUSTION check detects three classes of private storage exhaustion:

- Spike
- Leak
- Creep

The check is not able to detect exhaustion caused by fast increases that are on a machine-time scale or even are faster than one collection interval. Also, the check cannot detect all issues caused by fragmentation.

If PFA detects that there is a potential for the exhaustion of private virtual storage for one or more address spaces, PFA issues exception AIRH222E and provides a list of those address spaces in the report. During the analysis, this check writes the private storage usage data at intervals to a z/OS UNIX System Services file in comma-separated value (.csv) format.

PFA chooses the address spaces to track based on the following:

- Collection waits until one hour after IPL so that data from the IPL activity does not skew the results.
- Address spaces that match names in the INCLUDED_JOBS file for this check will have all six storage locations collected once they have been active for one full collection interval.
- All other address spaces are eligible for collection once they have reached the configurable usage for an individual storage location, they have reached the configurable active duration for the address space, and they have been active for a full collection interval.
 - The configurable usage parameter is COLL% which is the minimum percentage of the current usage to the current capacity required in order for a storage location to be collected. Note that an individual storage location must reach COLL% before that storage location is collected. Not all storage locations need to reach COLL% before an individual storage location is collected.
 - The configurable active duration parameter is COLLUPTIME which is the minimum number of minutes the address space must have been running in order for any storage location for this address space to be collected.
- Address spaces that start within one hour after IPL ignore the COLLUPTIME configuration parameter so that data can be collected sooner after an IPL.
- Address spaces that restart are considered new address spaces and must reach COLLUPTIME before being collected even if they retain the name and ASID of a previously active address space.
- Address spaces with duplicate names are considered different address spaces and are tracked individually.
- Modeling occurs only for storage locations within address spaces whose percentage of current usage to the current capacity has reached the configurable MOD% value.

PFA_PRIVATE_STORAGE_EXHAUSTION

- Modeling occurs for up to 15 address spaces for each of the six individual storage locations. The address spaces chosen are those whose percentage of current usage to current capacity has increased the most in the last hour.
- If potential future exhaustion is detected based on the modeled trend of the current usage into the future for a storage location within an address space, PFA issues a health check exception message.
- When an exception occurs, the check reports only the address spaces whose usage has the potential for future exhaustion.
- When no problem exists, the check reports the storage locations in address spaces that were modeled.
- This check supports supervised learning. For details, see [“Configure supervised learning” on page 91](#)

When PFA_PRIVATE_STORAGE_EXHAUSTION issues the exception AIRH222E message, it will take dumps of the address spaces that have potential for the exhaustion of private virtual storage. The dumps are only taken when PFA_PRIVATE_STORAGE_EXHAUSTION check has DEBUG ON (DEBUG=1). PFA also includes the PFA address space in dumps. PFA dumps when first AIRH222E is issued and will dump again at the next comparison if the problem still exists.

Reason for check:

If an address space runs out of private virtual storage, jobs and started tasks experience abends.

Best practice:

The best practice is to predict private virtual storage problems before they occur, determine the cause of the problem, and take the appropriate action.

When IBM Health Checker for z/OS issues exception message AIRH222E, PFA has predicted that the amount of private virtual storage allocated within an address space is in jeopardy of being exhausted. Use the following steps to determine the appropriate action:

1. Look at the exception report. An asterisk next to an area or region indicates that future exhaustion was detected for that area or region.
2. Use [Figure 39 on page 149](#) as a worksheet to calculate the values in the report for the entire area in which the exception occurred.

For example, if the exception occurred for the Extended User Region (EUSER), transfer the data in the exception report of the entire Extended User Private Area (ABOVE 16M) to [Figure 39 on page 149](#). This allows you to see where the cause might be.

Note that the cause of the future exhaustion may not be in the area or region for which the exception occurred. The area or region for which the exception occurred may be a victim of the overall behavior of the entire private area to which it belongs.

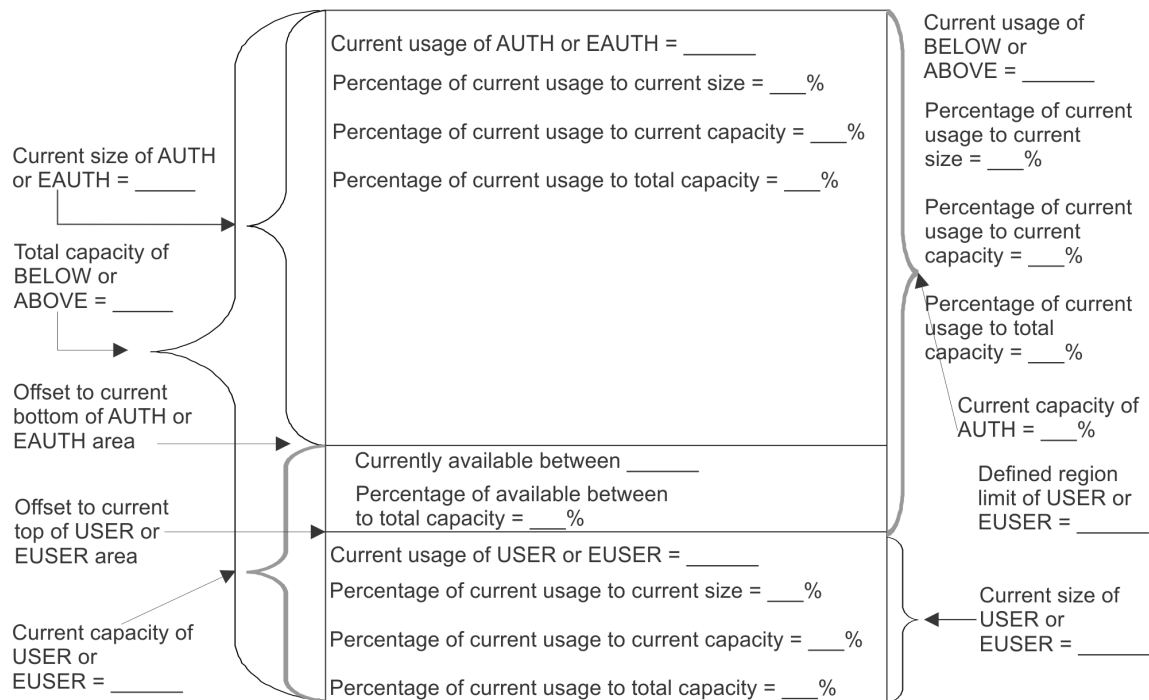


Figure 39. Calculating the cause of the private virtual storage exhaustion

3. In storage exhaustion scenarios, the **percentage of available storage compared to the total capacity** might be small. This is because one or both of the areas is taking a lot of storage or one (or both) of their dynamic boundaries is at an abnormal offset. To help determine the cause, look at the values you entered in [Figure 39 on page 149](#) and consider the following:

- a. First, look at the area or areas for which the exception was issued:

If the **percentage of current usage to current size** is relatively high, the space is being used well and does not have a lot of fragmentation or free space within it. If the **percentage of current usage to current capacity** is also high, future requests for storage may fail based on the current trends. If the **percentage of available usage compared to the total capacity** is relatively small, the problem may be in the corresponding area or region to this area or region.

If the **percentage of current usage to current size** is relatively low, there are gaps of free space available, but it may be fragmented and the current trend is detecting that the current usage will exhaust storage in the future.

- b. If the percentage of **available between to total capacity** is relatively low and the area for which the exception was issued is utilizing storage well, look at the other area or region that is in this private area. For example, if the exception was in the USER region, look at the AUTH area.

If the **percentage of current usage to current size** is relatively low, the offset to the top (or end) for this area or region may be located in a position within the region or area such that it is causing the **currently available between** to be very small.

For example, if the exception occurred in the USER region, but the **percentage of current usage to current size** for the AUTH area is not relatively close to 100%, some storage is used at the end of the AUTH area when it could likely be higher in that area now which would allow the **current available between** to be larger and allow the USER region to expand into that space.

For another example, if the exception occurred in the AUTH area, but the **percentage of current usage to current size** of the USER area is not relatively close to 100%, some storage is being used at the top of the USER region when it could likely be lower in that area now which would allow the **current available between** to be larger and allow the AUTH area to expand into that space.

- c. If the **percentage of current usage to current size** is relatively high for both AUTH|EAUTH and USER|EUSER and the **percentage of available between to total capacity** is relatively small,

storage exhaustion is likely to occur because overall private virtual storage is being allocated at a rate such that the trend is indication future exhaustion and the problem is not in just one area or region.

The bottom line is:

- If the **percentage of current used to current capacity** is relatively high, it indicates that you are approaching storage exhaustion for that area.
- If the **percentage of current usage to current size** is relatively low, it indicates that the storage area could be fragmented because it is holding much more space than it is actually using.

z/OS releases the check applies to:

z/OS V2R2 and later.

Type of check:

Remote

Parameters accepted:

Yes, as follows:

Table 16. PFA_PRIVATE_STORAGE_EXHAUSTION check parameters

Parameter name	Default value	Minimum Value	Maximum Value	Description
collectint	5 Minutes	1	360	This parameter determines the time (in minutes) to run the data collector that determines the amount of private virtual storage being used. The default is 5 minutes (5).
modelint	720 Minutes	4	1440	This parameter determines how often (in minutes) you want the system to analyze the data and construct a new private virtual storage usage model or prediction. Note that even when you set a value larger than 360, PFA performs the first model at 360 minutes (6 hours). By default, PFA analyzes the data and constructs a new model every 720 minutes (12 hours). The model interval must be at least four times larger than the collection interval. If necessary, modeling occurs more frequently.
collectinactive	1 (on)	0 (off)	1 (on)	Defines whether data will be collected and modeled even if the check is not eligible to run, (not ACTIVE(ENABLED), in IBM Health Checker for z/OS. The default is 1 (on) which means data will be collected even if comparisons will not be performed.
debug	0 (off)	0 (off)	1 (on)	This parameter (an integer of 0 or 1) is used at the direction of IBM service to generate additional diagnostic information for the IBM Support Center. When debug is set to 1 for this check, PFA dumps the address spaces causing the exception when an exception is issued. This debug parameter is used in place of the IBM Health Checker for z/OS policy. The default is off (0).
excdirdays	90 days	30	365	This parameter specifies the maximum number of days an exception directory will be kept for a check. The maximum number of exception directories kept is 30 regardless of age. Additional directories are deleted once they are this number of days old. The default is 90 days.

Table 16. PFA_PRIVATE_STORAGE_EXHAUSTION check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
forcemodel	no	00:00	23:59	<p>This parameter specifies whether a model is forced at a specific time even if the next model time is not reached. The NO value indicates that a model is not forced. Other values must be specified in a 24-hour time format and indicate the time of day to force the model. Modeling in a stable environment then proceeds the MODELINT minutes after this model. If the data becomes unstable for this metric, modeling occurs more often as needed after this time.</p> <p>For example, specifying FORCEMODEL (05:00) MODELINT (780) causes modeling to occur at 05:00 followed by the next model 13 hours later, at 18:00, in a stable environment. When using the values in an unstable environment, modeling occurs at 05:00 and at any point when modeling is needed because of instability. The FORCEMODEL time is not in effect until PFA models one time. PFA must establish a model, without delaying the first model, until the FORCEMODEL time, and PFA must have sufficient data for the first model. Also, if modeling occurred recently, such that there are no new collections at the FORCEMODEL time, the next model time is set to MODELINT minutes from the FORCEMODEL time.</p>
colluptime	180 Minutes	30	1440	<p>This parameter indicates the minimum number of minutes the address space must have been running (elapsed from address space start time) before collections will start. The default is 180 minutes.</p> <p>This parameter is ignored for jobs in the INCLUDED_JOBS file. This parameter is also ignored for jobs that start within the first hour after IPL.</p>
coll%	20 %	10	50	<p>This parameter indicates the minimum percentage of the current usage to the current capacity in order for the location for the address space to be collected. Setting this value lower will cause data collection for more address spaces. Setting this value higher will cause data collection for fewer address spaces.</p> <p>This parameter is ignored for jobs in the INCLUDED_JOBS file.</p>

Table 16. PFA_PRIVATE_STORAGE_EXHAUSTION check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
mod%	40 %	10	70	<p>This parameter indicates the minimum percentage of the current used space to the current capacity that must be used by any area before modeling will occur for address spaces where an area is being collected. For example, a value of 40 indicates that the currently used space must be at least 40 percent of the current capacity before modeling will be performed for that area. The default is 40.</p> <p>MOD% must be greater than or equal to COLL% because data must be collected before models can occur.</p>
comp%	100 %	75	125	<p>This parameter indicates the percentage of the current capacity used for the value that defines exhaustion. The capacity is adjusted to this percentage of capacity in calculations. Also, the prediction must be greater than or equal to this adjusted capacity in order for comparisons to be performed for modeled address spaces. For example, a value of 100 indicates that 100 percent of the capacity is used in the calculations using capacity and that the prediction must be at least 100 percent of the current capacity before comparisons will be performed. The default is 100.</p> <p>Set this value to be less than 100 if the system has abends caused by virtual storage exhaustion so that this check is more sensitive and comparisons are made when the predictions aren't indicating future usage that will be 100% of capacity. Set this value to be greater than 100 if false exceptions occur for multiple address spaces so that the check is less sensitive and comparisons aren't made until the prediction is larger than 100% of capacity.</p> <p>COMP% must be greater than or equal to MOD% because predictions must exist before comparisons can be performed.</p>

Table 16. PFA_PRIVATE_STORAGE_EXHAUSTION check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
e_high	180	0 or UNUSED	1440 or MAX	<p>This parameter indicates when a HIGH severity exception which results in a critical eventual action WTO is issued. The value specified is the maximum number of minutes before exhaustion is detected in order for a HIGH severity exception to be issued. The default is 180 which means that if exhaustion is predicted to occur between now and 180 minutes from now, a HIGH severity exception will occur.</p> <p>A value of 0 or UNUSED indicates that this dynamic severity is not used.</p> <p>A value of MAX or 1440 indicates that this dynamic severity is used to the end of the time the current model is active. It is recommended to specify MAX or 1440 for the lowest level of severity to be used. If none of the dynamic severity parameters specify MAX or 1440, the severity used for exception messages for the time range left after the last dynamic severity specified will be SYSTEM which may resolve to a higher severity than the last value specified.</p> <p>The values for E_HIGH, E_MED, E_LOW, and E_NONE cannot overlap.</p>
e_med	300	0 or UNUSED	1440 or MAX	<p>This parameter indicates when a MED severity exception which results in an eventual action WTO is issued. The value specified is the maximum number of minutes before exhaustion is detected in order for a MED severity exception to be issued. The default is 300 which means that if exhaustion is predicted to occur between the start of this dynamic severity time range and 300 minutes from now, a MED severity exception will occur.</p> <p>The start of this dynamic severity time range occurs when the time range for the dynamic severity for E_HIGH has expired if it is non-zero.</p> <p>A value of 0 or UNUSED indicates that this dynamic severity is not used.</p> <p>A value of MAX or 1440 indicates that this dynamic severity is used to the end of the time the current model is active. It is recommended to specify MAX or 1440 for the lowest level of severity to be used. If none of the dynamic severity parameters specify MAX or 1440, the severity used for exception messages for the time range left after the last dynamic severity specified will be SYSTEM which may resolve to a higher severity than the last value specified.</p> <p>The values for E_HIGH, E_MED, E_LOW, and E_NONE cannot overlap.</p>

Table 16. PFA_PRIVATE_STORAGE_EXHAUSTION check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
e_low	MAX	0 or UNUSED	1440 or MAX	<p>This parameter indicates when a LOW severity exception which results in an informational WTO is issued. The value specified is the maximum number of minutes before exhaustion is detected in order for a LOW severity exception to be issued. The default is 1440 which means that if exhaustion is predicted to occur between the start of this dynamic severity time range and the expiration of the prediction, a LOW severity exception will occur.</p> <p>The start of this dynamic severity time range occurs when the time range for the dynamic severity for E_HIGH and E_MED have expired if either or both are non-zero.</p> <p>A value of 0 or UNUSED indicates that this dynamic severity is not used.</p> <p>A value of MAX or 1440 indicates that this dynamic severity is used to the end of the time the current model is active. It is recommended to specify MAX or 1440 for the lowest level of severity to be used. If none of the dynamic severity parameters specify MAX or 1440, the severity used for exception messages for the time range left after the last dynamic severity specified will be SYSTEM which may resolve to a higher severity than the last value specified.</p> <p>The values for E_HIGH, E_MED, E_LOW, and E_NONE cannot overlap.</p>

Table 16. PFA_PRIVATE_STORAGE_EXHAUSTION check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
e_none	UNUSED	0 or UNUSED	1440 or MAX	<p>This parameter indicates when an exception with no severity is issued. The value specified is the maximum number of minutes before exhaustion is detected in order for an exception with no severity to be issued. The default is UNUSED which means that this severity level is not used by default.</p> <p>When a non-zero value is specified, the start of this dynamic severity time range occurs when the time range for the dynamic severity for E_HIGH, E_MED, and E_LOW have expired when they are non-zero.</p> <p>A value of 0 or UNUSED indicates that this dynamic severity is not used.</p> <p>A value of MAX or 1440 indicates that this dynamic severity is used to the end of the time the current model is active. It is recommended to specify MAX or 1440 for the lowest level of severity to be used. If none of the dynamic severity parameters specify MAX or 1440, the severity used for exception messages for the time range left after the last dynamic severity specified will be SYSTEM which may resolve to a higher severity than the last value specified.</p> <p>The values for E_HIGH, E_MED, E_LOW, and E_NONE cannot overlap.</p>

To determine the status of the private storage exhaustion check, issue `f pfa,display,check(pfa_private_storage_exhaustion),detail`. For the command example and more details, see [“PFA DISPLAY examples” on page 87](#). The following example shows the output written to message AIR018I in SDSF:

```
AIR018I 16:20:21 PFA CHECK DETAIL
CHECK NAME: PFA_PRIVATE_STORAGE_EXHAUSTION
ACTIVE : YES
TOTAL COLLECTION COUNT : 5
SUCCESSFUL COLLECTION COUNT : 5
LAST COLLECTION TIME : 09/01/2015 10:18:22
LAST SUCCESSFUL COLLECTION TIME : 09/01/2015 10:18:22
NEXT COLLECTION TIME : 09/01/2015 10:23:22
TOTAL MODEL COUNT : 1
SUCCESSFUL MODEL COUNT : 1
LAST MODEL TIME : 09/01/2015 10:18:24
LAST SUCCESSFUL MODEL TIME : 09/01/2015 10:18:24
NEXT MODEL TIME : 09/01/2015 22:18:24
CHECK SPECIFIC PARAMETERS:
COLLECTINT : 5
MODELINT : 720
COLLECTINACTIVE : 1=ON
DEBUG : 0=OFF
EXCDIRSDAYS : 90
FORCEMODEL : NO
COLLUPTIME : 180
COLL% : 20
MOD% : 40
COMP% : 100
E_HIGH : 180
E_MED : 300
E_LOW : MAX
E_NONE : UNUSED
```

Dynamic Severity Settings: Note that the exception message issued will be sent with the severity based on the dynamic severity parameters, E_HIGH, E_MED, E_LOW, and E_NONE if any are non-zero or any are a value other than UNUSED. The severity is based on the amount of time before the exhaustion is predicted to occur.

The parameters indicate the maximum number of minutes from the time the exception is issued to the predicted exhaustion. For example, the default values state that a HIGH severity message is issued if the exhaustion will occur within 180 minutes, a MED severity message is issued if the exhaustion will occur from over 180 minutes to 5 hours, a LOW severity message is issued if the exhaustion will occur from over 5 hours to the end of the model time. The severity of NONE is not used by default. If you do not specify MAX (or 1440) for one of the severity values, the system will choose the severity based on the severity in the WTOTYPE and SEVERITY parameters of the check. This system severity value may be larger than the lowest dynamic severity specified; therefore, it is recommended to specify MAX for one of the dynamic severity parameters when dynamic severity is being used.

For example, if HIGH severity is not desired and MED severity is desired for exhaustion predicted to occur from now until 6 hours from now and LOW severity is desired for anything farther than 6 hours from now to the end of the model time, use the following command or create a policy in HZSPRMxx to make the change persistent:

```
f hzsproc,update,check(ibmpfa,pfa_p*),parm('E_HIGH(UNUSED) E_MED(720)
E_LOW(MAX) E_NONE(UNUSED)')
```

For another example, if HIGH severity and MED severity defaults are desired, but you want no severity exceptions for everything beyond 5 hours rather than LOW severity exceptions, use the following command or create a policy in HZSPRMxx to make the change persistent:

```
f hzsproc,update,check(ibmpfa,pfa_p*),parm('E_HIGH(180) E_MED(300)
E_LOW(UNUSED) E_NONE(MAX)')
```

User override of IBM values:

The following shows keywords you can use to override check values on either a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. See Chapter 8, “Managing PFA checks,” on page 83. You can copy and modify this statement to override the check defaults:

```
ADDREPLACE POLICY[(polycyname)] [STATEMENT(name)]
UPDATE CHECK(IBM PFA,PFA_PRIVATE_STORAGE_EXHAUSTION)
ACTIVE
SEVERITY(MEDIUM)
INTERVAL(00:01)
PARMS=('COLLECTINT(5)', 'MODELINT(720)',
        'COLLECTINACTIVE(1)', 'DEBUG(0)', 'EXCDIRSDAYS(90)',
        'FORCEMODEL(NO)', 'COLLUPTIME(180)', 'COLL%(20)',
        'MOD%(40)', 'COMP%(100)', 'E_HIGH(180)', 'E_MED(300)',
        'E_LOW(MAX)', 'E_NONE(UNUSED)')
DATE(20140313)
REASON('The reason for the change.')
```

Verbose support:

The check does not provide additional details in verbose mode. The check uses the PFA DEBUG parameter (not the DEBUG parameter in IBM Health Checker for z/OS) for writing additional information to its log files for both exception and non-exception states.

Debug support:

The DEBUG parameter in IBM Health Checker for z/OS is ignored by this check. Rather, the debug parameter is a PFA check specific parameter. For details, see “Understanding how to modify PFA checks” on page 83.

Reference:

For more information about PFA, see the topic on “Overview of Predictive Failure Analysis” on page 73.

Messages:

The output is a message arrival rate prediction report that corresponds to the message issued. One of the following reports is generated:

- AIRH222E

For additional message information, see the topics on:

- [AIRH messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).
- [AIR messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).

SECLABEL recommended for MLS users:

SYSLOW

Output:

The output is a private storage exhaustion output report.

No problem report: When no exception is issued for the PFA_PRIVATE_STORAGE_EXHAUSTION check, PFA issues a report displaying the modeled storage locations.

```

Private Storage Exhaustion Prediction Report

Last successful model time      : 04/05/2012 07:08:01
Next model time                : 04/05/2012 19:08:04
Model interval                 : 720
Last successful collection time : 04/05/2012 08:22:15
Next collection time           : 04/05/2012 08:37:16
Collection interval            : 15

Total capacity of User Private Area (Below 16)      : 8168
Total capacity of Extended User Private Area (Above 16M): 1211392

Name      ASID Area      Current
Usage in  Current    Current    Percentage
Kilobytes Prediction in Kilobytes Capacity
Used to
Capacity

-----
JOB1      0032 EAUTH      475980      480000      830936      57%
JOB12     0035 BELOW      4668        5652        8168        57%
JOBZZZ    053D ABOVE     380456     390000     1211382     31%
JOB1XY    0532 USER       2292        2140        7792        29%
...
(maximum of 15)

```

Figure 40. Private storage exhaustion prediction no problem report

Exception report: When an exception is issued, only address spaces causing the exception are listed. The summary information is listed for the storage locations causing the exception followed by details for each address space with an exception. Any areas causing the exception will have an asterisk in front of them.

In the example that follows, the AUTH area caused the exception in address space PSETSTAU.

- The **percentage of current usage to current size** for the AUTH area is 100% which means that it is efficiently using all of the storage within the area down to the **offset to the current bottom of the AUTH area**.
- The **percentage of current usage to current capacity** is 69% which indicates that there is some storage available between in which the AUTH area can expand. The **percentage of available between to total capacity** is at 31% which confirms there is space available for the expansion.
- The **percentage of current usage to current size** for the USER region is also at 100% which means that it is efficiently using all of the storage within that region up to the **offset to the current top of the USER region**.
- Therefore, there isn't any fragmentation in either the AUTH area or the USER region that is causing the potential for storage exhaustion and the top of the USER region and the end of the AUTH area are in appropriate locations.
- The problem indicated by this exception is simply that the trend that was modeled for the AUTH area indicated that it has the potential to allocate more storage than the capacity in the future which will lead to abends.

Figure 41 on page 158 shows a simplified view of the extended user private area at high capacity but low usage, using the figures shown in the exception report shown in Figure 42 on page 159 and Figure 43 on page 159:

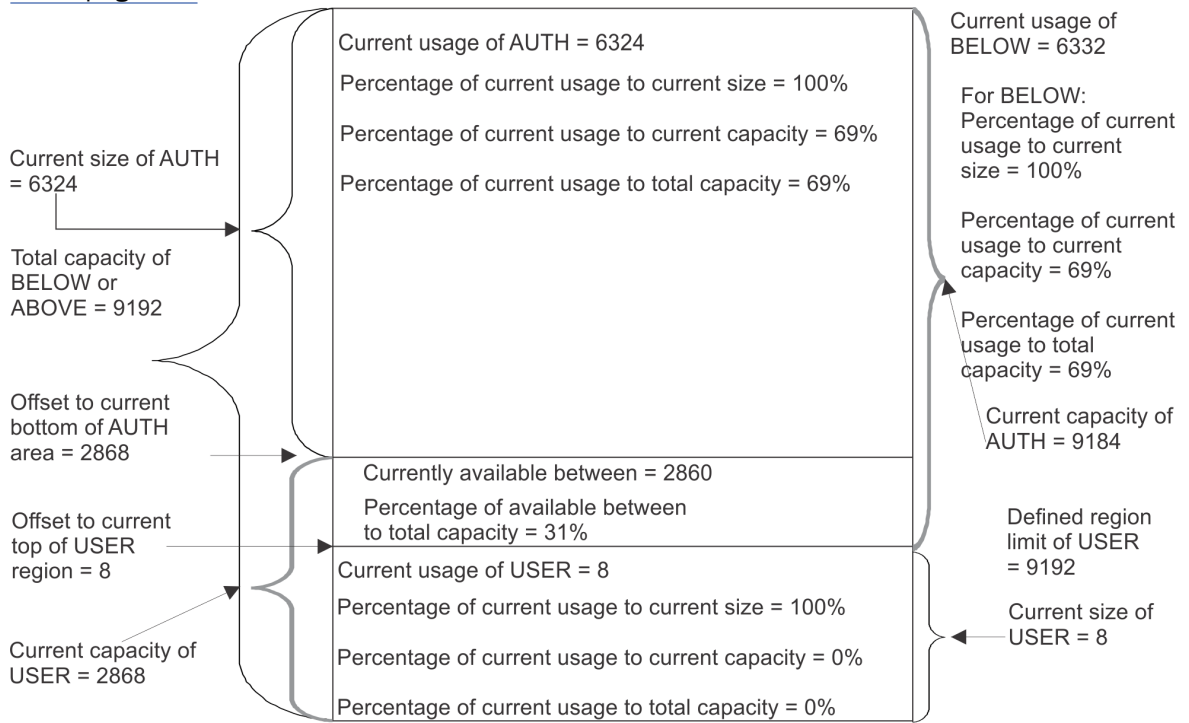


Figure 41. An example of extended user private area at high capacity but low usage

Figure 42 on page 159 shows details for one address space although four address spaces are listed in Figure 43 on page 159. In an actual exception report, the details are provided for each address space having an exception.


```

Job name: PSETSTAU
ASID      : 0029
Extended User Private Area (Above 16M line):
  Total capacity (K)          : 1474248
  Current size (K)           : 2004
  Current usage (K)          : 1912
  Percentage of current usage to current size : 95%
  Percentage of current usage to total capacity : 0%
  Percentage of current size to total capacity : 0%
  Predicted usage (K)        : Not modeled
  Currently available between EUSER and EAUTH (K) : 1472244
  Percentage of available between to total capacity: 100%
Extended LSQA/SWA/229/230 Area (EAUTH):
  Current capacity (K)       : 1474240
  Current size (K)           : 1996
  Current usage (K)          : 1904
  Percentage of current usage to current size : 95%
  Percentage of current usage to current capacity: 0%
  Percentage of current usage to total capacity : 0%
  Percentage of current size to total capacity : 0%
  Predicted usage (K)        : Not modeled
  Offset to current bottom of EAUTH area (K) : 1472252
Extended User Region (EUSER):
  Current capacity (K)       : 1472252
  Current size (K)           : 8
  Current usage (K)          : 8
  Percentage of current usage to current size : 100%
  Percentage of current usage to current capacity: 0%
  Percentage of current usage to total capacity : 0%
  Percentage of current size to total capacity : 0%
  Predicted usage (K)        : Not modeled
  Offset to current top of EUSER region (K) : 8
  Defined region limit (K)   : 1481728

```

Figure 42. Private storage exhaustion prediction no problem report (1 of 2)

```

*User Private Area (Below 16M line):
  Total capacity (K)          : 9192
  Current size (K)           : 6332
  Current usage (K)          : 6332
  Percentage of current usage to current size : 100%
  Percentage of current usage to total capacity : 69%
  Percentage of current size to total capacity : 69%
  Predicted usage (K)        : 12291
  Currently available between USER and AUTH (K) : 2860
  Percentage of available between to total capacity: 31%
*LSQA/SWA/229/230 Area (AUTH):
  Current capacity (K)       : 9184
  Current size (K)           : 6324
  Current usage (K)          : 6324
  Percentage of current usage to current size : 100%
  Percentage of current usage to current capacity: 69%
  Percentage of current usage to total capacity : 69%
  Percentage of current size to total capacity : 69%
  Predicted usage (K)        : 12283
  Offset to current bottom of AUTH area (K) : 2868
User Region (USER):
  Current capacity (K)       : 2868
  Current size (K)           : 8
  Current usage (K)          : 8
  Percentage of current usage to current size : 100%
  Percentage of current usage to current capacity: 0%
  Percentage of current usage to total capacity : 0%
  Percentage of current size to total capacity : 0%
  Predicted usage (K)        : Not modeled
  Offset to current top of USER region (K) : 8
  Defined region limit (K)   : 9192

```

Figure 43. Private storage exhaustion prediction no problem report (2 of 2)

The following fields apply to all reports:

- **USER:** An abbreviation for the user region in the private area.

- **EUSER:** An abbreviation for the user region in the extended private area.
- **AUTH:** An abbreviation for the LSQA/SWA/229/230 area in the private area.
- **EAUTH:** An abbreviation for the LSQA/SWA/229/230 area in the extended private area.
- **ABOVE:** The extended user private area (above 16M). This is also EUSER+EAUTH.
- **BELOW:** The user private area (below 16M). This is also USER+AUTH.
- **Total capacity:** The number of kilobytes that can be used by the user region and the LSQA/SWA/229/230 area combined (or the extended user region and the extended LSQA/SQA/229/230 area combined).
- **Current usage:** The number of kilobytes currently used by the region or area. This number may contain some free space within the area.
- **Predicted usage:** The prediction of this region or area at the last model, if modeled. The prediction is created using analysis of the trend created over time by the current usage by this region or area projected into the future. The percentage of the predicted usage to the current capacity must be at least COMP% in order for the comparisons to be performed.
- **Currently available between USER and AUTH (or EUSER and EAUTH):** The number of kilobytes unused between the top of the USER | EUSER region and the end of the AUTH | EAUTH area. This entire area is available for use by the AUTH | EAUTH. This entire area is available for use by the USER | EUSER unless the amount from the start of the region to the defined region limit offset is less than the start of the region to the offset to the current end of the AUTH | EAUTH area. (See also the definition for current capacity.)
- **Current size:** The current number of kilobytes located within the region or area. This number includes space used in that region or area as well as free space within that region or area.
- **Current capacity:** The current number of kilobytes that can be used by the USER or AUTH areas (or EUSER or EAUTH areas). This value is used as the capacity for the USER or AUTH areas (or EUSER or EAUTH areas) in comparisons to detect exhaustion.
 - For USER | EUSER, this number is the minimum of the amount from the start of the region to the offset to the current end of the AUTH | EAUTH area and the amount from the start of the region to the defined region limit offset.
 - For AUTH | EAUTH, this number is the sum of the current size of the area plus the amount currently available between USER and AUTH (or EUSER and EAUTH).

Note that the capacity value is dynamic because the amount currently available between the USER and AUTH (or EUSER and EAUTH) changes.

- **Defined region limit:** The region limit as defined to VSM for the address space. If this value is less than the offset to current end of AUTH | EAUTH area, it is used in calculating the capacity for USER | EUSER.
- **Offset to current top of USER | EUSER region:** The offset from the start of the virtual storage address space to the top of the user region (or extended user region).
- **Offset to current end of AUTH | EAUTH area:** The offset from the start of the virtual storage address space to the end of the LSQA/SWA/229/230 AUTH | EAUTH area.
- **Percentage used to capacity:** The ratio of the current usage and the current capacity multiplied by 100.
- **Percentage of current usage to total capacity:** The ratio of the current usage and the total capacity multiplied by 100.
- **Percentage of current usage to current size:** The ratio of the current usage and the current size multiplied by 100.
- **Percentage of current size to total capacity:** The ratio of the current size and the total capacity multiplied by 100.
- **Percentage of available between to total capacity:** The ratio of the currently available between USER and AUTH (or EUSER and EAUTH) and the total capacity multiplied by 100.

- **Not modeled:** This value is printed for the prediction if the address space was not modeled for the specified area or region.

Directories

Note: The content and names for these files and directories are subject to change and cannot be used as programming interfaces; these files are documented only to provide help in diagnosing problems with PFA.

pfa_directory

This directory contains all the PFA check subdirectories and is pointed to by the PFADIR= specified in the /etc/PFA/ini file, or, if you do not specify PFADIR=, the home directory of the PFA user.

pfa_directory/PFA_PRIVATE_STORAGE_EXHAUSTION/data

The directory for PFA private storage exhaustion that holds data and modeling results.

Results files

- `systemName.prediction` - This file is generated by the modeling code. It lists the address spaces and storage locations that were modeled along with their predictions, the private storage usage, and additional information required for PFA processing.
- `systemName.prediction.html` - This file contains an .html report version of the data found in the `systemName.prediction` file.

Data store files:

- `systemName.5day.data` - The data collection file for 5 days of data.
- `systemName.sum` - The summarized data for days 6-30. After 5 days, the data in the `5day.data` file is averaged and compressed. The data summarized is deleted from the `5day.data` file.

Intermediate files:

- `systemName.data` - The file that is used as input to modeling which contains the address spaces and their storage locations that have recently increased the most.
- `systemName.sumX` - This file is used during the summarization of the `5day.data` file.

This directory holds the following log files. Additional information is written to these log files when `DEBUG(1)`.

- `stderr` - This file only contains data if messages are generated by the JVM.
- `stdout` - This file only contains data if messages are generated by the JVM.
- `systemName.builder.log` - The log file generated by intermediate code that builds the files that are input to modeling with details about code execution.
- `systemName.launcher.log` - The log file generated by launcher code.
- `systemNameCONFIG.LOG` - The log file containing the configuration history for the last 30 days for this check.
- `systemNameCOLLECT.LOG` - The log file used during data collection.
- `systemNameMODEL.LOG` - The log file used during portions of the modeling phase.
- `systemNameRUN.LOG` - The log file used when the check runs.
- `systemName.cart.log` - The log file generated by modeling code that contains the execution details of modeling code.
- `systemName.tree` - This file is generated by the modeling code. It contains information about the model tree that is built based on collected data.

pfa_directory/PFA_PRIVATE_STORAGE_EXHAUSTION/EXC_timestamp

This directory contains all the relevant data for investigating exceptions issued by this check at the timestamp provided in the directory name. PFA keeps directories only for the last 30 exceptions and only if the exception is newer than the number of days specified in the EXCDIRDAYS parameter. Therefore at each exception, if directories exist that are too old, they will be deleted. In

PFA_PRIVATE_STORAGE_EXHAUSTION

addition, if more than 30 exception directories exist, the oldest directory is deleted so that only 30 exceptions remain after the last exception is issued.

- `systemNameREPORT.LOG` - The log file containing the same contents as the IBM Health Checker for z/OS report for this exception as well as other diagnostic information issued during report generation.

pfa_directory/PFA_PRIVATE_STORAGE_EXHAUSTION/config

This directory contains the configuration files for the check.

- `EXCLUDED_JOBS` - The file containing the list of excluded jobs for this check.
- `INCLUDED_JOBS` - The file containing the list of included jobs for this check.

PFA_SMF_ARRIVAL_RATE

Description:

When SMF is active on the system, this check determines if there is an abnormal SMF arrival rate per amount of CPU used. PFA examines only the SMF record types you set the SMFPRMxx parmlib member to generate. When the number of SMF records written per amount of CPU used is unusually high or low, PFA can provide an early indication of a problem and potentially prevent damage to an LPAR.

To avoid skewing the SMF arrival rate, PFA ignores the first hour of SMF data after IPL and the last hour of SMF data prior to shutdown. In addition, PFA attempts to track the same persistent jobs that it tracked prior to IPL or PFA restart if the same persistent jobs are still active. (The same persistent jobs must still be active for PFA to track the same jobs and there must have been ten jobs tracked previously.)

This check is not designed to detect performance problems caused by insufficient resources, faulty WLM policy, or spikes in work. However, it might help to determine if a performance problem detected by a performance monitor or WLM is caused by a damaged system.

Guideline: If you modify the SMF record types in the SMFPRMxx parmlib member, delete the files in the *PFA_SMF_ARRIVAL_RATE/data* directory to ensure PFA is collecting relevant information.

The SMF arrival rate check issues an exception using the following four types of comparisons. After the check issues an exception, it does not perform the next comparison type.

Note: If SMF is restarted, the data previously collected will be automatically discarded and the check will enter the phase where it collects data for a period of time to use in determining the jobs to track. This processing is done to reduce false positives and to ensure that data that was potentially collected using different SMF record types is not used by PFA after SMF is restarted.

1. **Top persistent jobs:** The SMF arrival rate check tracks the top persistent jobs individually. Jobs are considered persistent when they start within an hour after IPL. PFA determines which jobs to track individually based on the following:
 - If PFA previously ran on this system and the same 10 jobs that were previously tracked are active, PFA tracks the same jobs.
 - If PFA never ran on the system or the same jobs previously tracked are not all active, PFA collects data for a period of time to use in determining the jobs with the highest arrival rates. After this time, the jobs with the highest arrival rates for that period are individually tracked.

PFA performs this type of comparison to determine if the SMF arrival rate is higher than expected or lower than expected.

2. **Other persistent jobs:** The persistent jobs that are not individually track are considered "other persistent jobs". The SMF arrival rate check models predictions using the totals for this group. When PFA determines the SMF arrival rate is higher than expected, the comparisons are performed individually using a mathematical formula. When PFA determines the SMF rate is lower than expected, the comparisons are performed using the totals for the group.
3. **Non-persistent jobs:** The jobs that start over an hour after IPL are the non-persistent jobs. The predictions and comparisons are done using the totals for this group. PFA only performs this type of comparison to determine if the SMF arrival rate is higher than expected.
4. **Total system:** The predictions and the comparisons are done using the totals for the entire system. PFA performs this type of comparison to determine if the SMF arrival rate is higher than expected or lower than expected. All jobs are included in this group except those that match a job specified in the */config/EXCLUDED_JOBS* file for this check.

Reason for check:

The objective of this check is to determine if there is potential of damage to an LPAR by checking the arrival rate of SMF records per amount of CPU used on a system.

Best practice:

If an unexpectedly high number of SMF records was detected, the best practice is to review the SMF records being sent by the address spaces identified on the report and examine the system log to determine what is causing the increase in SMF activity.

If an unexpectedly low number of SMF records was detected, the best practice is to examine the report in SDSF for details about why the exception was issued. Use the Runtime Diagnostics output in the report to assist you in diagnosing and fixing the problem. For more information about Runtime Diagnostics, see [Chapter 5, “Runtime Diagnostics,” on page 35](#).

z/OS releases the check applies to:

z/OS V1R12 and later.

Type of check:

Remote

Parameters accepted:

Yes, as follows:

Table 17. PFA_SMF_ARRIVAL_RATE check parameters

Parameter name	Default value	Minimum Value	Maximum Value	Description
collectint	15 Minutes	15	360	This parameter determines how often (in minutes) to run the data collector that retrieves the current SMF arrival rate.
modelint	720 Minutes	60	1440	This parameter determines how often (in minutes) you want the system to analyze the data and construct a new SMF arrival rate model or prediction. By default, PFA analyzes the data and constructs a new model every “default value” minutes. The model interval must be at least four times larger than the collection interval. Note that, even when you set a value larger than 360, PFA performs the first model at 360 minutes (6 hours). By default, PFA analyzes the data and constructs a new model every 720 minutes (12 hours).
stddev	3	2	100	This parameter is used to specify how much variance is allowed between the actual SMF arrival rate per amount of CPU and the expected SMF arrival rate. It is used when determining if the actual SMF arrival rate has increased beyond the allowable upper limit. It also determines how much variance is allowed across the time range predictions. If you set the STDDEV parameter to a small value, an exception will be issued if the actual SMF arrival rate is closer to the expected SMF arrival rate and the predictions across the time ranges are consistent. If you set the STDDEV parameter to a larger value, an exception will be issued if the actual SMF arrival rate is significantly greater than the expected SMF arrival rate even if the predictions across the different time ranges are inconsistent.
collectinactive	1 (on)	0 (off)	1 (on)	Defines whether data will be collected and modeled even if the check is not eligible to run, not ACTIVE(ENABLED), in IBM Health Checker for z/OS.
trackedmin	3	0	1000	This parameter defines the minimum SMF arrival rate required for a persistent job in order for it to be considered a top persistent job that should be tracked individually.

Table 17. PFA_SMF_ARRIVAL_RATE check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
exceptionmin	1	0	1000	This parameter is used when determining if an exception should be issued for an unexpectedly high SMF arrival rate. For tracked jobs and other persistent jobs, this parameter defines the minimum SMF arrival rate and the minimum predicted SMF arrival rate required to cause a too high exception. For non-persistent jobs and the total system comparisons, this parameter defines the minimum SMF arrival rate required to cause a too high exception.
checklow	1 (on)	0 (off)	1	Defines whether Runtime Diagnostics is run to validate that the absence of SMF records is caused by a problem. If this value is off then exceptions will not be issued for conditions in which the SMF arrival rate is unexpectedly low.
stddevlow	2	2	100	This parameter is used to specify how much variance is allowed between the actual SMF arrival rate per amount of CPU and the expected SMF arrival rate when determining if the actual rate is unexpectedly low. <ul style="list-style-type: none"> • If you set the STDDEVLOW parameter to a smaller value, an exception is issued when the actual SMF arrival rate is closer to the expected SMF arrival rate. • If you set the STDDEVLOW parameter to a larger value, an exception is issued when the actual SMF arrival rate is significantly lower than the expected SMF arrival rate.
limitlow	3	1	100	This parameter defines the maximum SMF arrival rate allowed when issuing an exception for an unexpectedly low number of SMF records.
excdirdays	90	30	365	This parameter specifies the maximum number of days that an exception directory is kept for the check. The maximum number of exception directories kept is 30 regardless of age. Directories are deleted when they are this number of days old. The default is 90 days.

Table 17. PFA_SMF_ARRIVAL_RATE check parameters (continued)

Parameter name	Default value	Minimum Value	Maximum Value	Description
forcemodel	no	00	23:59	<p>This parameter specifies whether a model is forced at a specific time even if the next model time is not reached. The NO value indicates that a model is not forced. Other values must be specified in a 24-hour time format and indicate the time of day to force the model. Modeling in a stable environment then proceeds the MODELINT minutes after this model. If the data becomes unstable for this metric, modeling occurs more often as needed after this time.</p> <p>For example, specifying FORCEMODEL (05:00) MODELINT (780) causes modeling to occur at 05:00 followed by the next model 13 hours later, at 18:00, in a stable environment. When using the values in an unstable environment, modeling occurs at 05:00 and at any point when modeling is needed because of instability. The FORCEMODEL time is not in effect until PFA models one time. PFA must establish a model, without delaying the first model, until the FORCEMODEL time, and PFA must have sufficient data for the first model. Also, if modeling occurred recently, such that there are no new collections at the FORCEMODEL time, the next model time is set to MODELINT minutes from the FORCEMODEL time.</p>
debug	0 (off)	0 (off)	1 (on)	This parameter (an integer of 0 or 1) is used at the direction of IBM service to generate additional diagnostic information for the IBM Support Center. This debug parameter is used in place of the IBM Health Checker for z/OS policy. The default is off (0).

To determine the status of the SMF arrival rate check, issue `f pfa,display,check(pfa_smf_arrival_rate),detail`. For the command example and more details, see [“PFA DISPLAY examples” on page 87](#). The following example shows the output written to message AIR018I in SDSF:

```

AIR018I 02:22:54 PFA CHECK DETAIL
CHECK NAME:  PFA_SMF_ARRIVAL_RATE
ACTIVE                      : YES
TOTAL COLLECTION COUNT      : 5
SUCCESSFUL COLLECTION COUNT : 5
LAST COLLECTION TIME        : 02/05/2009 10:18:22
LAST SUCCESSFUL COLLECTION TIME : 02/05/2009 10:18:22
NEXT COLLECTION TIME        : 02/05/2009 10:33:22
TOTAL MODEL COUNT           : 1
SUCCESSFUL MODEL COUNT      : 1
LAST MODEL TIME              : 02/05/2009 10:18:24
LAST SUCCESSFUL MODEL TIME   : 02/05/2009 10:18:24
NEXT MODEL TIME              : 02/05/2009 22:18:24
CHECK SPECIFIC PARAMETERS:
COLLECTINT                   : 15
MODELINT                     : 720
COLLECTINACTIVE              : 1=YES
DEBUG                        : 0=NO
EXCDIRSDAYS                  : 90
FORCEMODEL                   : NO
DEBUG                        : 0=NO
STDDEV                       : 3
TRACKEDMIN                   : 3
EXCEPTIONMIN                 : 1
CHECKLOW                     : 1=YES
STDDEVLOW                    : 2
LIMITFLOW                    : 3

```


User override of IBM values:

The following shows keywords you can use to override check values on either a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. This statement can be copied and modified to override the check defaults:

```
ADDREPLACE POLICY[(policyname)] [STATEMENT(name)]
UPDATE CHECK(IBM PFA,PFA_SMF_ARRIVAL_RATE)
ACTIVE
SEVERITY(MEDIUM)
INTERVAL(ONETIME)
PARMS=('COLLECTINT(15)', 'MODELINT(720)', 'COLLECTINACTIVE(1)', 'DEBUG(0)',
       'EXCDIRSDAYS(90)', 'FORCEMODEL(NO)', 'STDDEV(3)', 'TRACKEDMIN(2)',
       'EXCEPTIONMIN(10)', 'CHECKLOW(1)', 'STDDEVLOW(4)', 'LIMITLOW(3)')
DATE(20140313)
REASON('The reason for the change.')
```

The SMF arrival rate check is designed to run automatically after every data collection. Do not change the INTERVAL parameter.

Verbose support:

The check does not provide additional details in verbose mode. The check uses the PFA DEBUG parameter (not the DEBUG parameter in IBM Health Checker for z/OS) for writing additional information to its log files for both exception and non-exception states.

Debug support:

The DEBUG parameter in IBM Health Checker for z/OS is ignored by this check. Rather, the debug parameter is a PFA check specific parameter. For details, see [“Understanding how to modify PFA checks”](#) on page 83.

Reference:

For more information about PFA, see the topic about [“Overview of Predictive Failure Analysis”](#) on page 73.

Messages:

The output is a SMF arrival rate prediction report that corresponds to the message issued. PFA generates one of the following reports:

- AIRH187E and AIRH208E - tracked jobs exception report
- AIRH188E and AIRH209E - other persistent jobs exception report
- AIRH191E - other non-persistent jobs exception report
- AIRH174E and AIRH175E - total system exception report

For additional message information, see the topics about:

- [AIRH messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).
- [AIR messages in z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#).

SECLABEL recommended for MLS users:

SYSLOW

Output:

The output is a variation of the SMF arrival rate prediction report. The values found in the SMF arrival prediction file are as follows:

Tracked jobs exception report: PFA issues the SMF arrival rate exception report for tracked jobs when any one or more tracked, persistent jobs cause an exception. These exceptions can be caused by a higher than expected SMF arrival rate or a lower than expected SMF arrival rate. Only the tracked jobs that caused the exception are included in the list of jobs on the report. If the report was generated due to a lower than expected SMF arrival rate, the report includes Runtime Diagnostics

output to can help you diagnose the behavior. The following example is the SMF arrival rate tracked jobs exception report for jobs that had a higher than expected SMF arrival rate (AIRH187E):

```
SMF Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 11:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 720
Last successful collection time : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval             : 15
Highest STDDEVLOW needed to avoid this exception: 21

Persistent address spaces with high rates:

Job Name      ASID      Enqueue Request Rate
                                     1 Hour      24 Hour      7 Day
-----
TRACKED1 001D      158.00      23.88      22.82      15.82
TRACKED3 0029      11.00       12.43      2.36       8.36
```

Figure 44. SMF arrival rate: tracked jobs higher than expected

The following example is the SMF arrival rate tracked jobs exception report for jobs that had a lower than expected SMF arrival rate (for AIH208E):

```

SMF Arrival Rate Prediction Report

Last successful model time      : 10/10/2010 11:08:01
Next model time                 : 10/10/2010 23:08:01
Model interval                  : 720
Last successful collection time : 10/10/2010 17:41:38
Next collection time            : 10/10/2010 17:56:38
Collection interval             : 15
Highest STDDEVLOW needed to avoid this exception: 21

Persistent address spaces with low rates:

Job Name      ASID      Enqueue Request Rate      Predicted Enqueue Request Rate
              ASID      Rate      1 Hour      24 Hour      7 Day
-----
JOBS4      0027      1.17      23.88      22.82      15.82
JOBS5      002E      2.01      8.34      11.11      12.11

Runtime Diagnostics Output:
Runtime Diagnostics detected a problem in job: JOBS4
EVENT 06: HIGH - HIGHCPU - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 0027 JOBNAME: JOBS4
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 07: HIGH - LOOP - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID: 0027 JOBNAME: JOBS4 TCB: 004E6850
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE APPEARS TO BE IN A LOOP.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
Runtime Diagnostics detected a problem in job: JOBS5
EVENT 03: HIGH - HIGHCPU - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 002E JOBNAME: JOBS5
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 04: HIGH - LOOP - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID: 002E JOBNAME: JOBS5 TCB: 004E6850
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE APPEARS TO BE IN A LOOP.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----

```

Figure 45. SMF arrival rate: tracked jobs lower than expected

Other persistent jobs exception report: PFA issues the SMF arrival rate other persistent jobs exception report when a comparison of a persistent job (that is not being individually tracked) causes an exception when compared to the totals of the other persistent jobs. The exception can be caused by a higher than expected SMF arrival rate or a lower than expected SMF arrival rate. The predictions listed on this report are the predicted rates for the total other persistent jobs group. The list of jobs is only those persistent jobs (not tracked individually) that have a problem and is only generated for a higher than expected SMF arrival rate. No predictions are given for these jobs because PFA does not model individual predictions for jobs that are not tracked individually. If there is more than one job with the same name, four asterisks **** are printed for the ASID in the report. If the report was generated due to a lower than expected SMF arrival rate, the report will include Runtime Diagnostics output which can help diagnose the behavior. The following example is the SMF arrival

rate other persistent jobs exception report for jobs that had an unexpectedly high SMF arrival rate (AIRH188E):

```

SMF Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 17:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval             : 15

Other persistent jobs group:
SMF arrival rate
  in last collection interval   : 502.72
Prediction based on 1 hour of data : 20.27
Prediction based on 24 hours of data : 27.98
Prediction based on 7 days of data  : 31.22

Highest STDDEVLOW needed to avoid this exception:          25

Persistent address spaces with high rates:

Job          Message
Name         ASID      Arrival
-----
PERS1       001E      83.22
PERS2       0038      75.52
PERS3       0039      47.47
    
```

Figure 46. SMF arrival rate: other persistent jobs higher than expected

This example is the SMF arrival rate other persistent jobs exception report for jobs that had a lower than expected SMF arrival rate (for AIRH209E):

```

SMF Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 17:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval             : 15

Other persistent jobs group:
SMF arrival rate
  in last collection interval   : 1.13
Prediction based on 1 hour of data : 20.27
Prediction based on 24 hours of data : 27.98
Prediction based on 7 days of data  : 31.22

Highest STDDEVLOW needed to avoid this exception:          28

Runtime Diagnostics Output:

EVENT 01: HIGH - HIGHCPU - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 0027 JOBNAME: PERS4
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 02: HIGH - LOOP - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID: 0027 JOBNAME: PERS4 TCB: 004E6850
STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID: ++++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE APPEARS TO BE IN A LOOP.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
    
```

Figure 47. SMF arrival rate: other persistent jobs lower than expected

Non-persistent jobs exception report: PFA issues the SMF arrival rate non-persistent jobs exception report when the non-persistent jobs (as a group) cause an exception. This exception is only issued for a higher than expected SMF arrival rate. The SMF arrival rate and predictions listed on this report are calculated for the total non-persistent jobs group. The list of jobs contains only three non-persistent jobs that have high arrival counts. No predictions are given for these jobs because PFA does not model individual predictions for jobs that are not tracked individually. The following example is the SMF arrival rate non-persistent jobs exception report (AIRH191E):

```

SMF Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 11:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time  : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval              : 15
Highest STDDEV needed to avoid this exception:          14

Non-persistent jobs group:
SMF arrival rate
  in last collection interval      : 265.49
  Prediction based on 1 hour of data : 20.27
  Prediction based on 24 hours of data : 27.98
  Prediction based on 7 days of data  : 31.22
Address spaces with high arrivals:

Job Name      ASID      SMF
              ASID      Arrival
              ASID      Counts
-----
NONPERS1 001F      83
NONPERS2 0048      52
NONPERS3 0049      47

```

Figure 48. SMF arrival rate: non-persistent jobs with high counts

No problem report: When no exception is issued or when a total SMF arrival rate exception is issued, the following report is generated. When there is no problem or when an exception occurs because of a higher than expected SMF arrival rate, the list of jobs contains the jobs being tracked individually and the list of jobs can vary from one to ten. The Runtime Diagnostics section is written in the report when the exception is issued because there is a lower than expected SMF arrival rate. The following example is the SMF arrival rate no problem report (AIRH176I) and total system exception report issued due to a high than expected SMF arrival rate (AIRH174E) showing three jobs.

```

SMF Arrival Rate Prediction Report

Last successful model time      : 01/27/2009 17:08:01
Next model time                 : 01/27/2009 23:08:01
Model interval                  : 360
Last successful collection time  : 01/27/2009 17:41:38
Next collection time            : 01/27/2009 17:56:38
Collection interval              : 15

SMF arrival rate
  at last collection interval      : 83.52
  Prediction based on 1 hour of data : 98.27
  Prediction based on 24 hours of data: 85.98
  Prediction based on 7 days of data  : 100.22

Top persistent users:

Job Name      ASID      Enqueue
              ASID      Request
              ASID      Rate
              ASID      1 Hour      24 Hour      7 Day
-----
TRACKED1 001D      18.55      23.88      22.82      15.82
TRACKED2 0028      10.78      0.34       11.11      12.11
TRACKED3 0029      10.36      12.43      2.36       8.36

```

Figure 49. SMF arrival rate: no problem

- **SMF arrival rate in last collection interval:** The actual SMF arrival rate in the last collection interval where the rate is defined to be the number of messages divided by the CPU milliseconds.
- **Predicted rates based on...:** The SMF arrival rates based on one hour, 24 hours, and seven days. If no prediction is available for a given time range, the line is not printed. For example, if the check has been running for 2 days, 7 days of data is not available therefore PFA does not print the "Prediction based on 7 days of data" line.
- **Highest STDDEV needed to avoid this exception:** The highest STDDEV required to avoid the exception issued for all resources in exception status with the values in effect at this comparison.
- **Highest STDDEVLOW needed to avoid this exception:** The highest STDDEVLOW required to avoid the exception issued for all resources in exception status with the values in effect at this comparison.
- **TOO HIGH:** A STDDEV or STDDEVLOW needed to avoid the exception is TOO HIGH if the parameter value required is higher than the maximum allowed.
- **IRRELEVANT:** A STDDEV or STDDEVLOW needed to avoid the exception is IRRELEVANT if the parameter's value was not relevant in the comparison algorithm used for the values in the current comparison.
- **Runtime Diagnostics Output:** Runtime Diagnostics event records to assist you in diagnosing and fixing the problem. See the topic on "[Specific Runtime Diagnostics event reports](#)" on page 41 in Chapter 5, "[Runtime Diagnostics](#)," on page 35.
- **Job Name:** The name of the job that has SMF arrivals in the last collection interval.
- **ASID:** The ASID for the job that has SMF arrivals in the last collection interval.
- **SMF Arrival Rate:** The current SMF arrival rate for the job.
- **SMF Arrival Count:** The SMF arrival rate, from the last interval report, for the non-persistent job.
Note: The "SMF Arrival Count" field is unique to the non-persistent jobs exception report.
- **Predicted SMF Arrival Rate:** The predicted SMF arrival rate based on one hour, 24 hours, and seven days of data. If PFA did not previously run on this system or the same jobs previously tracked are not all active, there will not be enough data for a time range until that amount of time has passed. Also, gaps in the data caused by stopping PFA or by an IPL might cause the time range to not have enough data available. After the check collects enough data for any time range, predictions are made again for that time range. If there is not enough data for a time range, INELIGIBLE is printed and comparisons are not made for that time range.

Directories

Note: The content and names for these files and directories are subject to change and cannot be used as programming interfaces; these files are documented only to provide help in diagnosing problems with PFA.

pfa_directory

This directory contains all the PFA check subdirectories and is pointed to by the PFADIR= specified in the /etc/PFA/ini file, or, if you do not specify PFADIR=, the home directory of the PFA user.

pfa_directory/PFA_SMF_ARRIVAL_RATE/data

The directory for SMF arrival rate that holds data and modeling results. PFA automatically deletes the contents of the PFA_SMF_ARRIVAL_RATE/data directory that could lead to skewed predictions in the future.

Guideline: If the use of the z/OS image is radically different after an IPL (for instance, the change from a test system to a production system) or if you modify the SMF record types in SMFPRMxx, delete the files in the PFA_SMF_ARRIVAL_RATE/data directory to ensure the check can collect the most accurate modeling information.

Results files

- `systemName.1hr.prediction` - This file is generated by the modeling code for the predictions made for one hour of historical data. It contains predictions for each of the tracked address spaces, the other persistent category, the non-persistent category, and the total system category. It also contains additional information required for PFA processing.
- `systemName.24hr.prediction` - This file is generated by the modeling code for the predictions made for 24 hours of historical data. It contains predictions for each of the tracked address spaces, the other persistent category, the non-persistent category, and the total system category. It also contains additional information required for PFA processing.
- `systemName.7day.prediction` - This file is generated by the modeling code for the predictions made for seven days of historical data. It contains predictions for each of the tracked address spaces, the other persistent category, the non-persistent category, and the total system category. It also contains additional information required for PFA processing.
- `systemName.1hr.prediction.html` - This file contains an .html report version of the data found in the `systemName.1hr.prediction` file.
- `systemName.24hr.prediction.html` - This file contains an .html report version of the data found in the `systemName.24hr.prediction` file.
- `systemName.7day.prediction.html` - This file contains an .html report version of the data found in the `systemName.7day.prediction` file.
- `systemName.prediction.stddev` - The file generated by the modeling code to list the standard deviation of the predictions across the time ranges for each job.

Data store files:

- `systemName.OUT` - The data collection file.

Intermediate files:

- `systemName.data` - The file is used as input to the modeling to track if enough data is available to model.
- `systemName.1hr.data` - The file used as input to modeling code. It contains one hour of historical data.
- `systemName.24hr.data` - The file used as input to modeling code. It contains 24 hours of historical data.
- `systemName.7day.data` - The file used as input to modeling code. It contains seven days of historical data.
- `systemName.1hr.holes` - The file is used to track gaps in data, caused by stopping PFA or by an IPL, for a one hour period.
- `systemName.24hr.holes` - The file is used to track gaps in the data, caused by stopping PFA or by an IPL, for a 24 hour time period.
- `systemName.7day.holes` - The file is used to track gaps in the data, caused by stopping PFA or by an IPL, for the seven day time period.

This directory holds the following log files. Additional information is written to these log files when `DEBUG(1)`.

- `stderr` - This file only contains data if messages are generated by the JVM.
- `stdout` - This file only contains data if messages are generated by the JVM.
- `systemName.1hr.cart.log` - The log file generated by modeling code with details about code execution while one hour of historical data was being modeled.
- `systemName.24hr.cart.log` - The log file generated by modeling code with details about code execution while 24 hours of historical data was being modeled.
- `systemName.7day.cart.log` - The log file generated by modeling code with details about code execution while seven days of historical data was being modeled.
- `systemName.builder.log` - The log file generated by intermediate code that builds the files that are input to modeling with details about code execution.

- `systemName.launcher.log` - The log file generated by launcher code.
- `systemName.1hr.tree` - This file is generated by the modeling code. It contains information about the model tree which was built based on the last one hour of collected data.
- `systemName.24hr.tree` - This file is generated by the modeling code. It contains information about the model tree which was built based on the last 24 hours of collected data.
- `systemName.7day.tree` - This file is generated by the modeling code. It contains information about the model tree which was built based on the last seven days of collected data.
- `systemNameCONFIG.LOG` - The log file containing the configuration history for the last 30 days for this check.
- `systemNameCOLLECT.LOG` - The log file used during data collection.
- `systemNameMODEL.LOG` - The log file used during portions of the modeling phase.
- `systemNameRUN.LOG` - The log file used when the check runs.

pfa_directory/PFA_SMF_ARRIVAL_RATE/EXC_timestamp

This directory contains all the relevant data for investigating exceptions issued by this check at the timestamp provided in the directory name. PFA keeps directories only for the last 30 exceptions and only if the exception is newer than the number of days specified in the EXCDIRSDAYS parameter. Therefore at each exception, if directories exist that are too old, they will be deleted. In addition, if more than 30 exception directories exist, the oldest directory is deleted so that only 30 exceptions remain after the last exception is issued.

- `systemNameREPORT.LOG` - The log file containing the same contents as the IBM Health Checker for z/OS report for this exception as well as other diagnostic information issued during report generation (such as Runtime Diagnostic event records).

pfa_directory/PFA_SMF_ARRIVAL_RATE/config

This directory contains the configuration files for the check.

- `EXCLUDED_JOBS` - The file containing the list of excluded jobs for this check.
- `INCLUDED_JOBS` - The file containing the list of included jobs for this check.

Part 4. Diagnosing by problem type

After you identify the problem type, use the following diagnosis procedures to identify the source and extract symptoms to build a search argument.

Chapter 10. Diagnosing an abend

- “[Overview of an abend](#)” on page 179 includes abend symptoms and examples
- “[Steps for diagnosing an abend](#)” on page 182 contains a flowchart and these steps to guide your diagnosis of an abend:
 1. “[Obtaining the abend and reason code](#)” on page 183
 2. “[Identifying the module and component](#)” on page 187
 3. “[Searching the problem reporting databases](#)” on page 190
 4. “[Gathering additional problem data for abends](#)” on page 191
 - “[Steps for gathering trace data for abends](#)” on page 191
 - “[Steps for collecting additional messages and logrec for abends](#)” on page 193
 - “[Steps for obtaining a dump for the error](#)” on page 194

Overview of an abend

The purpose of this chapter is to guide the diagnosis of an abnormal end (abend). Abends have an associated system completion code to describe the error and most have a reason code to further explain the error. These codes can be found by searching:

- [z/OS MVS System Codes](#)
- The documentation for the particular application that failed. For example:
 - For Language Environment completion codes, see [z/OS Language Environment Runtime Messages](#).
 - For RMF completion codes, see [z/OS RMF Messages and Codes](#).

An abend is classified as follows:

- **Software-detected:**
 - A system code in the form of three hexadecimal digits, possibly with a four byte reason code. For example, ABEND075. A system abend code is issued with the ABEND or CALLRTM macros used to terminate a task or address space when a system service or function detects an error.
 - A user code in the form of a four decimal digits, possibly with a four byte reason code. For example, ABENDU4094. A user code is issued using the ABEND macro to terminate a task or the entire job step. When the highest-level task in a job step ends abnormally, all related tasks or subtasks also terminate. When a subtask terminates, only work running on behalf of the subtask is affected, unless STEP=YES is specified.
- **Hardware-detected:**

Hardware might present a program interrupt or machine check on the execution of an instruction. The operating system detects these hardware problems and presents them as an abend.

Example: An instruction in an application running in storage key 7 references storage assigned to key 0. The difference in storage key causes a protection exception. This exception results in hardware presenting a program interruption code of 0004 to the operating system, which is externalized as ABEND0C4.

Related information:

- [z/Architecture® Principles of Operation, SA22-7832](#)

Symptoms of an abend: You can identify an abend by one or more of the following indicators:

- A symptom dump message on the console, in the system log, or job log can indicate a system or user abend.

Abend analysis

For example, message IEA995I is issued to the console:

System message indicating an abend

Notice the indication of a system completion code.

```
IEA995I SYMPTOM DUMP OUTPUT 731
SYSTEM COMPLETION CODE=EC6 REASON CODE=0000FD18
TIME=13.58.26 SEQ=00724 CPU=0000 ASID=0147
PSW AT TIME OF ERROR 070C4400 A90B111A ILC 2 INTC 78
NO ACTIVE MODULE FOUND - PRIMARY NOT EQUAL TO HOME
NAME=UNKNOWN
DATA AT PSW 290B1114 - 1F001F11 05EFEBEC D2640096
AR/GR 0: 00000000/00000000 1: 00000000/00000000
      2: 00000000/2ED9E4D0 3: 00000007/820A33B8
      4: 00000000/294F31D8 5: 00000000/2ED9E200
      6: 00000000/28AADD08 7: 00000002/00000100
      8: 00000007/7F29ECF8 9: 00000002/7F29EC18
      A: 00000000/00FD8B28 B: 00000000/000000E3
      C: 00000000/006FF390 D: 00000000/00F9C5B8
      E: 00000000/A90B111A F: 00000000/00FFB4C
```

- A system message indicating an SVC dump was requested for an error:

For example, here are some messages that might be issued when a SVC dump is taken for an error:

```
IEA794I SVC DUMP HAS CAPTURED: 357
DUMPID=002 REQUESTED BY JOB (OMVS)
DUMP TITLE=COMPON=BPX,COMPID=SCPX1,ISSUER=BPXMIPCE,MODULE=BPXLK
LCP+1DC2,ABEND=S0422,REASON=083A01A5
```

```
IEA611I COMPLETE DUMP ON DUMP.MVS06.D060320.T162245.S00034
DUMPID=034 REQUESTED BY JOB (ZFS )
FOR ASID (1001)
INCIDENT TOKEN: ORACLE MVS06 03/21/2006 00:22:45
ID = IOEDFS
```

```
IEA911E COMPLETE DUMP ON SYS1.DUMP08
DUMPID=001 REQUESTED BY JOB (RESOLVER)
FOR ASIDS(003B,0001)
INCIDENT TOKEN: CWYPLEX1 CPUX 04/21/2006 14:38:56
ERROR ID = SEQ00046 CPU00 ASID003B TIME09.38.56.1
```

- An application detects an error. One example is the following ISMF panel:

```
ISMF ABEND PANEL
COMMAND ==> _
*****
**
**
**          AN ABEND OCCURRED WHILE EXECUTING ISMF          **
**
**          SYSTEM ABEND CODE:  0C4                          **
**
**          ISMF CANNOT CONTINUE                              **
**
**          PRESS THE ENTER KEY OR USE END TO TERMINATE ISMF  **
**          USE HELP TO DISPLAY A LIST OF COMMON ABEND CODES  **
**
*****
*****
```

Another example is the ISPF panel:

```
-----ERROR RECOVERY-----
COMMAND ==>
* * * * *
```

```

* * * * *
* *
* *      ISPF PROCESSOR ENDED ABNORMALLY
* *
* *
* *
* *
* *      Task ABEND code 0C1
* *
* *
* *
* *
* *      Press ENTER to display primary option menu.
* *      Enter HELP command for list of common ABEND CODES.
* *
* *
* * * * *
* * * * *

```

- A component, function, subsystem or application message indicating an abend occurred through a message. For example, TSO/E message INMR030I that identifies an abend condition:

TSO/E message

```
INMR030I RECEIVE command terminated. ABEND abend_code.
```

- An error is recorded in SYS1.LOGREC record.

For example:

```

ERRORID: SEQ=11696 CPU=0040 ASID=00A1 TIME=12:48:20.3

SEARCH ARGUMENT ABSTRACT

PIDS/5752SCXMS RIDS/IEANUC01#L RIDS/IEAVXALA AB/S013E REGS/0D000 REGS/C009C
RIDS/IEAVXALR#R

SYMPTOM DESCRIPTION
-----
PIDS/5752SCXMS PROGRAM ID: 5752SCXMS
RIDS/IEANUC01#L LOAD MODULE NAME: IEANUC01
RIDS/IEAVXALA CSECT NAME: IEAVXALA
AB/S013E SYSTEM ABEND CODE: 013E
REGS/0D000 REGISTER/PSW DIFFERENCE FOR R0D: 000
REGS/C009C REGISTER/PSW DIFFERENCE FOR R0C:-009C
RIDS/IEAVXALR#R RECOVERY ROUTINE CSECT NAME: IEAVXALR

OTHER SERVICEABILITY INFORMATION

RECOVERY ROUTINE LABEL: IEAVXALR
DATE ASSEMBLED: 96270
MODULE LEVEL: HBB6603
SUBFUNCTION: ACCESS LIST ADD

TIME OF ERROR INFORMATION

PSW: 070C0000 80FF5D00 INSTRUCTION LENGTH: 02 INTERRUPT CODE: 0078
FAILING INSTRUCTION TEXT: 5DB88140 174458C0 022856C0

REGISTERS 0-7
GR: 01381495 00000000 00000002 7FFFBF00 00000000 00000000 7FFFBF10 00000000
AR: 00000000 00000000 00000000 00000002 00000000 00000000 00000002 00000000
REGISTERS 8-15
GR: 00FF3D88 01382494 81380496 00F9B700 80FF5D9C 80FF5D00 81380D58 01BFD620
AR: 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000

HOME ASID: 00A1 PRIMARY ASID: 0002 SECONDARY ASID: 00A1
PKM: 80C0 AX: 0001 EAX: FFFF

```

Often, a system completion code or a wait state code indicate an abend. However, there are some exceptions, so use the following table to help guide your diagnosis of an abend:

Abend analysis

If you receive wait state code	Code represents	Diagnose with	Notes®
X'071'	System failure or the operator initiated a restart.	Chapter 13, “Diagnosing a loop,” on page 225	Find complete explanations of wait state codes in z/OS MVS System Codes .
Abend X'122'	Operator canceled the job, requesting a dump.	Chapter 12, “Diagnosing a job or subsystem hang,” on page 215	This abend might also indicate a loop, see Chapter 13, “Diagnosing a loop,” on page 225 . Find complete explanations of wait state codes in z/OS MVS System Codes .
Abend X'222'	Operator canceled the job, without requesting a dump.	Chapter 12, “Diagnosing a job or subsystem hang,” on page 215	This abend might also indicate a loop, see Chapter 13, “Diagnosing a loop,” on page 225 . Find complete explanations of wait state codes in z/OS MVS System Codes .
X'322'	Job exceeded the time limit specified by the TIME option.	Chapter 13, “Diagnosing a loop,” on page 225	Find complete explanations of wait state codes in z/OS MVS System Codes .
All others		“Steps for diagnosing an abend” on page 182	To find the abend code and reason code, see “Obtaining the abend and reason code” on page 183 .

Steps for diagnosing an abend

About this task

Before you begin: You need to know how to use Interactive Problem Control System (IPCS) and have access to the following:

- SVC dump, SYSMDUMP, or stand-alone dump
- EREP of software records (TYPE=S) from SYS1.LOGREC or IPCS VERBX LOGDATA report from a dump.
- OPERLOG, SYSLOG, Job log, other message log for the time frame of the error

You should also be able to locate and use the following:

- IBM product documentation; see [Chapter 26, “Diagnosis information for z/OS base elements and features,” on page 333](#).

Use the following flowchart to guide diagnosis of an abend.

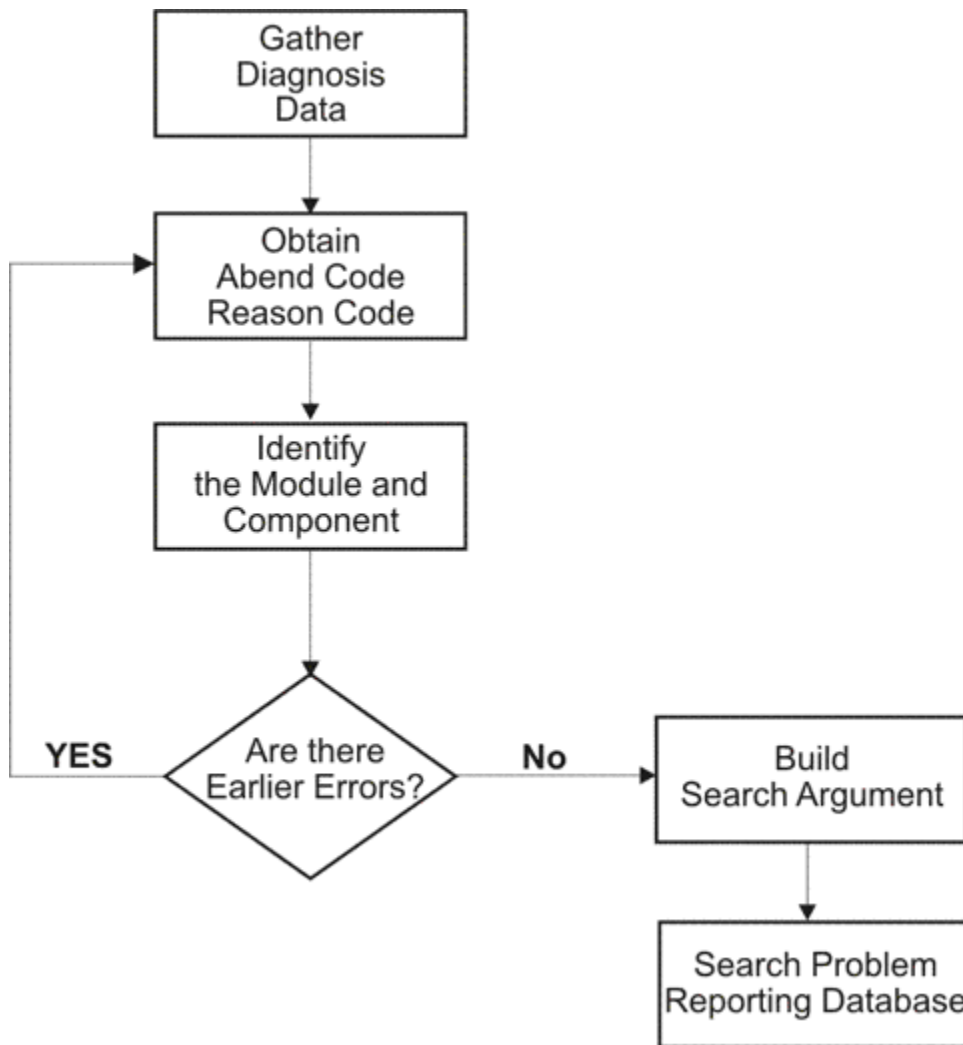


Figure 51. Flowchart for abend analysis

Use the following steps to guide diagnosis of an abend:

Procedure

1. [“Obtaining the abend and reason code” on page 183](#)
2. [“Identifying the module and component” on page 187](#)
3. [“Searching the problem reporting databases” on page 190](#)
4. [“Gathering additional problem data for abends” on page 191](#)
 - [“Steps for collecting additional messages and logrec for abends” on page 193](#)
 - [“Steps for obtaining a dump for the error” on page 194](#)
 - [“Steps for gathering trace data for abends” on page 191](#)

Obtaining the abend and reason code

The abend code indicates the nature of the problem. If you have the abend code, move on to [“Identifying the module and component” on page 187](#).

Steps for obtaining the abend code

Before you begin: You need to know how to access the following:

Abend analysis

- SVC dump, SYSMDUMP, and transaction dump
- Software record in SYS1.LOGREC from the time frame of the error.

If an SVC dump was taken for the error as indicated by IEA794I, IEA611I or IEE911E, use the following IPCS commands to format information and extract the abend and reason code (when available).

- IPCS STATUS FAILDATA

For the abend code, look for AB/S0hhh in the symptoms under the heading SEARCH ARGUMENT ABSTRACT, if present. For the reason code, look for PRCS/hhhhhhhh in the symptoms, or look in the register indicated in the abend code explanation.

Not every dump has error information associated with it. Some dumps, like those requested through the SLIP or DUMP command will not have STATUS FAILDATA output. There are occasions when STATUS FAILDATA will not provide the information you need. Check SYSTRACE and SUMMARY FORMAT (look for RTM2WA SUMMARY) for the abend. If you cannot find it, the dump was probably not requested for an abend. Check the dump title (L Title) and determine why it was not requested. Also, check SYS1.LOGREC, and the SYSLOG and job log for the time frame of the dump.

In the following example, fields AB/S005C and PRCS/00000214 indicate what to extract for the abend and reason code. The free-format search argument is: ABEND05C and RSN00000214.

```
* * * DIAGNOSTIC DATA REPORT * * *

SEARCH ARGUMENT ABSTRACT
RIDS/IEFW21SD#L RIDS/#UNKNOWN AB/S005C PRCS/00000214 REGS/0E01E REGS/0C6D8

Symptom                Description
-----                -
RIDS/IEFW21SD#L        Load module name: IEFW21SD
RIDS/#UNKNOWN          Csect name: #UNKNOWN
AB/S005C                System abend code: 005C
PRCS/00000214          Abend reason code: 00000214
REGS/0E01E             Register/PSW difference for ROE: 01E
REGS/0C6D8             Register/PSW difference for ROC: 6D8

SERVICEABILITY INFORMATION NOT PROVIDED BY THE RECOVERY ROUTINE

Program id
Recovery routine csect name
Recovery Routine Label
Date Assembled
Module Level
Subfunction

Time of Error Information

PSW: 071C2000 83AA2110 Instruction length: 02 Interrupt code: 000D
Failing instruction text: BFFFB148 0A0D98EC B08807FE

Breaking event address: 00000000_00000000
AR/GR 0-1 00000000/00000001_00000000 00000000/00000000_0405C000
AR/GR 2-3 00000000/00000000_7F10C128 00000000/00000000_00000024
AR/GR 4-5 00000000/00000000_006EA338 00000000/00000000_006A3648
AR/GR 6-7 00000000/00000000_006A2D64 00000000/00000000_7F363028
AR/GR 8-9 00000000/00000000_00000000 00000000/00000000_00000000
```

In the following STATUS FAILDATA output, the symptom string indicates a system abend code of X'0C4' with a reason code of X'00000010'

```
SEARCH ARGUMENT ABSTRACT

PIDS/5752SCPX1 RIDS/BPXINPVT#L RIDS/BPXVFPCT AB/S00C4 PRCS/00000010
REGS/C4E10 RIDS/BPXMIPCE#R

Symptom                Description
-----                -
PIDS/5752SCPX1        Program id: 5752SCPX1
RIDS/BPXINPVT#L      Load module name: BPXINPVT
RIDS/BPXVFPCT        Csect name: BPXVFPCT
AB/S00C4              System abend code: 00C4
PRCS/00000010        Abend reason code: 00000010
```

```
REGS/C4E10      Register/PSW difference for R0C:-4E10
RIDS/BPXMIPCE#R Recovery routine csect name: BPXMIPCE
```

OTHER SERVICEABILITY INFORMATION

```
Recovery Routine Label: BPXMIPCE
Date Assembled:        12/19/04
Module Level:          HBB7720
Subfunction:           OpenMVS
```

Time of Error Information

```
PSW: 47043000 80000000 00000000 27EA624C
Instruction length: 02  Interrupt code: 0010
Failing instruction text: 58E0D56C 18F10E0E B2190200
Translation exception address: 00000000_66831000
```

- EREP or VERBX LOGDATA

Use EREP to format software records (TYPE=S) recorded to SYS1.LOGREC for the time frame of the failure or error or format SYS1.LOGREC records from the in storage buffer using the VERBX LOGDATA command.

For information about EREP, see [EREP Reference](#).

To find software records that might be associated with the failure, follow these suggestions:

- Look for the general time frame
- Search on the job name or ASID involved
- Search for the failing component id.

For the abend code, look for AB/S0hhh in the symptoms under the heading SEARCH ARGUMENT ABSTRACT, if present.

For the reason code, look for PRCS/hhhhhhhh in the symptoms, or look in the register indicated in the abend code explanation.

Here is an example of a software logrec entry formatted by EREP or the IPCS VERBX LOGDATA command against an SVC dump. The free-format search argument generated for this entry is: ABEND05C RSN00000214.

```
TYPE: SOFTWARE RECORD REPORT: SOFTWARE EDIT REPORT DAY.YEAR
(SVC 13) REPORT DATE: 115.06
FORMATTED BY: IEAVTFDE HBB7703 ERROR DATE: 096.06
MODEL: 2084 HH:MM:SS.TH
SERIAL: 0D8B9F TIME: 12:14:42.66

JOBNAME: HSMVM0A SYSTEM NAME: VM0A
ERRORID: SEQ=03567 CPU=0000 ASID=003C TIME=12:14:42.6

SEARCH ARGUMENT ABSTRACT

PIDS/5752SC1B4 RIDS/IEFW21SD#L RIDS/IEFDB400 AB/S005C PRCS/00000214 REGS/0E01E
REGS/0C6D8 RIDS/IEFDB402#R

SYMPTOM DESCRIPTION
-----
PIDS/5752SC1B4 PROGRAM ID: 5752SC1B4
RIDS/IEFW21SD#L LOAD MODULE NAME: IEFW21SD
RIDS/IEFDB400 CSECT NAME: IEFDB400
AB/S005C SYSTEM ABEND CODE: 005C
PRCS/00000214 ABEND REASON CODE: 00000214
REGS/0E01E REGISTER/PSW DIFFERENCE FOR R0E: 01E
REGS/0C6D8 REGISTER/PSW DIFFERENCE FOR R0C: 6D8
RIDS/IEFDB402#R RECOVERY ROUTINE CSECT NAME: IEFDB402

OTHER SERVICEABILITY INFORMATION

RECOVERY ROUTINE LABEL: IEFAB4ED
DATE ASSEMBLED: 05223
MODULE LEVEL: UA20441
SUBFUNCTION: DYNAMIC ALLOCATION

TIME OF ERROR INFORMATION
```

Abend analysis

```
PSW: 071C2000 83AA2110 INSTRUCTION LENGTH: 02 INTERRUPT CODE: 000D
Failing INSTRUCTION TEXT: BFFFB148 0A0D98EC B08807FE
```

```
REGISTERS 0-7
GR: 00000000 0405C000 7F10C128 00000024 006EA338 006A3648 006A2D64 7F363028
```

- OPERLOG, SYSLOG, or job log or VERBEXIT MTRACE

Look for message IEA995I or other messages with an abend code in the message text. The message might also give a reason code. For example:

```
IEA995I SYMPTOM DUMP OUTPUT 694
USER COMPLETION CODE=4039 REASON CODE=00000000
TIME=05.07.41 SEQ=33565 CPU=0000 ASID=0247
PSW AT TIME OF ERROR 078D1000 99937C66 ILC 2 INTC 0D
ACTIVE LOAD MODULE ADDRESS=1987A4C0 OFFSET=000BD7A6
NAME=CEEPLPKA
DATA AT PSW 19937C60 - 00181610 0A0D58D0 D00498EC
AR/GR 0: 80C4BB3E/84000000 1: 00000000/84000FC7
2: 00000000/000A1E08 3: 00000000/0002000D
4: 00000000/1992CA78 5: 00000000/000A1954
6: 00000000/00000000 7: 00000000/198758E0
8: 00000000/000A1E08 9: 00000000/000A34C6
A: 00000000/000A1954 B: 00000000/19937B90
C: 00000000/0009EA58 D: 00000000/000A4608
E: 00000000/9992BBAE F: 00000000/00000000
END OF SYMPTOM DUMP
```

In message texts, an abend code can be called a *system completion code* or a *SYS CODE*. A message can show an abend code in the variable text without identifying it as an abend code; use the message explanation to understand the variable text.

- *VERBEXIT SYMPTOM output from a dump*: Format the dump completely, as described in step “1” on page 194 of “Gathering additional problem data for abends” on page 191. Look for AB/S0hhh and PRCS/hhhhhhhh in the symptoms.

In the following VERBEXIT SYMPTOM output, the primary symptom string indicates a system abend code of X'03C' and a return code of X'2D000810'.

Primary Symptom String:

```
RIDS/NUCLEUS#L RIDS/IARYTASS PIDS/5752SC1CR AB/S003C RIDS/IARRR#R
VALU/HC0099680 REGS/0E0B4 REGS/088FA PRCS/2D000810 VALU/CNAGEMENT
```

Symptom	Symptom data	Explanation
RIDS/NUCLEUS#L	NUCLEUS#L	Routine identifier
RIDS/IARYTASS	IARYTASS	Routine identifier
PIDS/5752SC1CR	5752SC1CR	Component identifier
AB/S003C	003C	ABEND code - system
RIDS/IARRR#R	IARRR#R	Routine identifier
VALU/HC0099680	C0099680	Error related hexadecimal value
REGS/0E0B4	0E0B4	Program register
REGS/088FA	088FA	Program register
PRCS/2D000810	2D000810	Return code
VALU/CNAGEMENT	NAGEMENT	Error related character value

The dump does not contain a secondary symptom string.

- Dump title

Look at the dump title; some titles contain the abend and reason codes. Use the DISPLAY DUMP,TITLE or DISPLAY DUMP,ERRDATA to display the dump title and any error information associated with captured dumps, or dumps written to pre-allocated or dynamically allocated data sets. In response to the DISPLAY command, message IEE853I or IEE854I are issued containing the requested information. Look in the IEE853I or IEE854I message replies for the abend code, reason codes and the registers. For example:

```
IEE853I 12.54.26 SYS1.DUMP TITLES 939
SYS1.DUMP DATA SETS AVAILABLE=000 AND FULL=002
CAPTURED DUMPS=0000, SPACE USED=00000000M, SPACE FREE=00000200M
DUMPO0 TITLE=ABEND=S0C4,RC=0010,COMPON=SDSF-ESTAE,COMPID=5647-A01
```

```
,ISSUER=ISFSTAE,SDSF ABEND ROUTINE
DUMP TAKEN TIME=17.49.48 DATE=nn/nn/nnnn
DUMP01 TITLE=ABEND=S0C4,RC=0010,COMPON=SDSF-ESTAE,COMPID=5647-A01
,ISSUER=ISFSTAE,SDSF ABEND ROUTINE
DUMP TAKEN TIME=nn.nn.nn DATE=nn/nn/nnnn
```

For a dump that has been copied from the SYS1.DUMPxx data set or for an SVC dump you are viewing in IPCS, use the IPCS LIST TITLE subcommand to obtain the dump title.

You know you are done when you locate the abend and reason codes. You can then look up a description of the abend code using the product documentation and follow the recommendations.

Related information:

- For IBM product documentation, see [Chapter 26, “Diagnosis information for z/OS base elements and features,”](#) on page 333.
- For information about the IPCS STATUS subcommand, see [z/OS MVS IPCS Commands](#).
- For information about the logrec data set, see [z/OS MVS Diagnosis: Tools and Service Aids](#).
- For information about abend codes, see the product documentation. For example:
 - [z/OS MVS System Codes](#)
 - [z/OS UNIX System Services Messages and Codes](#)
 - [z/OS Communications Server: IP Messages Volume 1 \(EZA\)](#)
 - [z/OS Communications Server: IP Messages Volume 2 \(EZB, EZD\)](#)
 - [z/OS Communications Server: IP Messages Volume 3 \(EZY\)](#)
 - [z/OS Communications Server: IP Messages Volume 4 \(EZZ, SNM\)](#).

Identifying the module and component

In some cases, the abend code indicates the source of the problem and how to repair it. However, there are cases when you cannot identify if the problem was caused by a z/OS, a vendor, or an installation problem. In this case, you must analyze the abend code to see which module or component was involved to conduct a more granular search for a known problem.

Steps for identifying the module and component

About this task

Before you begin: You need to know how to use IPCS and also have access to the following:

- SVC dump, SYSMDUMP, and SADUMP
- Job log, system log, OPERLOG or application error log.
- Master trace

You should also be able to locate and use:

- IBM product documentation, see [Chapter 26, “Diagnosis information for z/OS base elements and features,”](#) on page 333.

Perform the following steps to identify the module that caused the abend and the offset of the failing instruction into the module.

Procedure

1. Do one of the following, depending on the problem data available.
 - For an SVC dump, obtain the component name, component identifier, and module name from the dump title.

Abend analysis

- Obtain the PIDS/cccccccc and RIDS/cccccccc symptoms from the search argument. PIDS is the program identifier, which is the four character product identifier and the five character component identifier. RIDS identifies the module.

If the search argument in an SVC dump does not identify the program and module or if the problem involves multiple abends, analyze the dump for the failing task or service request. See *z/OS MVS Diagnosis: Tools and Service Aids* for information about analyzing an SVC dump.

- Obtain the address in the right half of the program status word (PSW) in STATUS FAILDATA dump output. The zArchitecture PSW is 128 bits (16 bytes). The combination of bits 31 and 32 show the current addressing mode. These are the low order bit of the left half and the high order bit of the right half of the PSW. The meanings are as follows:

- 00 - 24-bit mode
- 01 - 31-bit mode
- 10 - Invalid
- 11 - 64-bit mode

In some places the PSW is shown in a shorter 64-bit (8 bytes) form. This indicates that the addressing mode is 24-bit or 31-bit and the PSW is shown in an ESA/390 form. In that case bit 31, the low order bit in the first half, shows the addressing mode:

- 0 - 24-bit mode
- 1 - 31-bit mode

Subtract the instruction length from the PSW address to obtain the address of the failing instruction. Do not subtract the instruction length in the following cases; the failing instruction is at the PSW address.

- Program check interruptions for which the processing of the instruction identified by the old PSW is nullified. See *z/Architecture Principles of Operation* for the interruption action. Some examples are:
 - Page translation exception: interrupt code = 0011
 - Segment translation exception: interrupt code = 0010
 - Access register translation exception

The following interrupt codes result in the operation being nullified:

- LFX translation exception = 0026
 - LSX translation exception = 0027
 - ASX-translation exception = 0021
 - ASTE-validity exception = 002B
 - ASTE-sequence exception = 002C
 - LSTE sequence exception = 002E
 - ASTE instance exception = 002F
 - Region-first-translation exception
 - Region-second-translation exception
 - Region-third-translation exception
- Obtain the PSW and registers for the error from the STATUS CPU REGISTERS subcommand.

In the following STATUS CPU REGISTERS output, the address can be found in the second half of the PSW. Note that this presentation uses both the first (bit 32) and last (bit 63) bits in the PSW shown combine to indicate addressing mode. Bit 64 will be on when the PSW represents code running in 64 bit mode.

```
CPU STATUS:
PSW=070C1000 83D00B72 (RUNNING IN PRIMARY, KEY 0, AMODE 31, DAT ON)
DISABLED FOR PER
```

```

ASID(X'0015') 03D00B72. DATSVY02+03CA IN EXTENDED PRIVATE
ASCB21 at F9CD80, JOB(DAESVY01), for the home ASID
ASXB21 at 6FE038 for the home ASID. No block is dispatched
HOME ASID: 0015 PRIMARY ASID: 0015 SECONDARY ASID: 0015
GPR VALUES
  0-3 00000000 03D017B0 00000000 03D01A12
  4-7 03D00EC1 03D00CE8 006D4FF8 FD000000
  8-11 03D025BF 83D007A8 03D015C0 03D017A7
  12-15 03D01830 03D015C0 03D019EB 03D00DA9
IEA11015I The requested ALETs are zero.
CONTROL REGISTER VALUES
  0-3 5EB1EE40 00C0407F 002B5040 00800015
  4-7 00000015 01756540 FE000000 00C0407F
  8-11 00000000 00000000 00000000 00000000
  12-15 01F7C27F 00C0407F DF881755 7F704008
THE PRECEDING STATUS CPU INCLUDED THE REGS OPTION

```

Example for for 31-bit:

```

PSW: 040C0000 816B65A6 Instruction Length: 04 Interrupt Code: 0011
Failing instruction text: 58F0C030 50F0B222 5BF0B240
Translation Exception Address: 00000000_7F37B003

```

Example for 64-bit:

```

Time of Error Information:
PSW: 04046001 80000000 00000000 0178F356
Instruction length: 04 Interrupt code: 0004
Failing instruction text: 000A5023 00005032 00044172

```

STATUS CPU REGISTERS supplies the name of the failing module and its offset without having to do a separate WHERE subcommand.

2. Do one of the following:

- If analyzing the dump interactively, use the instruction address in an IPCS WHERE subcommand to obtain the name of the load module and the offset of the address into the load module. If the module name is not preceded with IEANUC01, then IPCS has given the load module name. If you enter the STATUS CPU REGISTERS subcommand, a WHERE is performed automatically.

Use the AMBLIST service aid to list the CSECTs in the load module. Use the offset into the load module to identify the CSECT containing the failing instruction. Then subtract the starting address of the CSECT from the instruction address to obtain the offset into the CSECT.

For instructions on using the AMBLIST service aid, see [AMBLIST](#) in *z/OS MVS Diagnosis: Tools and Service Aids*.

- If the WHERE command does not present a module name, follow this procedure:
 - Using the IPCS Browse panel, enter the PSW instruction address and ASID obtained from the time of error information. Browse backwards looking for the CSECT name eye-catcher. IBM module names are generally followed by an assembly date and a product identifier or PTF level, such as HBB7720; most eye-catchers are at the beginning of a module, but some are at the end.

3. The module prefix identifies the component, subsystem, or product, if provided by IBM. See [Identifying modules, components, and products](#) in *z/OS MVS Diagnosis: Reference*.

For example, using the information in the following output from you can determine what component was involved in an error from the module prefix. The ADY module prefix indicates a DAE-type error.

Module Prefix	Component Name	Product ID	Component ID
ADF	TSO and TSO/E session manager	5665	28505
ADY	Dump analysis and elimination (DAE)	5752	SC143

4. Continue diagnosis as follows:

- For a z/OS component, continue with [“Steps for searching the problem reporting databases” on page 190](#).

Abend analysis

- For an IBM subsystem or product, continue diagnosis with the subsystem or product. See [Chapter 26, “Diagnosis information for z/OS base elements and features,”](#) on page 333 for a listing of components and products.
- For an installation-provided program, including an installation exit routine, continue diagnosis with that program, using the dump for the abend.

Results

Related information:

- See [“Searching problem reporting databases”](#) on page 8 for more information.
- See [z/OS MVS Diagnosis: Tools and Service Aids](#) for dump analysis of a problem in an installation-provided program.
- See [z/OS MVS IPCS Commands](#) for information about IPCS subcommands.

Searching the problem reporting databases

Search arguments are used to search problem reporting databases. If the problem you are diagnosing was already reported and the symptoms are in the database, the search produces a match. Searching is an iterative process; you might need to gather additional data and continue your search.

Steps for searching the problem reporting databases

1. Search the problem reporting database to determine if the problem was previously reported. See [“Searching for a known problem”](#) on page 10.

Use the free-format search argument abstract as a symptom string to determine if the problem already exists. For example ABEND05C RSN00000241.

For example, the following search argument abstract would generate the free-format search argument: ABEND03C RSN2D000810:

```
SEARCH ARGUMENT ABSTRACT
```

```
PIDS/5752SC1CR RIDS/NUCLEUS#L RIDS/IARYTASS AB/S003C PRCS/2D000810 REGS/0E0B  
REGS/088FA RIDS/IARRR#R
```

2. If the search provides no match, remove some symptoms to broaden the search. If the search provides too many symptoms, trying adding some symptoms to limit the scope. Check to see which matches pertain to the operating system environment.
3. If the search finds that the problem was previously reported, request the problem fix. If not, continue with [“Gathering additional problem data for abends”](#) on page 191. Use the problem data gathered there to create more symptoms; use these symptoms in later searches.
4. If you still cannot find the cause of the abend or if the problem is new, report the problem to IBM.

Provide the information in [Chapter 28, “Problem diagnostic worksheet,”](#) on page 339, such as:

- Any changes made to the system recently, preceding when the problem began occurring (for example, PTFs or new products installed or new hardware).
- Problem type: abend
- Search argument
- Dump, formatted by IPCS, online or printed
- Failing input request: macro, command, or statement
- SDWAVRA keys, lengths, and contents
- Offset of the failing instruction into the module or CSECT
- Accompanying messages: identifiers and texts

- Logrec report, if used
- All printed output and output data sets related to the problem
- Data on any related problems
- Module name and level
- Name and level of the operating system(s) with a list of program temporary fixes (PTF) applied at the time of the problem and all installation modifications, exits, and products with other than Class A service
- Other problem data developed while using the diagnosis book for the component, subsystem, or program

Gathering additional problem data for abends

About this task

Gathering additional data will increase your chances of finding a match in the problem reporting databases. Use the procedures outlined in this section to create more symptoms; use these symptoms in later searches.

Steps for gathering additional data for abends

It is important to gather the following information:

- The impact of the problem to system or sysplex
- The names of jobs, functions or programs that were running at the time of the error
- The existence of any new software maintenance or hardware changes
- The associated messages from job log, SYSLOG, or OPERLOG
- The related component traces that are active.

This procedure is divided into the following sections:

- [“Steps for gathering trace data for abends” on page 191](#)
- [“Steps for collecting additional messages and logrec for abends” on page 193](#)
- [“Steps for obtaining a dump for the error” on page 194](#)

Before you begin any of these tasks: Complete the steps in [Chapter 10, “Diagnosing an abend,” on page 179](#).

Steps for gathering trace data for abends

Use the following steps to gather additional data from system trace table:

1. Analyze the system trace table, which is formatted by the SYSTRACE CURRENT ERROR subcommand. A system trace provides a record of system events. Use it to create a picture of the processing occurring at the time of the error.
2. Starting at the end of the trace, back up to the entry for the abend being diagnosed:
 - *SVC D or *SVCE D in the IDENT CD/D columns
 - The abend code in the right 3 bytes in the UNIQUE-3 column
 - The reason code in the UNIQUE-1 column

Example: In the following SYSTRACE output, the *SVC D indicates that an abend code has been loaded for processor 02. When examining system trace output, look for RCVY entries that represent entry into a recovery routine following an error or interruption.

```

PR ASID TCB-ADDR IDENT CD/D PSW----- ADDRESS-...
.
.
.
02 000D 006F8E88 SSRV 12D          8120BFD8 006F8E88 000B0000 00000000
00000000
01 002E 006BEE88 SVC 30 070C3000 827FAF36 00000000 00000001 072CFBF4
01 002E 006BEE88 PC ... 0 81157326 00100
01 002E 006BEE88 PT ... 0 811B782C 002E
02 000D 006F8E88 DSP 070C0000 8101A9B0 00000000 0001035C 000295A8
02 000D 006F8E88 *SVC D 070C0000 8101A9B2 80D12090 0001035C 000295A8
02 000D 006F8E88 PC ... 0 811D7626 00506
01 002E 006BEE88 SVCR 30 070C3000 827FAF36 00000000 00000028 0080002E
01 002E 006BEE88 SVC 38 070C3000 827FAF64 00000000 00000028 072CFBF4
01 002E 006BEE88 SVC A 070C1000 80F8146A 00000000 FD000236 80F81468

```

For system trace, when viewing program checks, look for the PGM entry just before a RCVY entry.

```

06 00C4 009FF540 PGM 004 078D3400 A930BD12 00040004 00000000 00000000
06 00C4 009FF540 *RCVY PROG 940C4000 00000004 00000000

```

3. In the entry, note the processor in the PR column, the address space identifier in the ASID column, and the task control block (TCB) address in the TCB-ADDR column. The ASID should be the same as the ASID identified in STATUS FAILDATA or STATUS CPU output.
4. Continue backing up, looking for the following entries:
 - The entry for the system service being processed when the abend occurred, if the abend occurred during system processing. This entry will have SVC, SSRV, or SVCE in the IDENT column and the same ASID as the abend entry.
 - Problem entries, which have an asterisk (*) before the identifier in the IDENT column.
 - Other entries for the same processor, PR column.
 - Other entries for the same address space, ASID column.
 - Repeated requests by a program for one or more system services. This pattern indicates a loop. An enabled loop has multiple types of entries before the pattern repeats. Continue diagnosis with the program requesting the system services.

You should now be able to determine the source of the abend or have more information to search the problem reporting database.

Related information

- For information about IPCS subcommands: SELECT, SUMMARY, VERBEXIT LOGDATA, and VERBEXIT MTRACE, see [z/OS MVS IPCS Commands](#).
- For the format and contents of the CDE, RB, RTM2WA, SDWA, VRAMAP (VRA keys), and TCB, see the version of *z/OS MVS Data Areas* that corresponds to the release you are running in your environment. See the [z/OS Internet library \(www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary\)](http://www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary).
- For the formats of system trace entries, see [z/OS MVS Diagnosis: Tools and Service Aids](#).
- For information about the SLIP command, see [z/OS MVS System Commands](#).
- See [EREP User's Guide](#) for formatting logrec records.
- See [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#) for message IEA995I.
- See [z/OS MVS System Messages, Vol 9 \(IGF-IWM\)](#) for the IOS messages.
- For the PGM parameter, see [z/OS MVS JCL Reference](#).
- For interactive TSO/E commands, see [z/OS TSO/E Command Reference](#).

Steps for collecting additional messages and logrec for abends

Use the following steps to gather additional messages and logrec:

1. Collect and analyze messages and logrec records about the problem. Use the ERRORID from the dump message and time stamps to select messages and software, symptom, and hardware logrec records related to the problem. Look in the following:
 - The job log
 - A TSO/E user's ISPF transaction log or session manager log
 - The hardcopy log, also known as the system log (SYSLOG)
 - VERBEXIT MTRACE dump output, which shows the buffer for system messages
 - VERBEXIT LOGDATA dump output, which formats the logrec buffer
 - The logrec data set, formatted by EREP
2. Look for the following:
 - Symptom dump message IEA995I for a previous, related abend
 - Messages identifying a failing program with a nonzero return code
 - I/O error messages
3. Identify the program being processed when the abend occurred by obtaining the job name from the following:
 - SUMMARY output
 - VERBEXIT LOGDATA output
 - Messages in the job log
 - Messages in VERBEXIT MTRACE output
 - SELECT output

For example, in the following SELECT output, the job name NVASt in address space 0073 contains an error.

ASID	JOBNAME	ASCBADDR	SELECTION CRITERIA
0073	NVASt	00F6B600	CURRENT ERROR

4. If a batch job was being processed, obtain the program name from the PGM parameter on the JCL EXEC statement.

In the following output, the PGM parameter of JCL statement indicates that the program name is UNIONE.

```
//BANK1 EXEC PGM=UNIONE,PARM=@PLANID=1,10S,SHR',
//          REGION=1024K,COND=(8,LE)
//BANKLOG DD DSN=NULLFILE,DISP=SHR
```

5. If interactive work was being processed, use the command being processed to identify the program.
6. Analyze the problem data for multiple problems. Collect data for related problem ERRORIDs that occur in a similar time frame. You can find this data in the logrec data set. The time stamps are a few seconds before or after the time stamp for the abend being diagnosed. The data involves the following:
 - The same job step
 - The same task (TCB) or service request (SRB)
 - The same home address space (ASID)
 - The address spaces involved in cross-memory mode processing

Abend analysis

- The same processor (CPU), if the problem occurred while the system was disabled for input/output (I/O) and external interrupts (EXT), as indicated in STATUS CPU dump output.

Look for the following:

- In the output from IPCS SUMMARY TCBERROR, look at the task completion codes in job step program TCB CMP fields; a nonzero completion code indicates an abend. You are looking at the correct abend if it has an associated RTM2WA. If a related task abended seconds before the abend being diagnosed, check the task's CDE, RTM2WA, and SDWA control blocks for the module name and other data about the abend. The output contains one RTM2WA for each abend being processed.

In the following SUMMARY TCBERROR output, the nonzero CMP field of the TCB indicates an error.

```
TCB: 009F3E88
  CMP..... 940C9000  PKF..... 80          LMP..... FF          DSP..... FF
  TSFLG.... 00          STAB..... 009FD200  NDSP..... 00002000
  JSCB.... 009FF40C  BITS..... 00000000  DAR..... 00
  RTWA.... 00000000  FBYT1.... 00
  Task non-dispatchability flags from TCBFLGS5:
  Secondary non-dispatchability indicator
  Task non-dispatchability flags from TCBNDSP2:
  SVC Dump is executing for another task
```

- In VERBEXIT LOGDATA output or the logrec reports and in messages in all locations, look for previous abends and symptom records for earlier problems that did not cause abends. The previous abend or the earlier problem might have led to the abend being diagnosed.
- Look for the name of the program that called the abending module.
 - The address of the calling program can be in the second half of the PSW stored in the caller's RB, which will precede the running RB, except for branch entries.
 - Determine the linkage conventions of individual save areas. The calling program's address might be in a save area.
 - In register 14 of the first RB.
 - If a command or macro was being processed, obtain the name of the module issuing the command or macro. The name is in the NAME field of the CDE for a request block (RB) for the abending module's TCB.

Check for problems in the calling program. The calling program might have caused the abend being diagnosed.

Investigate the following:

- Many abends relating to the same area of the system.
- Many TCBs with the same abend code.

You should now be able to identify the program causing the abend. If not, go to the next step.

Steps for obtaining a dump for the error

Use the following steps to obtain a dump and collect additional data using IPCS:

1. If a dump was not written for the abend, recreate the problem and obtain a dump by doing one of the following:
 - Set a SLIP to obtain an SVC dump. For example:

```
SLIP SET,C=0C9,JOBNAME=RMF,A=SVCD,END
```

This SLIP trap will request a dump when an ABEND0C9 occurs in the RMF address space. For more information, see the topic on the [SLIP command](#) in *z/OS MVS System Commands*.

- Insert a SYSABEND, SYSUDUMP or SYSMDUMP DD statement in the JCL for the abending job step to obtain an ABEND dump. For more information, see [z/OS MVS JCL Reference](#).

2. Use IPCS to look at the dump. Use IPCS subcommands in the order indicated by the following list. If using IPCS interactively for an SVC dump, respond yes to the IPCS message that asks if summary data can be used by dump access.
 - a. STATUS FAILDATA
 - b. STATUS SYSTEM

Example: In the following STATUS SYSTEM output, AMDSADMP indicates that this dump was scheduled. Also note the date and time the dump was taken.

```

SYSTEM STATUS:
Nucleus member name: IEANUC01
I/O configuration data:
  IODF data set name: IODF.IODF12
  IODF configuration ID: TC4SYST
  EDT ID: 00
Sysplex name: ENGTEST2
TIME OF DAY CLOCK: BE5E67AF 7ED6370E 02/15/2006 14:33:32.124515 GMT
TIME OF DAY CLOCK: BE5E24A1 5B96370E 02/15/2006 09:33:32.124515 local
Program Producing Dump: SADUMP
Program Requesting Dump: AMDSADMP

```

- c. STATUS CPU REGISTERS DATA CONTENTION
- d. STATUS WORKSHEET
- e. SUMMARY FORMAT
- f. VERBEXIT LOGDATA
- g. VERBEXIT SUMDUMP
- h. SYSTRACE
- i. VERBEXIT MTRACE

Example: In the following VERBEXIT MTRACE output, message IEF450I indicates a system abend of X'522' with a reason code of X'00'.

```

0001 007A5F54 N 0000000 AN03 93039 10:26:08.31 00000281
IEA989I SLIP TRAP ID=X13E MATCHED
0001 007A5F54 N 0000000 AN03 93039 10:26:08.34 00000281
IEA989I SLIP TRAP ID=X13E MATCHED
0001 007A5F54 N 0000000 AN03 93039 10:26:08.43 00000281
IEA989I SLIP TRAP ID=X13E MATCHED
0001 007A5F54 N 0000000 AN03 93039 10:26:08.49 00000281
IEA989I SLIP TRAP ID=X13E MATCHED
0001 007A5F54 N 4000000 AN03 93039 10:26:09.21 TSU05807 00000091
IEF450I LASSEC2 AAIRACF AAIRACF - ABEND=S522 U0000
REASON=00000000
0001 007A7430 N 4000000 AN03 93039 10:26:09.45 TSU06038 00000091
IEF450I LAMMLF AAIUSER AAIUSER - ABEND=S522 U0000
REASON=00000000
0001 007A7430 M 4000000 AN03 93039 10:26:09.59 TSU05807 00000090
IEF377I LASSEC2 AAIRACF AAIRACF
0001 007A5F54 E 064 00000090
LASSEC2.SPFL0G1.LIST NOT CATLGD 2
0001 007EC02C N 4000000 AN03 93039 10:26:09.66 TSU05807 00000090
/HASP395 LASSEC2 ENDED
0001 007A79C0 N 0200000 AN03 93039 10:26:10.06 TSU05807 00000081
/HASP250 LASSEC2 IS PURGED

```

- j. Subcommand selected from the list in [Table 19 on page 196](#).
 - k. VERBEXIT SYMPTOM
3. Before the VERBEXIT SYMPTOM subcommand, add other IPCS subcommands, depending on the problem indicated in the abend explanation or accompanying messages. Pick the subcommands from the following list:

Table 19. Summary of IPCS dump subcommands by problem

Problem involves	IPCS dump command
Allocation/unallocation of jobs	VERBEXIT ALCMWAIT
Asynchronous operations manager (AOM)	VERBEXIT AOMDATA 'TRCDUMP'
Auxiliary storage	ASMCHECK VERBEXIT ASMDATA
Availability management	VERBEXIT AVMDATA
Callable service requests	CBFORMAT addr STRUCTURE(CSRCPOOL) CBSTAT addr STRUCTURE(CSRCPOOL)
Communications	COMCHECK
Cross-system coupling facility (XCF)	COUPLE SUMMARY ALL COUPLE DETAIL ALL COUPLE EXCEPTION ALL
Data-in-virtual	DIVDATA SUMMARY ALL
Data lookaside facility of VLF	DLFDATA
Global resource serialization	VERBEXIT GRSTRACE
Input/output	IOSCHECK
JES2	VERBEXIT JES2
JES3	VERBEXIT JES3
Language Environment	VERBEXIT LEDATA
MVS message service (MMS)	VERBEXIT MMSDATA
z/OS UNIX System Services (OMVS)	OMVSDATA SUMMARY OMVSDATA DETAIL OMVSDATA EXCEPTION
Real storage manager (RSM)	RSMDATA SUMMARY RSMDATA EXCEPTION
System resources manager (SRM)	VERBEXIT SRMDATA
Storage Management Subsystem (SMS)	VERBEXIT SMSDATA
Time sharing option (TSO)	VERBEXIT TSODATA
Virtual storage manager (VSM)	VERBEXIT VSMDATA
Virtual lookaside facility (VLF)	VLFDATA

4. Use the RSMDATA SUMMARY output to get a summary of real storage usage in the system. Use the RSMDATA EXCEPTION report to determine where errors might have occurred. The following is an example of RSMDATA SUMMARY output:

```

R S M   S U M M A R Y   R E P O R T

Tot real      Prf real      Below Prf      ...
-----
In configuration . . . . .      131,072      ...      98,234  4,096  4,026      ...
Available for allocation      126,084      93,247  4,093  4,023      ...
Allocated . . . . .      54,127      53,253      184      118      ...
Percent usage . . . . .           42           57           4           2      ...
Common fixed frames . . .      3,291      3,145      19      19      ...
Percent of available . . .           2           3           0           0      ...
Total fixed frames . . .      8,283      -      28      -      ...
Percent of available . . .           6           -           0           -      ...

V=R Region:
First frame number X'00006'
Last frame number X'0004B'
Size (in frames)           70

Total disabled reference (DREF) pages in real:           2,309

Number of shared data pages:
Valid and fixed in real . . .           3
Valid and pageable in real . . .      1,356
On auxiliary storage . . . . .           0

Number of 64-bit common memory pages:
Backed in real . . . . .           513
Fixed in real . . . . .           144
DREF in real . . . . .           256
On auxiliary storage . . . . .           0

```

Figure 52. RSMDATA SUMMARY report

5. Examine the VRADATA output in the STATUS FAILDATA, VERBX LOGDATA or EREP report for an error for additional clues about the error. For some components, the data consists of a key, a length, and the contents.

Example: In the following Variable Recording Area from STATUS FAILDATA output, the VRA key is X'1A' and the length is X'94'.

```

VARIABLE RECORDING AREA (SDWAVRA)

+000  Key: 1A      Length: 94
+002  02000000    08004000    00040001    1D00E610    | .....W. |
+012  E0001300    00000000    00000000    00000000    | \..... |
+022  00000000    00000000    00000000    00000000    | ..... |
+032  00000000    00000000    00000000    0000002B    | ..... |
+042  00100000    00000000    00000000    00000000    | ..... |
+052  00000000    00000000    00000000    FF000000    | ..... |
+062  012DBD64    00000000    01A9CD24    00000C80    | .....Z. |
+072  00000000    019D6690    012DC698    00000000    | .....Fq. |
+082  01B42000    00000000    00000000    00000000    | ..... |
+092  00000000

```

STATUS FAILDATA will not format an SDWA for a dump requested by SLIP. If SDWA data is not in the dump, obtain problem data from STATUS CPU REGISTERS or view the SDUMP 4K SQA buffer. (See *Reading the SDUMPX 4K SQA buffer in z/OS MVS Diagnosis: Tools and Service Aids.*)

You should now have extracted enough problem data to do a search for a known problem, identify the source of the problem, or report the new problem to IBM or the appropriate vendor.

Related information

- For information about the SYSMDUMP DD, see

Abend analysis

- *z/OS MVS Diagnosis: Tools and Service Aids*
- *z/OS MVS IPCS Commands*
- *z/OS UNIX System Services Planning*
- *z/OS MVS JCL Reference*
- *z/OS MVS System Commands*

Chapter 11. Diagnosing a system hang or wait state

- [“Overview of a hang or wait” on page 199](#) includes descriptions of system hangs and wait states and the symptoms you might encounter
- [“Steps for diagnosing a system hang” on page 200](#) contains a flowchart and these steps to guide you through diagnosis of a hang or wait state:
 1. [“Collecting the problem description” on page 202](#)
 2. [“Diagnosing a hang or wait during IPL” on page 203](#)
 3. [“Diagnosing an enabled wait state” on page 205](#)
 4. [“Diagnosing a coded disabled wait state” on page 207](#)
 5. [“Diagnosing a system partitioned from a sysplex because of status update missing” on page 209](#)
 6. [“Searching the problem reporting databases” on page 210](#)
 7. [“Gathering additional data for hangs and waits” on page 212](#)

Overview of a hang or wait

A system hang or wait can occur gradually as a resource contention problem or abruptly when a disabled wait state is loaded for a critical software-detected error. Externally, the following list of symptoms might be noticed:

- A disabled coded wait state is loaded
- A hang during IPL or system initialization
- The consoles can be locked
- There can be contention for system resources
- The system code can be looping.

Note: If the problem is contention or system code looping, use Runtime Diagnostics to diagnose and possibly solve the problem before continuing these steps. See [“Specific Runtime Diagnostics event reports” on page 41](#) in Chapter 5, [“Runtime Diagnostics,” on page 35](#).

When there is a system failure or outage, a stand-alone dump must be taken for problem diagnosis. OPERLOG, SYSLOG, and EREP reports from the time frame of the system outage are also important.

This section will only discuss system hangs and waits. When a job or subsystem is hung, see [Chapter 12, “Diagnosing a job or subsystem hang,” on page 215](#).

Symptoms of a wait or hang: The system enters a wait or the entire system hangs. The terms hang and wait are used synonymously in these procedures. Some symptoms of a hang:

- No response occurs on the user's or system operator's console.
- No communication with the system through the console occurs.
- No response from subsystems (TSO/E, CICS, IMS, DB2, and others) occurs.
- The system does not issue or receive messages on the console.
- A series of messages that indicate waits followed by bursts of activity.
- A message indicating a wait appears on the system console.
- The program status word (PSW) contains X'070E0000 00000000'.
- The job entry subsystem does not respond to any commands. For example, in a JES2 system, enter a \$DI1 command and JES2 does not respond.

There are two types of wait states: enabled and disabled.

Enabled wait

The system stops processing without issuing a wait state code when the dispatcher did not find any work to be dispatched.

A special type of enabled wait is called a `no work wait` or a `dummy wait`. An indication of a dummy wait or no work wait is a PSW of X'070E0000 00000000' and GPRs containing all zeroes. Diagnosis is required for this type of wait only when the system does not resume processing.

The most common causes of an enabled wait are that the system is waiting for:

- Work – the system has no active jobs to process or all active jobs are swapped out.
- Action – an operator reply or other action.
- Missing interrupts – the system is waiting for a critical device, which is busy, not ready, reserved by another system, or has a mount pending. If the system residence (SYSRES) or paging (PAGE) volumes have missing interrupts, the operator may not get a message.
- System resource – work is waiting for a resource, which can be a lock, queue, input/output (I/O) device, page, or device allocation.

Disabled wait with a wait state code

The system issues a wait state code and stops. The operator can see the wait state code on the system console. This wait is called a `coded wait state` or a `disabled wait`. There are two types of disabled wait state codes:

restartable wait state

You can restart the system.

A restartable wait is one of the following:

- An attempt by the operating system to communicate with the operator. When the system cannot send a message to a console, the system can use a restartable wait state to contact the operator and obtain a response.
- A way to preempt processing. For a SLIP trap with an action of wait, the system will issue a message, then enter a restartable wait.
- A symptom of another problem.

non-restartable

You cannot restart the system. After capturing a stand-alone dump, you must reIPL the system.

Steps for diagnosing a system hang

About this task

You must know how to use IPCS and have access to the following types of information:

- Stand-alone dump
- EREP report of SYS1.LOGREC
- OPERLOG or SYSLOG
- The level of z/OS operating system. Use the IPCS CBFORMAT CVT command to find the level of the z/OS.

The following is an example of the CVT output:

```
***** TOP OF
DATA*****
CVT: 00FD4938

-0028  PRODN... SP7.0.6  PRODI... HBB7709  VERID...          MDL..... 2064  RELNO...
038
+0000  TCBP..... 00000218  0EF00... 00FF24EC  LINK..... 00FD48B4  AUSCB... 00FD4F20  BUF..... 00000000  XAPG.....
00FDE310
+0018  0VL00... 00FF63DE  PCNVT... 00FE0CD4  PRLTV... 00FE0B04  LLCB.... 018E50F0  LLTRM... 8146E288  XTLER...
00FE6D10
+0030  SYSAD... 00EEA898  BTERM... 00FEF820  DATE.... 0106114F  MSLT.... 00FD4F48  ZDTAB... 00DAD000  XITP.....
```

```
00FF9740
+0048 0EF01.... 00FF250C VSS..... 0000      VPSM..... 0000      EXIT..... 0A03      BRET..... 07FE      SVDCB....
```

- The state of the system. Use `IPCS STATUS CPU` and note the PSW and mode of each CPU. For example:

```
CPU(X'00') STATUS:
PSW=07060000 00000000 00000000 00000000
      No work wait
      CPU is in no work wait
```

Normally, a wait state code appears in the program status word (PSW) when the operating system enters a wait state. Use this code and the associated reason code to diagnose and fix the problem. Explanations for wait state codes are found in *z/OS MVS System Codes*.

The following steps will guide you through diagnosing a hang or wait:

1. [“Collecting the problem description” on page 202](#)
2. [“Diagnosing a hang or wait during IPL” on page 203](#)
3. [“Diagnosing an enabled wait state” on page 205](#)
4. [“Diagnosing a coded disabled wait state” on page 207](#)
5. [“Diagnosing a system partitioned from a sysplex because of status update missing” on page 209](#)
6. [“Searching the problem reporting databases” on page 210](#)
7. [“Gathering additional data for hangs and waits” on page 212](#)

Use the following flowchart to guide diagnosis of a system hang:

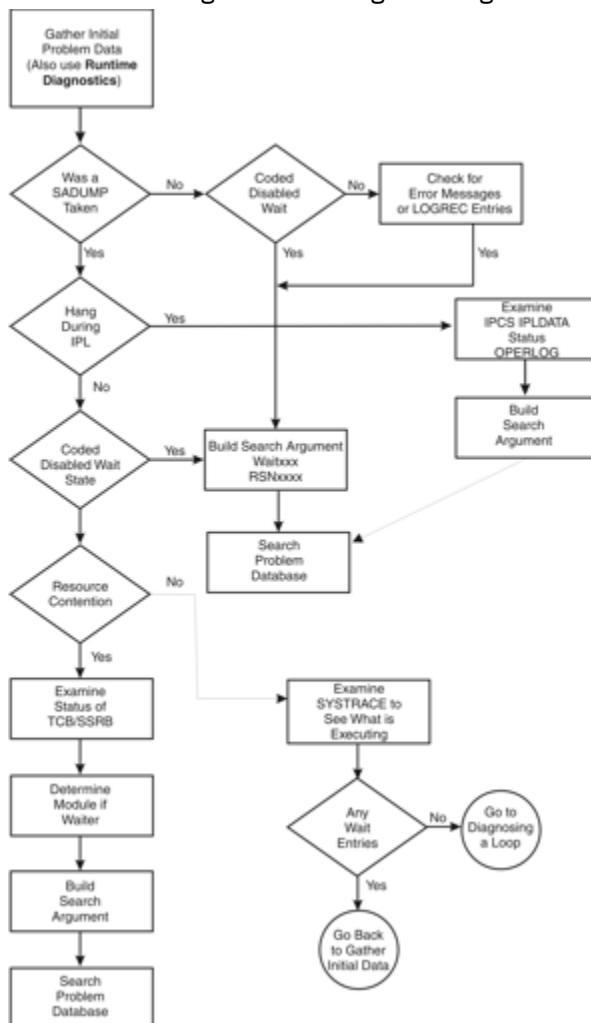


Figure 53. Flowchart for system hang analysis

Collecting the problem description

About this task

The problem descriptions found in [“Gathering diagnosis data” on page 5](#) and [Table 1 on page 7](#) indicate you have a hang, a disabled wait state, or an enabled wait state that needs diagnosis. Before using this procedure, if possible, use Runtime Diagnostics to identify and solve the problem (see [Chapter 5, “Runtime Diagnostics,” on page 35](#)).

Steps for collecting the problem description

Perform the following steps to collect the problem description:

1. Ensure that the symptom descriptions in [“Overview of a hang or wait” on page 199](#) and in [Table 1 on page 7](#) identify the problem is a hang.
If you see the system activity on the console is high and no jobs are being processed, the problem is a loop. Use the procedure in [Chapter 13, “Diagnosing a loop,” on page 225](#).
2. Describe what was happening on the system prior to the hang or wait and record this information in [Chapter 28, “Problem diagnostic worksheet,” on page 339](#). This includes:
 - What is the status for the system on the Hardware Management Console (HMC)?
 - What jobs were started just prior to the hang?
 - What commands were entered and responses received?
 - What recovery procedures did you attempt?
 - What error messages were received?
 - Were there environmental changes? For example, was a new device installed or software maintenance applied.
 - Was the impact to a subsystem like DB2 or the entire system workload?
3. Did the hang during IPL or system initialization? If yes, go to [“Diagnosing a hang or wait during IPL” on page 203](#).
4. Determine the state of the system by entering the state `IPCS STATUS CPU` command. Note the PSW for each CP.
 - a. If every CP is showing a no work wait, go to [“Diagnosing an enabled wait state” on page 205](#). For example:

```
CPU(X'00') STATUS:  
PSW=07060000 00000000 00000000 00000000  
No work wait  
CPU is in no work wait
```

- b. If any CP is showing a disabled coded wait state, go to [“Diagnosing a coded disabled wait state” on page 207](#). For example:

```
CPU(X'01') STATUS:  
PSW=000A0000 800200A2  
Disabled wait state code 00A2 SUPPLMNT INFO 80020  
Wait occurred because system monitor control  
information cannot be read or written.  
ASCB6 at F42700, JOB(XCFAS), for the home ASID  
ASXB6 at 5FDE88 and TCB6G at 5FF500 for the home ASID  
HOME ASID: 0006 PRIMARY ASID: 0006 SECONDARY ASID: 0006
```

- c. If the CP is not showing a no work wait or a coded disabled wait, start diagnosis by checking for resource contention. Go to [“Diagnosing an enabled wait state” on page 205](#).

Diagnosing a hang or wait during IPL

If a hang or wait occurs during IPL or early on during system initialization, obtain a stand-alone dump, SYSLOG, and note the last message issued to the screen or log. The objectives for analyzing the output of a stand-alone dump are:

- Gather symptom data.
- Determine the state of the system.
- Analyze the preceding system activity.
- Find the failing module and component.

Steps for diagnosing a hang or wait during IPL

1. Enter the IPCS **IPLDATA STATUS** command to determine how far along the system is in the IPL or nucleus initialization program (NIP) processing. There is an entry for each initialization routine. The last entry indicates the last initialization routine to run. Use the module name in a search for a known problem. The following example indicates IEAIPL99 ... Page frame table and cleanup as the last entry.

```

*** IPL Statistics ***

IEAIPL10 00:00:00.000 ISNIRIM - Read SCPINFO
IEAIPL20 00:00:01.688 Test Block storage to 2G
IEAIPL11 00:00:00.018 Fast FIND service
IEAIPL31 00:00:00.002 LOAD service
IEAIPL30 00:00:00.000 IPLWTO service
IEAIPL46 00:00:00.164 Read SCHIBs into IPL workspace
IEAIPL49 00:00:00.000 Process Load and Default parameters
IEAIPL50 00:00:00.774 IPL parmlib - process LOADxx and NUCLSTxx
IEAIPL51 00:00:00.019 System architecture
IEAIPL43 00:00:00.032 Find and Open IODF data set
IEAIPL60 00:00:00.008 Read NCRs from IODF
IEAIPL70 00:00:00.208 UIM environment - load CBD and IOS services
IEAIPL71 00:00:00.120 Build DFT for each device
IEAIPL08 00:00:00.007 Read EDT information from IODF
IEAIPL40 00:00:00.093 Read MLTs from nucleus
IEAIPL42 00:00:00.018 Read NMLs from nucleus (IEANynnn modules)
IEAIPL41 00:00:01.388 Read PDS directory entries and CESD records
IEAIPL05 00:00:01.056 Build and sort NUCMAP
IEAIPL02 00:00:03.779 Load nucleus modules
IEAIPL04 00:00:00.020 Allocate PFT and SQA/ESQA
IEAIPL14 00:00:00.000 Build LSQA/ELSQA for Master
IEAIPL06 00:00:00.000 IARMI - RSM blocks, master SGT
IEAIPL09 00:00:00.054 IAXMI - PFT, master RAB, etc.
IEAIPL07 00:00:00.037 Update AMODE for nucleus resident SVCs
IEAIPL03 00:00:00.027 Build UCBs, ULUT, etc.
IEAIPL18 00:00:00.172 Copy and relocate EDT to ESQA
IEAIPL99 00:00:00.465 Page frame table and cleanup

Total IPL Time: 00:00:10.162

```

Figure 54. IPL statistics example

The following NIP example indicates IEAVNPF6 ... Loadwait/Restart as the last entry.

```

*** NIP Statistics ***

IEAVNIP0 00:00:00.024 NIP Base
IEAVNIPM 00:00:00.077 Invoke NIP RIMs
IEAVNPE6 00:00:03.358 Service Processor Interface
IEAVNPF6 00:00:00.023 Loadwait/Restart

```

Hang and wait analysis

2. If the report is complete, the last entries indicate the master scheduler initialization is complete with the total times as in the following example:

```
FINSHMSI 00:00:00.001 Wait for attached CMDs
IEEMB860 00:05:17.024      Uncaptured time: 00:00:00.810

Total Time: 00:07:13.473
```

Enter the IPCS **SELECT ALL** command to verify which system address spaces are active and have completed initialization. For example:

```
ASID JOBNAME  ASCBADDR  SELECTION CRITERIA
-----
0001 *MASTER* 00FD3400 ALL
0002 PCAUTH   00F8DE80 ALL
0003 RASP     00F8DD00 ALL
0004 TRACE   00F8DB80 ALL
0005 DUMPSRV 00F8DA00 ALL
0006 XCFAS   00F81E80 ALL
0007 GRS     00F81D00 ALL
0008 SMSPDSE 00F80400 ALL
0009 CONSOLE 00F80280 ALL
000A WLM     00F4F300 ALL
000B ANTMAIN 00F4F180 ALL
000C ANTAS000 00F4F000 ALL
000D OMVS    00F4DE80 ALL
000F IEFSCHAS 00FC6E80 ALL
0010 JESXCF  00F8B500 ALL
0011 ALLOCAS 00F8B380 ALL
0012 IOSAS   00F97280 ALL
0013 IXGLOGR 00F97100 ALL
0014 JES2    00FC0D80 ALL
```

3. Do a search using symptoms that include the last:

- Initialization routine to run
- Message that was issued (to the log or screen)
- Address space to initialize

Table 20. Common wait states that occur during IPL. Where possible, this table contains an example of the wait state code, reason code, explanation, example, and where to find more information.

Wait state code	Reason code	Explanation	Find information in:
X'064'	X'005'	Indicates an ABEND was issued during NIP. To diagnose using a SADUMP, enter the SYSTRACE ALL command, go to the end of the output and enter the FIND *SVC PREV command to locate the ABEND issued.	Chapter 10, "Diagnosing an abend," on page 179
Example: WAIT X'064' RSNX'005'			
<pre>00 0001 00000000 *SVCE D 040C1000 814E1EE2 00000010 84000000 84878000 10000201 00000000 0001 0001 BEC5A7BA2DAC255B 00800004 00400000</pre>			
X'064'	X'009'	Indicates a program check occurred during NIP. To diagnose using a SADUMP, format the LCCA by entering the CBFORMAT LCCAx command (where x is CP ID).	Chapter 10, "Diagnosing an abend," on page 179

Table 20. Common wait states that occur during IPL. Where possible, this table contains an example of the wait state code, reason code, explanation, example, and where to find more information. (continued)			
Wait state code	Reason code	Explanation	Find information in:
Example: WAITX'064' RSNX'009'			
<pre> LCCA: 00F81000 +0000 LCCA..... LCCA CPUA..... 0041 CAFM..... 4000 PGR1..... 00000000 00000000 00000000 00000000 00000000 +001C 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 +0048 PGR2..... 00000000 08256D70 07BCBCF8 07DFC098 00FCC990 00000903 07BB6678 07E0717C 00000005 00000005 00000000 +0074 07EB2B5F 012B2AE0 08256F90 08256000 08257148 PPSW..... 070C1000 812B2B58 PINT..... 00040011 +0094 PVAD..... 08257000 CR0..... 00000000 PGR3..... 00000000 00000000 00000000 00000000 00000000 00000000 00000000 +00BC 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 P2A0..... 010D061C +00E4 P2A1..... 00000002 P2A2..... 00000000 P2A3..... 00000000 P2A4..... 00000000 P2A5..... 00000000 P2A6..... 00000000 +00FC P2A7..... 00000000 P2A8..... 00000000 P2A9..... 00000000 P2AA..... 00000000 P2AB..... 00000000 P2AC..... 00000000 +0114 P2AD..... 00000000 P2AE..... 00000000 P2AF..... 7F281180 +0120 RSGR..... 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 +014C 00000000 00000000 00000000 00000000 00000000 DSA2..... 00000000 P2C0..... 5FB1EE40 P2C1..... 0331B07F +016C P2C2..... 037B9E00 PX2K..... 8040 PX2S..... 000A PX2A..... 0001 PX2P..... 000A P2C5..... 04EC6280 +017C P2C6..... FE000000 P2C7..... 0331B07F P2C8..... 00000000 P2C9..... 00000000 P2CA..... 00000000 P2CB..... 00000000 +0194 P2CC..... 02FCA403 P2CD..... 0331B07F P2CE..... DF884EC5 P2CF..... 7F70A0B8 PSW3..... 00000000 00000000 </pre>			
X'040'	Not applicable	Indicates an ABEND was issued during NIP. Gather additional information by entering the IPCS STATUS CPU:	From this example, a search argument including symptoms WAIT040 ABEND878 RC10 IEAVNPB2 would be built to check for a known problem.
Example:			
<pre> CPU (X'00') STATUS: PSW=00020000 80000000 00000000 00878040 (Running in PRIMARY, key 0, AMODE 31, DAT OFF) Disabled for PER I/O MCH NIP RIM IEAVNPB2 has failed ABEND=878 REASON=00000010 Register values 0-3 84000000 84878000 03A5EF90 0000E676 4-7 00FCC498 00FCDC98 0393F428 00FCC000 8-11 0187C518 00001000 00000000 00000030 12-15 00000001 00000000 FE000424 00000010 ASCB1 at FCC000, JOB(*MASTER*), for the home ASID ASXB1 at FCC370 and a local SRB for the home ASID HOME ASID: 0001 PRIMARY ASID: 0001 SECONDARY ASID: 0001 </pre>			

Diagnosing an enabled wait state

When the IPCS STATUS CPU command does not show a *no work wait* or a *coded disabled wait*, start diagnosis by checking for resource contention using the following steps.

Steps for diagnosing an enabled wait state

1. Verify that IPCS STATUS CPU report shows every CPU in a no work wait:

```

CPU(X'00') STATUS:
PSW=07060000 00000000 00000000 00000000
No work wait
CPU is in no work wait

```

2. Enter the IPCS ANALYZE EXCEPTION command to look for resource contention.

```
CONTENTION EXCEPTION REPORT
JOBNAME=PMIMTAPE ASID=0065 TCB=007DD0F8
JOBNAME=PMIMTAPE HOLDS THE FOLLOWING RESOURCE(S):
  RESOURCE #0003: There are 0025 units of work waiting for this resource
  NAME=MAJOR=SYSIEFSD MINOR=Q4 SCOPE=SYSTEM
  STATUS FOR THIS UNIT OF WORK:
  This address space is on the SRM IN queue.
JOBNAME=PCICBDTS ASID=0266 TCB=008723A0
JOBNAME=PCICBDTS HOLDS THE FOLLOWING RESOURCE(S):
  RESOURCE #0002: There are 0022 units of work waiting for this resource
  NAME=MAJOR=SYSIEFSD MINOR=Q10 SCOPE=SYSTEM
  STATUS FOR THIS UNIT OF WORK:
  This address space is on the SRM IN queue.
***** END OF DATA *****
```

If resource contention exists, use the IPCS FIND command on the TCB or SSRB address that is identified as the holder of a resource in the analyze exception report to see if the TCB is waiting for any other resources. If found waiting, get the TCB or SSRB address of the holder of that resource and repeat the process until the bottom of the contention chain is reached.

3. Enter an IPCS SUMMARY FORMAT JOBNAME(xyz) for the holder of the resource in contention.
4. Use the IPCS FIND command to locate the TCB or SSRB that is identified as the holder of the resource. If found waiting, get the TCB/SSRB address of the holder of that resource and repeat the process, until the end of the contention chain is reached.
 - a. If the holder of the resource is a TCB, go to [“Examining the TCB status”](#) on page 220.
 - b. If the holder is an SSRB, either find the SSRB in the SUMMARY FORMAT output or format the SSRB control blocks with the IPCS CBFORMAT `srb address STR(SRB)` command. Using the PSW address from the CPSW field, use the IPCS WHERE command or browse storage to find the module name that determines where the SRB was last running.

For example:


```

SSRB:
02451200

+0000 ID..... SSRB    FLNK..... 02452900  ASCB..... 00F72D00  CPAF..... 0000    PASI..... 006C
PTCB..... 007E2250
+0014 EPA..... 00000000  RMTR..... 8142F3D8  PARM..... 00000000  WEB..... 01DF0598  PKF..... 00
FLGS..... 08
+0026 HLHI..... 00    FLGS..... 00    FRR.....
00000000
+0030 FPRS..... 00000000  00000000  00000000  00000000  00000000  00000000  00000000
TRAN..... 00000000
+0054 SAFN..... 0000    TYPE.....
0C
+0058 GPR0..... 00000000  GPR1..... 00000000  GPR2..... 00000041  GPR3..... 01DDE4C8
GPR4..... 00000000
+006C GPR5..... 00000000  GPR6..... 00000000  GPR7..... 00000000  GPR8..... 00000000  GPR9..... 00000000
GPR..... 00000000
+0084 GPRB..... 014228AF  GPRC..... 814218B0  GPRD..... 7F01E388  GPRE..... 0186926E  GPRF.....
00000000

+0098 CPSW..... 070C0000  81421ED0    CPUT..... 00FFFFFF  FF4AF580    TIME..... 00000000
009C49C0
+00B0 XSB..... 02451838  ORMT..... 18FEF6C8  LSA1..... 0206D998  LAA..... 000A0068  LSDP..... 0206DA18
ALOV..... 00000000

```

```

Register
values
0-3 00000000 00000000 00000000
00000000
4-7 00000000 00000000 00000000
00000000
8-11 00000000 00000000 00000000
00000000
12-15 00000000 00000000 00000000
00000000

+0108 DUCT..... 00000000 00000000 00000000 00000000 0007BF00 00000000 00000000 00000000 00000000 00000000
00000000
+0134 00000000 00000000 00000000 00000000
00000000
+0634 00    SSD..... 19413000  OPAS..... C240    OPTC..... 00000000  SUPF..... 80000069
SUSP..... 00
+064A SUSP..... 00000000 00000000 00000000 00000000    SUSP..... 00000000  SUSP..... 00000000
SYNC..... 00000000
+0668 SYNC..... 00000000  SYNC..... 00000000  SYNC..... 00000000  SYNC..... 00000000  SYNC..... 00000000
SUSP..... 00000000
+0680 00000000 00000000  SUSP..... 00000000  SUSP..... 00000000  AFPR..... 00000000  00000000  00000000
00000000
+06AC 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
00000000
+06D8 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
FPFL..... 00
+0740 TRNE..... 00000000 00000000

```

5. Use the module name to search for a known problem.
6. If the search reveals no contention, use the `IPCS SYSTRACE ALL` command to examine the system trace table for the ASIDs that are executing.
7. Use the `FIND WAIT` command in the SYSTRACE report to check for any WAIT type system trace entries.
8. If no WAIT entries exist, there might be contention on CPU resources. Scroll through SYSTRACE noting the ASIDs associated with the entries. If all the entries are for a couple of ASIDs and they are mostly CLKC or EXT type entries, it might indicate a loop. Go to [Chapter 13, “Diagnosing a loop,”](#) on page 225.
9. If WAIT entries are found, then there is no contention on any CPU resource. Talk to the operator to get more specific information on what appears to be hung from an operations perspective. If a specific job or class of jobs are hung (like Batch or TSO users), then get the job name or the specific TSO userid and go to [Chapter 12, “Diagnosing a job or subsystem hang,”](#) on page 215.

Diagnosing a coded disabled wait state

When the `IPCS STATUS CPU` command shows a *disabled coded wait state*, use the following steps to start diagnosis.

Steps for diagnosing a coded disabled wait state

1. Obtain the disabled wait state code from the from the stand-alone dump, if obtained or the Hardware Management Console (HMC). The IPCS STATUS WORKSHEET report from a stand-alone dump often contains a corresponding wait state message. For example:

```

MVS Diagnostic Worksheet

Dump Title: WAIT 0A2 REASON 15A ZOS 1.7 INSTALAC SYSPLEX

CPU Model 2066 Version 00 Serial no. 0123B4 Address 00
Date: 05/22/2006      Time: 20:48:49.817141 Local

Wait State Message Issued at 20:43:48 on Day 142 of 2006:
IXC436W THIS SYSTEM HAS LOST TIME SYNCHRONIZATION WITH THE OTHER
SYSTEMS IN THE SYSPLEX AND
HAS BEEN PLACED INTO A NON-RESTARTABLE
WAIT STATE CODE: 0A2 REASON CODE: 15A

SYSTEM RELATED DATA

CVT  SNAME (154) IEASYSFI  VERID (-18)
      CUCB   (64) 00FD5140  PVTP  (164) 00FFB108    GDA  (230) 0210A278
      RTMCT (23C) 00F4FB20  ASMVT (2C0) 00FD75D8    RCEP (490) 01907F18

CSD  Available CPU mask: 8000  Alive CPU mask: 80000000  00000000
      Number of active CPUs: 00000001

PROCESSOR RELATED DATA

NAME          OFFSET | CPU 00
-----|-----
PSW at time of dump | 00020000
                   | 80000000
                   | 00000000
                   | 0015A0A2
  
```

Figure 55. IPCS STATUS WORKSHEET report from a stand-alone dump

2. Find the wait state code using [z/OS MVS System Codes](#). If there is no recommended action:
 - a. Look up the wait state code in the [System completion code to module table](#) in [z/OS MVS System Codes](#).
 - b. Use the module name to identify the component using the [Identifying modules, components, and products table](#) in [z/OS MVS Diagnosis: Reference](#).
3. Perform a search using the wait state code, reason code, module name and component identifier to look for a known problem. If you cannot find a match, report the problem to IBM.

Table 21. Common disabled wait states . This table contains examples where possible for illustrative purposes.			
Wait state code	Reason code	Explanation	More information:
X'01B'		<p>This is a restartable wait state that is loaded when an active SLIP trap requests an action of WAIT. The conditions specified on the SLIP command are met. The system enters a wait state, as requested. Information about the environment when the slip trap matches is presented in message IEE844W. For example:</p> <pre> Message IEE844W: IEE844W SLIP TRAP 0001 MATCHED. ACTION=WAIT TYPE=PER. PER INFO: 207F: 010FEED2 PSW: 440cc000 810FEED6 CR 3-4 80000001 00010001 GR 0-3 00000000 01DA9600 00000041 01F993B8 4-7: 01EA9600 00000000 010FEE80 00F8C000 8-B: 00FD6F80 010FDA8 00FCD680 9598947F C-F: 01F998E0 014E7040 01EA9140 0000000C AR 0-3: 00000000 00000000 00000000 00000000 4-7: 00000000 00000000 00000000 00000000 8-B: 00000000 00000000 00000000 00000000 C-F: 00000000 00000000 00000000 00000000 RESTART THE SYSTEM TO CONTINUE </pre>	<p>The topic on SLIP problem data in the SLIP work area in z/OS MVS Diagnosis: Tools and Service Aids.</p>
X'040'		The system ended a task during nucleus initialization program (NIP) processing.	Chapter 10, "Diagnosing an abend," on page 179
X'064'	X'005'	Indicates that an ABEND was issued during NIP. To diagnose using a SADUMP, do a SYSTRACE ALL, max PF8 to the bottom of the output and do a 'F *SVC PREV' to locate the ABEND issued.	
EXAMPLE: WAITX'064' RSNX'005'			
<pre> 00 0001 00000000 *SVCE D 040C1000 814E1EE2 00000010 84000000 84878000 10000201 00000000 0001 0001 BEC5A7BA2DAC255B 00800004 00400000 </pre>			
X'064'	X'009'	<p>A program check occurred during nucleus initialization program (NIP). Message IEA304W explains the wait state and entry code.</p> <p>If the message does not appear on the console, you can find the message in the wait state message area (WSMA). The WSMA is described in <i>z/OS MVS Data Areas</i> in the <i>z/OS Internet library</i> (www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary).</p> <p>To diagnose using a stand-alone dump, format the LCCA by entering the CBFormat LCCAx command (where x is the CPU ID).</p>	Chapter 10, "Diagnosing an abend," on page 179
X'0A2'	X'004'	The operator entered the VARY XCF, sysname, OFFLINE command to remove the system from the sysplex.	Ask the operator or system programmer why the system was varied out of the sysplex before continuing with diagnosis.
	X'104'	I/O is prevented because a system is fenced.	Go to "Diagnosing a system partitioned from a sysplex because of status update missing" on page 209
	X'10C'	Cross-system coupling facility (XCF) or cross-system extended services (XES) encountered an unrecoverable error and stopped the system. The system also issues this wait state in response to an operator request to stop the system. For information about diagnosing sysplex problems see, z/OS MVS Diagnosis: Reference .	"Diagnosing an enabled wait state" on page 205

Diagnosing a system partitioned from a sysplex because of status update missing

A system that is partitioned from the sysplex because of status update missing is really indicating a system hang or wait. A system in a sysplex indicates its health by updating a timestamp value on the *SYSPLEX* couple dataset every second. If the timestamp is not updated for the failure detection interval as

defined in the COUPLExx parmlib member in use (85 seconds is the recommended value), the system will be partitioned from the sysplex. Use the following steps to guide your diagnosis:

Steps for diagnosing a system partitioned because of status update missing

When partitioned, the system is be put into a X'0A2' wait state with one of the following reason codes:

- X'104' if a System Failure Management (SFM) Policy is active and IO is fenced.
- X'10C' if the status update missing.

1. Verify the system was partitioned from the sysplex as indicated by message IXC101I, which is in the OPERLOG:

```
IXC101I SYSPLEX PARTITIONING IN PROGRESS FOR SYS22 REQUESTED BY
XCFAS. REASON: SFM STARTED DUE TO STATUS UPDATE MISSING
```

2. Verify the system is in a X'0A2' wait state by entering the IPCS STATUS CPU command:

```
CPU(X'01') STATUS:
PSW=000A0000 800200A2
  Disabled wait state code 00A2  SUPPLMNT INFO 8002
Wait occurred because system monitor control information cannot be read or written.
ASCB6 at F42700, JOB(XCFAS), for the home ASID
ASXB6 at 5FDE88 and TCB6G at 5FF500 for the home ASID
HOME ASID: 0006 PRIMARY ASID: 0006 SECONDARY ASID: 0006
```

3. Enter the IPCS COUPLE SYSPLEX EXCEPTION to identify the reason the system is being partitioned.
4. If the reason is that the system entered a coded disabled wait state prior to sysplex partitioning, go to [“Diagnosing a coded disabled wait state”](#) on page 207 otherwise, go to [“Diagnosing an enabled wait state”](#) on page 205.

Searching the problem reporting databases

Search arguments are used to search problem reporting databases. If the problem you are diagnosing was already reported and the symptoms are in the database, the search produces a match. Searching is an iterative process; you might need to use the procedures in [“Gathering additional data for hangs and waits”](#) on page 212 to gather additional data and continue your search.

Steps for searching the problem reporting databases

About this task

Use the following steps to search the problem reporting databases and determine if the problem was previously reported:

Procedure

1. Develop a free-format search argument using the symptoms obtained from the analysis performed. The free-format search argument can include any of the following symptoms:
 - WAITxxx RSNyyyyyyy (where xxx is the disabled wait state code and yyyyyyy is the associated reason code)
 - Module or CSECT name
 - Resource name that was found to be in contention
 - Message ID
 - Component ID
2. If an argument is not provided, use the primary symptom string in VERBEXIT SYMPTOM output, if available, or use the following symptoms:
 - Program identifier: PIDS/cccccccc

- CSECT name or module name: RIDS/cccccccc
 - Wait state:
 - If a disabled wait, with a wait state code: WS/DOhhh
 - If an enabled wait: WS/E0000
 - If ANALYZE EXCEPTION output indicates a lockout: PCSS/LOCKOUT
 - Input request (call, command, macro, statement), if one is associated with the problem: PCSS/cccccccc
 - Symptoms created from information in the STATUS CPU output
3. Select the problem type on the search tool panel, based on STATUS CPU output:

Table 22. Selecting the problem type for STATUS CPU output

Problem Type	STATUS CPU Output
Disabled wait	DISABLED WAIT STATE CODE
Enabled wait	NO WORK WAIT
Enabled wait	DISABLED FOR ccc (not I/O or EXT)
Hang	None of the above

4. If the search finds that the problem was previously reported, request the problem fix.
- Searching is an iterative process. If the search finds no match, you can remove some symptoms or change the symptoms and search again. Continue searching for matches by adding, removing, and changing symptoms using the steps in [“Gathering additional data for hangs and waits”](#) on page 212.
5. If you still cannot find the cause of the hang or wait or if the problem is new, report the problem to IBM using the procedure in [Chapter 27, “Reporting problems to IBM,”](#) on page 335. Record the following problem data in [Chapter 28, “Problem diagnostic worksheet,”](#) on page 339:
- Problem type: disabled wait, enabled wait, or hang
 - Search argument
 - Dump formatted by IPCS online
 - SMF records, if obtained
 - Accompanying messages: identifiers and texts
 - Hard-copy log, beginning 15 to 30 minutes before the problem, or master trace, if not wrapped between the problem and dump
 - Logrec records, beginning 15 to 30 minutes before the problem and edited using the SPOTCHK and TIMESEQ parameters
 - All output data sets related to the problem
 - Data on any related problems
 - Module name and level
 - Name and level of the operating system(s) with a list of program temporary fixes (PTF) applied at the time of the problem and all installation modifications, exits, and products with other than Class A service
 - Other problem data developed while using the procedures in this document or other diagnosis books for the component, subsystem, or program

Results

You know you are done when you find a match for your problem or report the problem.

Related information:

Hang and wait analysis

- See [“Searching problem reporting databases”](#) on page 8 for more information on developing search arguments.
- See [z/OS MVS IPCS Commands](#) for the IPCS subcommands:
 - ANALYZE
 - STATUS
 - VERBEXIT SYMPTOM
 - STATUS CPU

Gathering additional data for hangs and waits

Gathering additional data will increase your chances of finding a match in the problem reporting databases. Use the procedures outlined in this section to gather additional data and continue searching the problem reporting databases.

Steps for gathering messages and logrec for hangs

About this task

Use the following steps to collect and analyze messages and logrec records about the problem.

Procedure

1. Use time stamps to select messages and software, symptom, and hardware logrec records related to the problem. Look in the following:
 - OPERLOG or SYSLOG
 - VERBEXIT MTRACE dump output, which shows the buffer for system messages
 - VERBEXIT LOGDATA dump output, which formats the logrec buffer
 - Logrec data set, formatted by EREP
2. Use the COPYCAPD command to check for any SVC dumps captured in dataspaces that did not have a chance to get written out to a dataset prior to the system hang. For example:

```
Number  Time stamp      Title
-----  -
1  04/10/2006 15:41:18  END OF MEMORY RESOURCE MANAGER HANG DETECTED: TCB = 009BBE88, NAME = .....
2  04/10/2006 15:42:24  END OF MEMORY RESOURCE MANAGER HANG DETECTED: TCB = 009D64B0, NAME = .....
3  04/10/2006 15:44:43  IXC431I XCF STALLED GROUP
4  04/10/2006 15:46:35  JES2/XCF Env on current & other sysplex members via IEADMCJ2
5  04/10/2006 15:51:35  END OF MEMORY RESOURCE MANAGER HANG DETECTED: TCB = 009683F0, NAME = .....
6  04/10/2006 15:52:43  END OF MEMORY RESOURCE MANAGER HANG DETECTED: TCB = 009D17A8, NAME = .....
7  04/10/2006 15:53:51  END OF MEMORY RESOURCE MANAGER HANG DETECTED: TCB = 009A4998, NAME = .....
8  04/10/2006 15:55:00  END OF MEMORY RESOURCE MANAGER HANG DETECTED: TCB = 009A4758, NAME = .....
9  04/10/2006 15:56:21  END OF MEMORY RESOURCE MANAGER HANG DETECTED: TCB = 00985CF8, NAME = .....
```

These SVC dumps can be extracted into dump data sets using the COPYCAPD command and then diagnosed individually. Go to [Chapter 10, “Diagnosing an abend,”](#) on page 179.

Results

Related information:

- See [z/OS MVS IPCS Commands](#) for the IPCS subcommands.
- For formatting of logrec records, see [Recording logrec error records in z/OS MVS Diagnosis: Tools and Service Aids](#).
- See [z/OS MVS Diagnosis: Reference](#) for logrec reports.
- For explanations of the messages, see:
 - [z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#)
 - [z/OS MVS System Messages, Vol 2 \(ARC-ASA\)](#)

- *z/OS MVS System Messages, Vol 3 (ASB-BPX)*
- *z/OS MVS System Messages, Vol 4 (CBD-DMO)*
- *z/OS MVS System Messages, Vol 5 (EDG-GLZ)*
- *z/OS MVS System Messages, Vol 6 (GOS-IEA)*
- *z/OS MVS System Messages, Vol 7 (IEB-IEE)*
- *z/OS MVS System Messages, Vol 8 (IEF-IGD)*
- *z/OS MVS System Messages, Vol 9 (IGF-IWM)*
- *z/OS MVS System Messages, Vol 10 (IXC-IZP)*
- *z/OS MVS Dump Output Messages*
- The message book for a subsystem or program

You know you are done when your search produces a match.

Chapter 12. Diagnosing a job or subsystem hang

Overview of a hang or wait

When a job or subsystem hang occurs, you might notice part of the system is not functioning or that a job is in the system for a long time without processing.

Symptom of a job hang:

- A job remains in the system for a long time and does not end.

Symptom of a subsystem hang:

- A subsystem does not respond to any commands. For example, in a JES2 system, enter a \$DI1 command and JES2 does not respond.

Steps for diagnosing a job or subsystem hang

The following procedures will guide you through diagnosing a job or subsystem hang:

1. [“Gathering additional data for a job or subsystem hang” on page 216](#)
2. [“Determining the status of a hung job or subsystem” on page 217](#)
3. [“Determining if a job is waiting for resources” on page 217](#)
4. [“Determining address space dispatchability” on page 218](#)
5. [“Examining the SRB status” on page 220](#)
6. [“Examining the TCB status” on page 220](#)
7. [“Examining why a job is not running” on page 222](#)

Use the following flowchart to guide diagnosis of a job or subsystem hang:



Figure 56. Flowchart for job or subsystem hang analysis

Gathering additional data for a job or subsystem hang

About this task

It is important gather information about what the job or subsystem was doing at the time of the hang. This includes answering the following questions:

- What recovery procedures did you attempt? For example, was the MODIFY, CANCEL or FORCE command entered?
- What error messages were received at the time of the hang?
- What commands were entered and responses received?
- Were there environmental changes? For example, a new device installed or a software maintenance upgrade.

Record this information in [Chapter 28, “Problem diagnostic worksheet,”](#) on page 339.

Step for gathering additional data

Before you begin: You must know how to use IPCS and understand how to take an SVC dump. For complete information on SVC dump, see using the IEADMCxx parmlib member in the topic on [SVC dump in z/OS MVS Diagnosis: Tools and Service Aids](#).

Request an SVC dump in one of the following ways:

- Use the DUMP command for the hung job and any other related ASIDs.
- If this is a subsystem or system address space, use IEADMCxx, the DUMP command parmlib member. IEADMCxx allows you to specify the collection of dump data without having to remember and identify all the systems, address spaces and data spaces involved. IEADMCxx from SYS1.SAMPLIB defines the dump options for specific jobs, functions, subsystem, as documented in [z/OS MVS Initialization and Tuning Reference](#).
- Specify the following:

```
DUMP COMM=(dumptitle)
Rxx,ASID=(1,xx),SDATA=(GRSQ,SQA,CSA,RGN,TRT,COUPLE,XESDATA,NUC),END
```

Where xx is the ASID associated with the hung job. You can also use JOBNAME=(xyz).

If using a IEADMCxx parmlib member, enter:

```
DUMP COMM=(title),PARMLIB=xx
```

You must also determine the level of the z/OS operating system:

- Use the IPCS CBFORMAT CVT command. For an example, see [“Steps for diagnosing a system hang” on page 200](#).

Determining the status of a hung job or subsystem

Sometimes a loop can appear to be a hang. Use the following steps to identify if the job is running:

Steps for determining the status of a hung job or subsystem

About this task

1. In the SVC dump, look for entries in the system trace table by issuing IPCS **SYSTRACE JOBNAME(name)**. If there are no entries, continue with [“Steps for determining if a job is waiting for resources” on page 218](#).
2. If there are entries, summarize to determine if:
 - All the entries are in SRB mode or all the entries are for one or two TCB's.
 - The PSW addresses are all within a certain range and repeating. (If the PSW addresses in CLKC and EXT type trace entries are all within a specific range, go to [Chapter 13, “Diagnosing a loop,” on page 225](#).)

If both are true, use the IPCS WHERE command to identify which module that the PSW address is pointing to.
3. Use the module name and the symptom *LOOP* to build a search argument to check for a known problem. For details, see [“Extracting problem symptoms and search arguments” on page 8](#).

Determining if a job is waiting for resources

Use the following steps to identify if the job or subsystem is waiting for resources:

Steps for determining if a job is waiting for resources

About this task

1. Enter the IPCS ANALYZE RESOURCE command.
2. Look in the report for the job name that is hung to determine if it is waiting for system resources. If it is waiting for resource, note the resource name and examine the status of the unit of work holding the resource. For example:

```
JOBNAME=SDSF      ASID=003D  SSRB=02FC4900 --> unit of work holding resource
JOBNAME=SDSF      HOLDS THE FOLLOWING RESOURCE(S):

RESOURCE #0002:  There are 0018 units of work waiting for this resource
NAME=LOCAL LOCK FOR ASID 003D
DATA=INTERRUPTED AND NOW DISPATCHABLE
```

3. Do a FIND on the TCB or SSRB address of the holder to check if it is waiting on another resource. Repeat until the end of the contention chain is reached.
4. Choose from the following:
 - a. If the unit of work that is holding the resource is not dispatchable, go to [“Examining the SRB status” on page 220](#) or [“Examining the TCB status” on page 220](#) to determine why.
 - b. If the unit of work that is holding the resource is dispatchable, go to [“Steps for examining why a job is not running” on page 223](#).
 - c. If the job is not holding or waiting for a resource, go to [“Steps for examining address space dispatchability” on page 218](#).

Determining address space dispatchability

If you are working with a complete SVC dump, as indicated by messages IEE911E or IEA611I, examine address space dispatchability using the following steps.

Steps for examining address space dispatchability

Procedure

1. Enter IPCS SUMMARY FORMAT JOBNAME(xyz) to obtain a report for the hung job.
2. Locate the ASCB (F ASCB). For example:

```
ASCB: 00F2A280
+0000 ASCB..... ASCB      FWDP..... 00F2A100  BWDP..... 00F2A400  LTCS..... 00000000  SVRB..... 009FD598  SYNC..... 000001DE
+0018 IOSP..... 00000000  R01C..... 0000      WQID..... 0000      SAWQ..... 04CBAAF0  ASID..... 003D      R026..... 0000
+0028 LLS..... 00      HLHI..... 01      DPH..... 00EF      TCBE..... 00000000  LDA..... 7FF18EA0  RSMF..... C0
+0035 FLG3..... 00      R036..... 0000      CSCB..... 18E89D00  TSB..... 00000000  EJST..... 00000000  22CF88A0
+0048 EWST..... BEC51B1A  4660A320  JSTL..... 000141DE  ECB..... 809FDC80  UBET..... BE89F19D  TLCH..... 00000000
+0060 DUMP..... 009FFE88  AFFN..... FFFF      RCTF..... 01      FLG1..... 00      TMCH..... 00000000  ASXB..... 009FDE88
+0070 SWCT..... 00A7      DSP1..... 80      FLG2..... 00      RSV..... 0000      SRBS..... 000E      LLWQ..... 00000000
+007C RCTP..... 009FE0A8  LOCK..... 4FFFFFFF  LSWQ..... 04CBAAF1  QECB..... 00000000  MECB..... 40000000  OUCB..... 03292100
+0094 OUXB..... 0384E6F8  FMCT..... 0000      LEVL..... 03      FL2A..... 80      R09C..... 00000000  IQEA..... 00000000
+00A4 RTMC..... 00000000  MCC..... 00000000  JBNI..... 00000000  JBNS..... 00FAF60C  SRQ1..... 00      SRQ2..... 00
+00B6 SRQ3..... 00      SRQ4..... 00      VGTT..... 00000000  PCTT..... 00000000  SSRB..... 0000      SMCT..... 00
+00C3 SRBM..... 07      SWTL..... 00000D30  SRBT..... 00000001  301DBF80  LTCB..... 03748628  LTCN..... 00000005
+00D8 TCBS..... 00000000  LSQT..... 00000000  WPRB..... 009FE950  NDP..... EF      TNDP..... FF      NTSG..... FF
+00E7 IODP..... FD      LOCI..... 00000000  CMLW..... 044AA6B8  CMLC..... 00000000  SSO1..... 000000      SS04..... 00
+00F8 ASTE..... 03F05F40  LTOV..... 7FF95000  ATOV..... 7FFF0128  ETC..... 0001      ETCN..... 0001      LXR..... 0001
+010A AXR..... 0000      STKH..... 009FE960  GOEL..... 19021778  GSYN..... 00000000  XTCB..... 009F1598
+0120 CS1..... C0      CS2..... 00      R122..... 0000      GXL..... 03105000  EATT..... 00000000  034077E0
+0130 INTS..... BE89F19D  91A9EA80  LLL1..... 00      LL2..... 00      LLL3..... 00      LL4..... 00
+013C RCMS..... 00000000  IOSC..... 00000067  PKML..... 0000      XCNT..... 01F4      NSQA..... 00000000  ASM..... 03896540
+0150 ASSB..... 03AA4280  TCME..... 00000000  GQIR..... 00000000  R15C..... 00000000  CREQ..... 00000000
+016C RSME..... 03807290  AVM1..... 00      AVM2..... 00      AGEN..... 0000      ARC..... 00000000  RSMA..... 03807128
+017C DCTI..... 0000062E
```

3. Check the following fields in the ASCB control block:

DSP1 (+X'72')

Address space non-dispatchability bits:

- If X'80', an SVC dump is in progress, which does not indicate a problem.
- If X'40', an address space is failing. Check field MCC (ASCB+X'A8') for the memory termination completion code (MCC=8840D000). If so, this job is terminating with an ABEND40D. Extract the associated reason code from field ARC (ASCB+X'174').

- If X'10' or X'32', the address space is logically or physically swapped out. To determine the reason for the swap, issue the IPCS VERBX SRMDATA and FIND the job name (F jobname).

LOCK (+X'80')

Address space local lock word:

- If the word contains X'7FFFFFFF', a TCB or SSRB is suspended holding the local lock. To locate the PSW and registers for the lock holder, enter FIND IHSA. Use the IPCS WHERE command on the address from CPSW (current PSW) in the IHSA to identify the lock holder. Use this module name in a search.
- If the word contains X'0000004X', this identifies that the lock holder is currently executing on CP 4x at the time of the dump. Do an IPCS SYSTRACE CPU(4x), go to the end of the output and scroll back examining the trace entries to see what unit of work is executing. For information on interpreting system trace, see [System trace in z/OS MVS Diagnosis: Tools and Service Aids](#).
- If the word contains X'4FFFFFFF' or X'FFFFFFFF' (the SRB or TCB ready to run id), then the unit of work is on the dispatching queue and waiting to get dispatched. Go to [“Steps for examining why a job is not running” on page 223](#).

Examining the SRB status

Steps for examining the SRB status

1. Enter the IPCS CBFORMAT *srb* address STR(SRB) to format the SRB or SSRB and related control blocks that describe the environment of the system request block (SRB). For example:

```

SSRB:
02451200

+0000 ID..... SSRB      FLNK..... 02452900  ASCB..... 00F72D00  CPAF..... 0000      PASI..... 006C
PTCB..... 007E2250
+0014 EPA..... 00000000  RMTR..... 8142F3D8  PARM..... 00000000  WEB..... 01DF0598  PKF..... 00
FLGS..... 08
+0026 HLHI..... 00      FLGS..... 00      FRR.....
00000000
+0030 FPRS..... 00000000  00000000  00000000  00000000  00000000  00000000  00000000
TRAN..... 00000000
+0054 SAFN..... 0000      TYPE.....
0C
+0058 GPR0..... 00000000  GPR1..... 00000000  GPR2..... 00000041  GPR3..... 01DDE4C8
GPR4..... 00000000
+006C GPR5..... 00000000  GPR6..... 00000000  GPR7..... 00000000  GPR8..... 00000000  GPR9..... 00000000
GPRA..... 00000000
+0084 GPRB..... 014228AF  GPRC..... 814218B0  GPRD..... 7F01E388  GPRE..... 0186926E  GPRF.....
00000000

+0098 CPSW..... 070C0000  81421ED0      CPUT..... 00FFFFFF  FF4AF580      TIME..... 00000000
009C49C0
+00B0 XSB..... 02451838  ORMT..... 18FEF6C8  LSA1..... 0206D998  LAA..... 000A0068  LSDP..... 0206DA18
ALOV..... 00000000

Register
values
0-3 00000000 00000000 00000000
00000000
4-7 00000000 00000000 00000000
00000000
8-11 00000000 00000000 00000000
00000000
12-15 00000000 00000000 00000000
00000000

+0108 DUCT..... 00000000 00000000 00000000 00000000 0007BF00 00000000 00000000 00000000 00000000
00000000 00000000
+0134      00000000 00000000 00000000 00000000
00000000
+0634      SSD..... 19413000  OPAS..... C240      OPTC..... 00000000  SUPF..... 80000069
SUSP..... 00
+064A SUSP..... 00000000 00000000 00000000 00000000      SUSP..... 00000000  SUSP..... 00000000
SYNC..... 00000000
+0668 SYNC..... 00000000  SYNC..... 00000000  SYNC..... 00000000  SYNC..... 00000000  SYNC..... 00000000
SUSP..... 00000000
+0680      00000000  SUSP..... 00000000  SUSP..... 00000000  AFPR..... 00000000  00000000  00000000
00000000 00000000
+06AC      00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
00000000 00000000
+06D8      00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
FPFL..... 00
+0740 TRNE..... 00000000 00000000

```

2. Use the PSW address from the CPSW field in the IPCS WHERE command or browse storage to look for a module name to determine where the SRB was last executing.
3. Use the module name and hang to search for a known problem.

Examining the TCB status

Steps for examining the TCB status

1. Enter the IPCS SUMMARY FORMAT JOBNAME(xyz) to generate a report of address space related control blocks representing the job, which includes the ASCB, SRB's, TCB's, RB's, and others.

2. Locate the TCB.

- a. If you have identified a TCB as holding or waiting for a resource, enter the FIND TCB: **xxxxxxx** where **xxxxxxx** is the virtual address of the TCB.
- b. If you have not identified a TCB, scroll to the end of the SUMMARY FORMAT report to locate the last TCB on the chain.

3. Examine the TCB non-dispatchability bits, ignoring the SDUMP non-dispatchability indicator:

```
TCB:
008FF2A0
+0000 RBP..... 008FF218 PIE..... 00000000 DEB..... 00000000 TIO..... 008D1FE8 CMP..... 00000000
TRN..... 40000000
+0018 MSS..... 7FFFE700 PKF..... 80 FLGS..... 00000004 01 LMP..... FF
DSP..... FF
+0024 LLS..... 00000000 JLB..... 008FCDD8 JPQ.....
008FD000
.....

Register
values
0-3 00000001 25400F88 00000040
00000000
4-7 008DFAB0 0000000A 008D1FE0
25401048
8-11 008FC138 008FF560 00000000
A5400D0A
12-15 830048B2 00006F60 00000311
808FF218
+012C EAE..... 7FF4E8F8 ARC..... 00000000 GRES..... 00000000 STCB..... 7F6F8A68 TTIME.... 00000000
007A1440
+0144 CELAP.... 00006FF0 R148..... 0000 RBYT1.... 00 LEVEL.... 03 BDT..... 00000000 NDAXP....
00000000
+0154 SENV.....
00000000
Task non-dispatchability flags from
TCBFLGS4:
Top RB is in a wait --->
important
Task non-dispatchability flags from
TCBFLGS5:
Secondary non-dispatchability
indicator
Task non-dispatchability flags from
TCBNDSP2:
SVC Dump is executing for another task ----> ignore, always on in an svcdump
```

4. If RB is in a WAIT or with no flags on:

- a. Do a FIND ACTIVE to locate the RB control blocks (these can be PRBs, SVRBs, IRBs or SIRBs).
- b. Look at the last RB formatted. Verify that the wait or suspend count in the high order byte of the RBLINK is greater than zero. Extract the PSW address from the OPSW field. For example:

```
ACTIVE RBS
PRB: 008FF218
-0020 XSB..... 7FFFDCA0 FLAGS2... 80 RTPSW1... 00000000 00000000 RTPSW2... 00000000
00000000
-0008 FLAGS1... 40800003 WLIC..... 00020001
+0000 RSV..... 00000000 00000000 SZSTAB... 00110082 CDE..... 008FD000 OPSW..... 070C0000
A5400F00
+0018 SQE..... 00000000 LINK..... 018FF2A0
+0020 GPR0-3... FD000008 00006FF8 00000040 008DFAD4
+0030 GPR4-7... 008DFAB0 008FD0C8 008D1FE0 FD000000
+0040 GPR8-11.. 008FC138 008FF560 00000000 008FD0C8
+0050 GPR12-15. 830048B2 00006F60 008FC184 008FC168
64-Bit GPRs from the RB/XSB

Left halves of all registers contain zeros
0-3 FD000008 00006FF8 00000040 008DFAD4
4-7 008DFAB0 008FD0C8 008D1FE0 FD000000
8-11 008FC138 008FF560 00000000 008FD0C8
12-15 830048B2 00006F60 008FC184 008FC168
+0060 RSV..... E2C8D9F0 F4D3D740
```

- c. Enter the IPCS **WHERE xxxxxxxx** , using the address from the RBOPSW field with the high order addressing mode bit off, to identify which module or CSECT was last in control. For example :

```
Command ==> ip w 25400f00
***** TOP OF DATA *****
ASID(X'0020') 25400F00. SHR04LP+01F8 IN EXTENDED PRIVATE
ASID(X'0020') 25400F00. AREA(Subpoo1251Key08)+0F00 IN EXTENDED PRIVATE
***** END OF DATA *****
```

- If PSW points into ISGGWAIT, the TCB is waiting on an ENQ resource. Check the storage pointed to from GPR1 saved in this SVRB. It will point to storage containing the ENQ parmlist. The ENQ parmlist +4 points to the major name and +8 points to the minor name associated with the resource. For a mapping of the parmlist, see [SVC 38 \(OA26\)](#) in *z/OS MVS Diagnosis: Reference*.
- If the PSW points to IEAVEWAT, a program call (PC) entered the WAIT as requested. To determine who the requester was, locate the current linkage stack entry for this TCB by entering F LSE: PREV. Enter the IPCS WHERE command on the PSW address from field PSWE to determine who performed the WAIT (PC 30D). Use this module name when doing a search for a known problem. For example:

```
LSE: 7F58B140
GENERAL PURPOSE REGISTER VALUES
00-01... 00000000 00000001 00000000 80095238
02-03... 00000000 00FCB818 00000000 018944B8
04-05... 00000000 7FF6ADB0 00000000 03ECC378
06-07... 00000000 00FCB818 00000000 98F27718
08-09... 00000000 18F29088 00000000 00FEC410
10-11... 00000000 03ECC0D0 00000000 18F24747
12-13... 00000000 18F23748 00000000 7F4FF100
14-15... 00000000 0000030D 00000000 18F27712
ACCESS REGISTER VALUES
00-03... 00000000 00000000 00000000 00000000
04-07... 00000000 00000000 00000000 00000000
08-11... 00000000 00000000 00000000 00000000
12-15... 00000000 00000000 00000000 00000000
PKM..... 8040 SASN..... 0006 EAX..... 0000 PASN..... 0006
PSW..... 07041000 80000000 PSWE..... 00000000 18F27752---> this PSW address indicates
TARG..... 00000000 0000030D MSTA..... 00000000 00000000 who did a PC 30D WAIT
TYPE..... 0D
PC STATE ENTRY
RFS..... 0D80 NES..... 0000
```

5. If the *abnormal wait* non-dispatchability flag is on, check for a subtask that is also abending. Do a FIND on TCB examining the CMP field in the TCB for one that is non-zero. Next, repeat step 3 examining the TCB non-dispatchability bits.
6. If any other TCB non-dispatchability flags are on, determine the name to the TCB non-dispatchability flag set by examining the value of TCBFLGS field, bytes 4 and 5, and field TCBNDSP. Interpret bit settings using the mapping of the TCB in the version of *z/OS MVS Data Areas* that corresponds to the release of z/OS you are running in your environment. See the [z/OS Internet library \(www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary\)](http://www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary). Use the TCB non-dispatchability bit name in a search for a known problem.
7. Build a symptom string including:
 - The word WAIT
 - The module name
 - The component ID identified
 - Any other symptoms. For example, the resource waited on (SYSZTIOT or CMSEQDQ lock) or the TCB non-dispatchability bit name (TCBSTP).

Search for a known problem.

Examining why a job is not running

When a job is not running, but the TCBs or SRBs are dispatchable, determine what is preventing the job from being dispatched. This often occurs when higher priority work is monopolizing the CPs. Use the following steps to examine the activity in the system trace.

Steps for examining why a job is not running

1. Enter the IPCS **SYSTRACE ALL** command.
2. Enter `FIND WAIT` in the SYSTRACE report and identify WAIT type trace entries.
 - a. If you find WAIT entries, there are intervals when the dispatcher did not find any dispatchable work to dispatch. Review the previous sections of [Chapter 12, “Diagnosing a job or subsystem hang,”](#) on [page 215](#) to ensure nothing overlooked.
 - b. If you do not find any WAIT entries, examine which ASIDs are running by scrolling through the IPCS SYSTRACE ALL report.
3. In a IPCS `VERBX SRMDATA` report, compare the service classes and periods of the jobs executing in SYSTRACE to those that are hung.

For example:

```
JOB   ZFS
ASID  01F5
OUCB  02700780 IN      QUEUE
          +11 (NSW) NONSWAPPABLE
          +11 (PVL) PRIVILEGED PROGRAM
          (ASCBRSME) RAX ADDRESS IS 020848C8
                   SERVICE CLASS = SYSSTC
                   WORKLOAD = SYSTEM
                   INTERNAL CLASS= $SRMG00D
                   PERIOD = 01
```

If these are of higher priority, examine where they are running by using the IPCS `WHERE` on several of the PSW addresses from the SYSTRACE report.

4. Search for a known problem using the module, CSECT, or both.

Chapter 13. Diagnosing a loop

Note: Use Runtime Diagnostics to diagnose and possibly solve the loop condition before continuing these steps. See “HIGHCPU - Address spaces with high processor usage (CPU analysis)” on page 43 in Chapter 5, “Runtime Diagnostics,” on page 35.

Overview of a loop

A loop is a repetitive set of instructions being performed by a job or unit of work. A job or function that is looping can appear to be hung or can use a high amount of CP resource and lock out other work from getting service.

There are three types of loops and symptoms for each type:

Disabled loop

A disabled loop is repetitive execution, usually in system level code, with the IO and EXT type interrupts prevented with a PSW mask of X'x4' in the high order byte of the PSW. A disabled loop is bound to one CP in the system. If in a multi-processor environment and resources are held, a spin loop is detected. If on a uniprocessor, a disabled loop will result in a system outage.

Enabled loop

An enabled loop occurs under a unit of work (TCB or SRB) that is executing on behalf of a job or function. It is executing with a PSW that is enabled for I/O and external interrupts with a mask of X'x7' in the high order byte. A unit of work that is looping enabled, is interrupted periodically for IO, EXT or CLKC type interrupts, which are traced in the system trace table.

Spin loop

A spin loop is a timed disabled loop in system code controlled by the installation with specifications in the EXSPATxx (excessive spin condition actions) parmlib member. The system can spin or loop disabled waiting for a resource, such as a lock, to be released by another CP in a multi-processing environment. See *z/OS MVS Initialization and Tuning Reference* for more information on the EXSPATxx (excessive spin condition actions) parmlib member.

Disabled loop symptoms

Disabled loops are easier to identify than enabled loops. Symptoms include:

- System CP usage increases for unexplained reasons.
- There is no communication with the system through the master and alternate consoles.
- Console communications are locked out. To check for communication with the console, enter DISPLAY T command and the system will not respond.

Enabled loop symptoms

Enabled loops allow some or all interrupts. The loops are usually caused by an error in an application program. All or most of the loop is in code running in problem state, but the loop can include system code if any instructions in the loop request system services. An enabled loop can run on more than one central processor. The loop will uselessly consume resources and might take over all system operation.

Additional symptoms include:

- A bottleneck, indicating that the system slows down periodically, thus creating a performance problem.
- A job stays in the system for a long time without changing status or ending.
- Low priority work slows down or stops running (a result of a higher priority enabled loop).

Loop analysis

- System CP usage increases for unexplained reasons or CP usage of an address space is much higher than normal.

Spin loop symptoms

A spin loop occurs when one processor in a multiprocessor environment is unable to communicate with another processor or requires a resource currently held by another processor. The processor that has attempted communication is the *detecting* or *spinning* processor. The processor that has failed to respond is the *failing* processor.

The detecting processor continuously attempts its communication with the failing processor until either:

- It is successful.
- A specified time interval has passed.

When the communication is not successful within this interval, an excessive spin loop *time out* exists. The detecting processor then initiates recovery processing for the condition.

MVS processing for excessive spin-loop conditions can provide recovery without any operator prompts or actions required. The following recovery actions can be defaulted to or specified in the EXSPATxx parmlib member:

SPIN

Continue spinning for another interval to allow the event to complete

ABEND

End the current unit of work on the failing processor but allow the recovery routines to retry

TERM

End the current unit of work on the failing processor and do not allow the recovery routines to retry

ACR

Invoke alternate CP recovery (ACR) to take the failing processor offline.

- The system chooses the appropriate action without requiring any decision or action. If an action taken in response to an occurrence of an excessive spin loop does not resolve the condition, the system takes the next action when the next excessive spin loop time out occurs. The default order in which the system takes the actions is SPIN, ABEND, TERM, and ACR.
- An installation can change the order of the actions, except the first one, that the system takes.
- For hardware-related errors that formerly caused message IEA490A, the system immediately initiates ACR processing without working through the sequence of actions and without requiring any intervention.
- There is a default spin loop time-out interval. You can change this interval through the combination of a parameter in the EXSPATxx (excessive spin condition actions) parmlib member and entering the SET command.
- To avoid unnecessary recovery actions, system functions that can validly exceed the interval are exempt from excessive spin-loop processing. If they exceed the time out interval, the system functions do cause an excessive spin loop record to be written to the log recording data set (LOGREC) data set.
- The installation can still control excessive spin loop recovery through operator actions.

See EXSPATxx (excessive spin condition actions) in *z/OS MVS Initialization and Tuning Reference*.

Steps for diagnosing a loop

The following steps guide you through diagnosing a loop:

1. [“Gathering additional data for a loop” on page 227](#)

2. [“Analyzing the dump to determine the type of loop”](#) on page 228
3. [“Diagnosing a disabled loop”](#) on page 229
4. [“Diagnosing an enabled loop”](#) on page 230
5. [“Diagnosing an excessive spin \(spin loop\)”](#) on page 232
6. [“Analyzing a logrec error record ”](#) on page 235
7. [“Searching the problem reporting databases”](#) on page 235

Use the following flowchart to guide diagnosis of a loop:

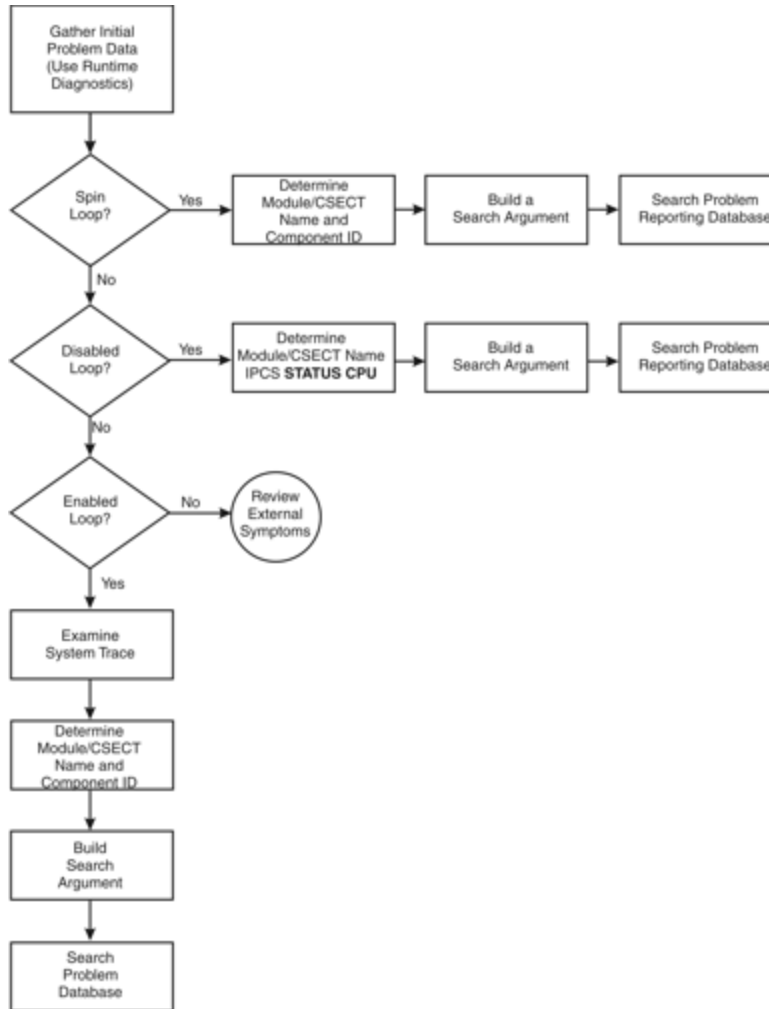


Figure 57. Flowchart for diagnosis of a loop

Gathering additional data for a loop

By gathering the correct data you can determine what recovery actions are necessary.

Steps for gathering loop data

Gather the following types of data:

1. The description of the external symptoms, including any software or hardware changes.
2. What recovery actions were attempted? Was a MODIFY, CANCEL or FORCE command entered?

Loop analysis

3. Request an SVC dump using the system DUMP command for the job or jobs involved in the loop. If this is a subsystem or system address space, define the dump options using the applicable dump parmlib member, IEADMCxx (DUMP command parmlib) documented in [z/OS MVS Initialization and Tuning Reference](#).

Otherwise, specify the following:

```
DUMP COMM=(dumptime) Rxx,ASID=(1,xx),SDATA=(GRSQ,SQA,CSA,RGN,TRT,COUPLE,XESDATA,NUC),END
```

Where xx is the ASID associated with the looping job. You can also specify it with JOBNAME=(xyz). If an SVC dump is not possible or the system is hung, request a stand-alone dump.

4. If the DUMP command cannot be entered because the system is hung, or does not complete, request a stand-alone dump of the system. See the topic on [Stand-alone dump](#) in [z/OS MVS Diagnosis: Tools and Service Aids](#).
5. If this is an enabled loop and there is a SYSMDUMP DD coded in the JCL, you can enter a CANCEL command for the looping job with the DUMP option.

Related information

- See [z/OS MVS Diagnosis: Tools and Service Aids](#) for information about stand-alone dumps, SVC, SYSMDUMP, and stand-alone dumps.
- See [z/OS MVS JCL Reference](#) for the SYSMDUMP DD statement.
- See [z/OS MVS System Commands](#) for information about the DUMP and CANCEL commands.

Analyzing the dump to determine the type of loop

You can determine what type of loop you have by analyzing the dump.

Step for analyzing the dump for loop type

Before you begin: You need access to IPCS.

Analyze any dump for the type of loop.

Format the dump with an IPCS STATUS CPU subcommand. Under the heading CPU(X'nn') STATUS, look for the following:

```
DISABLED FOR ccccccc
```

- System processing was disabled for one or more types of interrupts for the module running at the time of the dump. The system can be disabled for program event recording (PER), I/O, external interrupts (EXT), and machine checks (MCH).

The type is:

- Disabled loop: If the system was disabled for I/O or EXT or both.
- Enabled loop: If the system was not disabled for I/O or EXT.

In this example, the statement DISABLED FOR PER indicates an ENABLED loop.

```
Example: STATUS CPU Subcommand Output
CPU STATUS:
Warnings regarding STRUCTURE(ASCB) at ASID(X'0001') 00FD5F00:
Located via STRUCTURE(ASVT) at ASID(X'0001') 00F336D0
  Storage not in dump

PSW=070C2000 8AC2CB2A (RUNNING IN PRIMARY, KEY 0, AMODE 31, DAT ON)
  DISABLED FOR PER
  ASID(X'0001') 0AC2CB2A. AREA(PRIVATEX)+02CB2A IN EXTENDED PRIVATE
  ASCB76 at F52080, JOB(CATALOG), for the home ASID
```

Related information:

See [z/OS MVS IPCS Commands](#) for the STATUS subcommand.

Diagnosing a disabled loop

A disabled loop is not visible in the system trace output because disabled routines do not take interrupts. Normally, a disabled loop results in a spin loop in a multiprocessor environment. When analyzing a stand-alone dump for a disabled loop, use the stored status data to determine the module involved in the loop. Disabled loops often result in a system outage if they persist. Therefore, you usually be working with a stand-alone dump.

Steps for diagnosing a disabled loop

1. In an IPCS session, enter the IPCS STATUS CPU command to examine status of each CP. For example:

```
CPU(X'02') STATUS:
PSW=04042000 80000000 00000000 011E8592
(Running in PRIMARY, key 0, AMODE 31, DAT ON)
Disabled for PER I/O EXT
NOCPU ASID(X'0001') 011E8592. IEANUC01.IEAVELKX+073A IN READ ONLY NUCLEUS
ASCB1 at FD3400, JOB(*MASTER*), for the home ASID
ASXB1 at FD3598 and a local SRB for the home ASID
HOME ASID: 0001 PRIMARY ASID: 0001 SECONDARY ASID: 0001

CLTE: 01F76020
+0000 BLSL..... 00000000 XDS..... 00000000 XRES..... 00000000 XQ..... 00FD3018 ESET..... 00FD3028 IXSC..... 00000000
+0018 IXSH..... 00FD3040 IXDS..... 00000000 IXLL..... 00000000 ULUT..... 00FD3030 IXRE..... 00000000 WLMR..... 00000000
+0030 WLMQ..... 00FD3050 REGS..... 00000000 CNTX..... 00FD3060 SSD..... 00000000
HOLDING LOCK(S): CPU
CURRENT FRR STACK IS: NORMAL
CPU(X'04') STATUS:
PSW=04040000 80000000 00000000 017D5126
(Running in PRIMARY, key 0, AMODE 31, DAT ON)
Disabled for PER I/O EXT
NOCPU ASID(X'0001') 017D5126. IEANUC01.IAXRC+034E IN READ ONLY NUCLEUS
ASCB1 at FD3400, JOB(*MASTER*), for the home ASID
ASXB1 at FD3598 for the home ASID. No block is dispatched
HOME ASID: 0001 PRIMARY ASID: 0003 SECONDARY ASID: 0367

CLTE: 01B19000
+0000 BLSL..... 00000000 XDS..... 00000000 XRES..... 00000000 XQ..... 00FD3018 ESET..... 00FD3028 IXSC..... 00000000
+0018 IXSH..... 00FD3040 IXDS..... 00000000 IXLL..... 00000000 ULUT..... 00FD3030 IXRE..... 00000000 WLMR..... 00000000
+0030 WLMQ..... 00FD3050 REGS..... 00000000 CNTX..... 00FD3060 SSD..... 00000000
HOLDING LOCK(S): CPU RSM RSMCM RSMST RSMAD
CURRENT FRR STACK IS: NORMAL
```

In the preceding example, the PSW for CP 2 indicates that it is executing disabled in CSECT **IEAVELKX**, the lock manager. The status for CP 4 indicates several locks are held, therefore, this is a normal spin in **IEAVELKX**. The loop that is causing the problem is in IAXRC.

2. Use the name of the CSECT executing on CP 4, IAXRC and the symptom *LOOP* to check for a known problem.
3. From the module prefix, identify the component, subsystem, or product, if provided by IBM.

Use the module name to query the SMP/E zone for a module entry with that module name. If the search does not find a match, the module is not an IBM module. If the search indicates a match, use the FMID to positively identify the product.

4. Continue diagnosis as follows:
 - a. If all the addresses are in components of z/OS, continue with [“Searching the problem reporting databases” on page 235](#).
 - b. If all the addresses are in an IBM subsystem or product, continue diagnosis with the diagnosis publication for the subsystem or product. See Chapter 26, [“Diagnosis information for z/OS base elements and features,” on page 333](#) for the correct publication.
 - c. If all the addresses are in components of z/OS and in an IBM subsystem or product, continue with [“Searching the problem reporting databases” on page 235](#) and with the diagnosis book for the subsystem or product. See Chapter 26, [“Diagnosis information for z/OS base elements and features,” on page 333](#) for the correct publication.

- d. If any of the addresses are in an installation-provided program, including an installation exit routine, continue diagnosis with that program, using the dump.

If some addresses are in the program or routine and some in system modules, the loop is probably in the program or routine and includes one or more requests for system services.

Diagnosing an enabled loop

Enabled loops are often quite large and can include several distinct operations, such as I/O, SVCs, and module linkages. Because the loop is enabled, it is interrupted, preempted, and resumed many times. This makes the loop pattern difficult to recognize. “Steps for diagnosing an enabled loop” on page 230 can help make that identification easier.

For an enabled loop you should have either an SVC dump or a stand-alone dump.

Steps for diagnosing an enabled loop

Verify the system is in an enabled loop by examining the activity in the system trace table.

1. Enter the IPCS SYSTRACE ALL TIME (LOCAL) command. The following is an example of the system trace table output:

```

----- SYSTEM TRACE TABLE -----
-----
PR ASID WU-ADDR- IDENT CD/D PSW----- ADDRESS- UNIQUE-1 UNIQUE-2 UNIQUE-3 PSACLHS- PSALOCAL PASD SASD TIMESTAMP-LOCAL
UNIQUE-4 UNIQUE-5 UNIQUE-6 PSACLHSE DATE-05/08/2006
05-004C 00000000 EXT 1005 077C6000 951AFBD0 00001005 00000000 00000000 0301 004C 11:00:24.184339 00000000
05-004C 00000000 EXT 1005 077C6000 951AFBB2 00001005 00000000 00000000 0301 004C 11:00:24.184681 00000000
05-004C 031F6478 SSRB 077C6000 951AFBB2 00000000 04979298 00 00000000 0301 004C 11:00:24.184683 00000000
05-004C 00000000 EXT 1005 077C6000 951AFBA6 00001005 00000000 00000000 0301 004C 11:00:24.185024 00000000
05-004C 00000000 EXT 1005 077C6000 951AFBD8 00001005 00000000 00000000 0301 004C 11:00:24.185366 00000000
05-004C 00000000 EXT 1005 077C6000 951AFBD0 00001005 00000000 00000000 0301 004C 11:00:24.185708 00000000
05-004C 031F6478 SSRB 077C6000 951AFBD0 00000000 04979298 00 00000000 0301 004C 11:00:24.185710 00000000
05-004C 00000000 EXT 1005 077C6000 951AFBC0 00001005 00000000 00000000 0301 004C 11:00:24.186051 00000000
05-004C 00000000 EXT 1005 077C6000 951AFBD0 00001005 00000000 00000000 0301 004C 11:00:24.186393 00000000
05-004C 00000000 EXT 1005 077C6000 951AFBAC 00001005 00000000 00000000 0301 004C 11:00:24.186734 00000000
05-004C 031F6478 SSRB 077C6000 951AFBAC 00000000 04979298 00 00000000 0301 004C 11:00:24.186737 00000000
05-004C 00000000 EXT 1005 077C6000 951AFBD0 00001005 00000000 00000000 0301 004C 11:00:24.187082 00000000
05-004C 00000000 EXT 1005 077C6000 951AFBAC 00001005 00000000 00000000 0301 004C 11:00:24.187423 00000000
05-004C 00000000 EXT 1005 077C6000 951AFBB2 00001005 00000000 00000000 0301 004C 11:00:24.187765 00000000
05-004C 031F6478 SSRB 077C6000 951AFBB2 00000000 04979298 00 00000000 0301 004C 11:00:24.187767 00000000
05-004C 00000000 EXT 1005 077C6000 951AFBD4 00001005 00000000 00000000 0301 004C 11:00:24.188108

```

Figure 58. System trace table entry

The columns represent:

- The PR, ASID, and TCB-ADDR columns tell you who is executing
 - The IDENT CD/D columns tell you the type of event
 - The PSW column tells you where
 - The UNIQUE columns tell you how
 - The PASD tells you which ASID is being addressed
 - The TIMESTAMP tells you when.
2. To diagnose an enabled loop, look for patterns of repetitive entries in system trace. Typically, the bounds of an enabled loop are identified by the PSW addresses in EXT, CLKC, and IO events. In the preceding example, the bounds of the loop appear to go from PSW address 151AFBA6 to 151AFBD0.
 3. Enter the IPCS WHERE command on these PSW addresses to determine which module or CSECT the loop is in or browse storage at that location. For example, under IPCS Option 1:

```

Command ==>

CURRENT DEFAULTS:
Source ==> DSNAME('H44IPCS.PMR59356.B019.DUMP')
Address space ==> ASID(X'0001')

OVERRIDE DEFAULTS:
Source ==> DSNAME('H44IPCS.PMR59356.B019.DUMP')

```



```
Address space ==> ASID(X'301')
```

```
Password ==>
POINTER:
Address ==> 151AFBA6
Remark ==> ASID(X'0301')
```

Browse backward (PF 7) looking for an module or CSECT identifier.

```
Command ==>
151AEEA0 1606B1F8 00000000 00000000 00000000 | ...8..... |
151AEEB0 00000000 04000005 00000000 00000000 | ..... |
151AEEC0 :151AEECF. LENGTH(X'10')--All bytes contain X'00'
151AEEED0 00000000 80000000 00000000 00000000 | ..... |
151AEEEE0 00000000 00000000 03002100 30004001 | ..... |
151AEEF0 :151AEEF0F. LENGTH(X'20')--All bytes contain X'00'
151AEEF10 40000000 00000000 00000000 00000000 | ..... |
151AEEF20 :151AEEFFF. LENGTH(X'E0')--All bytes contain X'00'
151AF000 90ECDD84 A7C50046 151B0688 C3D8D4D6 | ...xE....hCQMO |
151AF010 D9C5C4C3 40404040 4040F0F3 61F0F261 | REDC 03/02/ |
151AF020 F0F6F1F9 4BF4F740 40404040 40404040 | 0619.47 |
151AF030 40404040 4040D6E2 40404040 F2F2F040 | OS 220 |
151AF040 D7D2F2F0 F6F0F040 F5F6F9F7 60C9F0F3 | PK20600 5697-I03 |
151AF050 404D835D 40D99683 9285A340 E29686A3 | ..... |
151AF060 A6819985 6B40C995 834B40F1 F9F9F96B | ..... |
151AF070 40F2F0F0 F640C193 9340D989 8788A3A2 | ..... |
151AF080 40D985A2 8599A585 844B4000 00000000 | ..... |
```

Browse backward (PF 7) looking for an module or CSECT identifier.

4. Obtain the PSW addresses from the system trace entries involved in the loop. For ANALYZE output and SUMMARY KEYFIELD CURRENT output, use the TCB address.

If using a PSW address, ignore the leftmost bit of the leftmost digit. The leftmost bit of the leftmost digit denotes addressing mode and is not part of the address.

5. Do one of the following actions, for each address in the loop:
 - If analyzing the dump interactively, use the address in a WHERE subcommand to obtain the name of the load module.
 - If analyzing printed output, find the address:
 - In dump output from the LIST or VERBEXIT SUMDUMP subcommand. Look for the CSECT name eye-catcher. IBM module names are generally followed by an assembly date and a product identifier or PTF level, such as HBB7720 or UY01234; most eye-catchers are at the beginning of a module, but some are at the end.
 - In a module listed for the LPAMAP or VERBEXIT NUCMAP subcommand. LPAMAP will list load modules. Use AMBLIST to obtain the offsets of CSECTS within those load modules. NUCMAP lists CSECTS with offsets, but can only be used for modules within the nucleus.

In the following example, STATUS CPU output and WHERE subcommand, the PSW identifies the address as X'13206AA'.

```
CPU STATUS:
```

```
PSW=040C2000 813206AA (RUNNING IN PRIMARY, KEY 0, AMODE 31, DAT ON)
DISABLED FOR PER I/O EXT
ASID(X'0006') 013206AA. IEANUC01.IXLM2SP+07AA IN READ ONLY NUCLEUS
ASCB1 at FD1780, JOB(*MASTER*), for the home ASID
ASXB1 at FD1A30 for the home ASID. No block is dispatched
HOME ASID: 0001 PRIMARY ASID: 0006 SECONDARY ASID: 0006
```

Using the WHERE subcommand, the load module name is IXLM2SP plus an offset of 7AA.

```
ASID(X'0006') 013206AA. IEANUC01.IXLM2SP+07AA IN READ ONLY NUCLEUS
```

6. From the module prefix, identify the component, subsystem, or product, if provided by IBM.

Loop analysis

Use the module name to query the SMP/E zone for a module entry with that module name. If the search does not find a match, the module is not an IBM module. If the search indicates a match, use the FMID to positively identify the product.

7. Continue diagnosis as follows:

- If all the addresses are in components of z/OS, continue with [“Searching the problem reporting databases” on page 235](#).
- If all the addresses are in an IBM subsystem or product, continue diagnosis with the diagnosis book for the subsystem or product.
- If all the addresses are in components of z/OS and in an IBM subsystem or product, continue with [“Searching the problem reporting databases” on page 235](#) and with the diagnosis book for the subsystem or product.
- If any of the addresses are in an installation-provided program, including an installation exit routine, continue diagnosis with that program, using the dump.

If some addresses are in the program or routine and some in system modules, the loop is probably in the program or routine and includes one or more requests for system services.

Related information:

- See [Identifying modules, components, and products in z/OS MVS Diagnosis: Reference](#).
- See [z/OS MVS IPCS Commands](#) for the IPCS subcommands.
- See [z/OS MVS Diagnosis: Tools and Service Aids](#)
- See [z/OS MVS Diagnosis: Reference](#) for the SVC summary.
- For the format and contents of the EED and TCB, see [z/OS MVS Data Areas](#) in the [z/OS Internet library](#) (www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary).

Diagnosing an excessive spin (spin loop)

To avoid unnecessary recovery actions, system functions that can validly exceed the interval are exempt from excessive spin-loop processing, so that they will not cause any recovery actions. If they exceed the time-out interval, these system functions do cause an excessive spin loop record to be written to the logrec data set.

You can examine the in-storage logrec buffer for entries that recovery routines have made, but which were not written to the logrec data set because of a system problem. Very often it is these records that are the key to the problem solution. See the topic on [Obtaining information from the logrec data set in z/OS MVS Diagnosis: Tools and Service Aids](#).

When the system writes a dump, the dump includes the records in the logrec buffer in storage; the buffer records have been either written to the logrec data set or are queued to be written to the logrec data set. When you begin to diagnose a dump for a system problem, you can use IPCS to view the system records in the logrec recording control buffer.

Before you begin You need to have IPCS installed and the following information:

- An EREP report of SYS1.LOGREC
- Any SVC dumps, stand-alone dumps showing the ABEND071, or SYSMDUMP dump
- SYSLOG and OPERLOG. Look for spin loop messages like IEE331A or IEE178I. Note that IEE331A is rare; IEE178I is more likely to occur. For example:

```
00 IEE331A PROCESSOR (0) IS IN AN EXCESSIVE
DISABLED SPIN LOOP WAITING FOR CPU IN
STOPPED STATE
REPLY U OR SPIN TO CONTINUE SPIN
REPLY ABEND TO TERMINATE WORK ON
PROCESSOR (1) WITH RETRY,
REPLY TERM TO TERMINATE WORK ON
PROCESSOR (1) WITHOUT RETRY,
OR STOP PROCESSOR (1) AND
```

```
REPLY ACR
(AFTER STOPPING THE PROCESSOR, DO NOT START IT)
```

```
IEE178I AUTOMATIC RECOVERY IS IN PROGRESS NO OPERATOR ACTION IS REQUIRED
PROCESSOR (00) DETECTED AN EXCESSIVE DISABLED SPIN LOOP
WAITING FOR LOCK RELEASE FROM PROCESSOR (03).
AUTOMATIC RECOVERY ACTION IS SPIN
```

Steps for diagnosing an excessive spin

Use the following steps:

1. In an EREP report of SYS1.LOGREC or an IPCS VERBX LOGDATA report in a dump, enter FIND S0071 (X'071'). You might find multiple records for S0071 with different reason codes.
2. The following example is for an ABEND071 RSN10, which is the first record written to logrec for a spin condition:

Example: VERBEXIT LOGDATA Output

SEARCH ARGUMENT ABSTRACT

PIDS/5752SC1CM RIDS/IEANUC01#L RIDS/IEAVTEXS AB/S0071 PRCS/00000010
REGS/0E68E REGS/0C7AA RIDS/IEAVTEXS#R

SYMPTOM	DESCRIPTION
-----	-----
PIDS/5752SC1CM	PROGRAM ID: 5752SC1CM
RIDS/IEANUC01#L	LOAD MODULE NAME: IEANUC01
RIDS/IEAVTEXS	CSECT NAME: IEAVTEXS
AB/S0071	SYSTEM ABEND CODE: 0071
PRCS/00000010	ABEND REASON CODE: 00000010
REGS/0E68E	REGISTER/PSW DIFFERENCE FOR R0E: 68E
REGS/0C7AA	REGISTER/PSW DIFFERENCE FOR R0C: 7AA
RIDS/IEAVTEXS#R	RECOVERY ROUTINE CSECT NAME: IEAVTEXS

Look for a SOFTWARE RECORD for abend X'071' with the following reason codes:

- Reason code X'10': Recovery action of SPIN was taken, and a logrec entry was written to provide information about the CSECT/module that was detected as looping. No attempt was made to abend the unit of work.
- Reason code X'20': Recovery action of ABEND was taken, targeting the looping program with a retryable ABEND071.
- Reason code X'30': A recovery action of TERM was taken, targeting the looping program with a non-retryable ABEND071.

In a logrec record for an abend X'071', look in the SEARCH ARGUMENT ABSTRACT for the symptom RIDS/IEAVTEXS. If found, look under the heading VARIABLE RECORDING AREA (SDWAVRA) for the following problem data:

- EX SPIN RECORD in the EBCDIC text
- An array of up to 16 pointers: addresses of the functional recovery routines (FRR) on the interrupted FRR stack
- A binary number: index into the FRR stack of the current FRR
- Array of processors causing the spin from SVTSPCP
- Spin loop timeout interval
- Spin loop recovery actions
- Control registers

In the following output, the SDWAVRA indicates an excessive spin record and contains an array of pointers to the FRR stack.

Example: VERBEXIT LOGDATA Output -- Variable Recording Area

Loop analysis

VARIABLE RECORDING AREA (SDWAVRA)					
+000	KEY: 39	LENGTH: 0E			
+002	C5E740E2	D7C9D540	D9C5C3D6	D9C4	EX SPIN RECORD
+010	KEY: 37	LENGTH: 04			
+012	C6D9D9E2				FRRS
+016	KEY: 38	LENGTH: 40			
+018	811F43F8	811D74ED	81319571	81086D05	A..8A...A.N.A..
+028	811D74ED	00000000	00000000	00000000	A.....
+038	00000000	00000000	00000000	00000000
+048	00000000	00000000	00000000	00000000

- a. Enter the IPCS WHERE command on the PSW address at the time of error from the logrec entries found for an S0071. This should identify the CSECT/module that is excessively looping.
 - b. Search for a known problem using the symptoms: ABEND071 and the module or CSECT name identified. If IPCS WHERE fails to identify a CSECT/module name, browse storage preceding the psw address looking for a module eye catcher using IPCS Option 1.
3. Identify the modules containing each instruction in the loop, as follows: For automatic spin loop recovery, if the dump is an SVC dump: Use a STATUS WORKSHEET subcommand to obtain the dump title. Obtain the component name, component identifier, or module name from the title.

In the following example, STATUS WORKSHEET output, the dump title is for the XES component.

```
MVS Diagnostic Worksheet
Dump Title:  COMPON=IXL,COMPID=5752SCIXL,ISSUER=IXLM1REC,MODULE=IXLM2SP
             ,ABEND=S0071,REASON=00000030

CPU Model 9021 Version A6 Serial no. 300359 Address 03
Date: 03/30/93      Time: 10:32:38 Local

Original dump dataset: SYS1.DUMP32
```

4. From the module prefix, identify the component, subsystem, or product, if provided by IBM.
- Use the module name to query the SMP/E zone for a module entry with that module name. If the search does not find a match, the module is not an IBM module. If the search indicates a match, use the FMID to positively identify the product.
5. Continue diagnosis as follows:
- If all the addresses are in components of z/OS, continue with [“Searching the problem reporting databases”](#) on page 235.
 - If all the addresses are in an IBM subsystem or product, continue diagnosis with the diagnosis book for the subsystem or product.
 - If all the addresses are in components of z/OS and in an IBM subsystem or product, continue with [“Searching the problem reporting databases”](#) on page 235 and with the diagnosis book for the subsystem or product.
 - If any of the addresses are in an installation-provided program, including an installation exit routine, continue diagnosis with that program, using the dump.
- If some addresses are in the program or routine and some in system modules, the loop is probably in the program or routine and includes one or more requests for system services.

Related information:

- See [z/OS MVS IPCS Commands](#) for the IPCS subcommands.
- See [Identifying modules, components, and products in z/OS MVS Diagnosis: Reference](#).
- See [z/OS MVS Diagnosis: Tools and Service Aids](#) for information about analyzing an SVC dump for a problem in an installation-provided program.

Analyzing a logrec error record

Analyze a logrec error record for a disabled loop. The excessive-spin logrec error record can identify the module running on the processor causing the spin condition.

Steps for analyzing a logrec error record

Before you begin: You need access to the logrec error record.

Use the following steps to identify the module running on the processor causing the spin condition, as follows:

1. Locate the 16 FRR addresses from the stack that was current when the target processor was restarted. These addresses appear after the identification text EX SPIN RECORD at the start of the VRA.
2. Identify the current FRR on the stack from the INDEX=x value that follows the sixteen addresses. The value of x can be 0 through 16.

If x is 0, the stack contains no current FRRs. Otherwise x is an index indicating which of the 16 addresses points to the current FRR. For example, if x is 2, the second address points to the current FRR.

3. Use a storage map to identify the component that owns the FRR at this address.

For example: In the following output example, the current FRR is 81319571, which is the third FRR in the stack. This is the current FRR because INDEX=03.

VERBEXIT LOGDATA output — FRR Stack

+012	C6D9D9E2				FRRS	
+016	KEY: 38	LENGTH: 40				
+018	811F43F8	811D74ED	81319571	81086D05	A..8A...A.N.A..	
+028	811D74ED	00000000	00000000	00000000	A.....	
+038	00000000	00000000	00000000	00000000	
+048	00000000	00000000	00000000	00000000	
+058	KEY: 37	LENGTH: 06				
+05A	C9D5C4C5	E77E			INDEX=	
+060	KEY: 38	LENGTH: 01				
+062	03				.	

Searching the problem reporting databases

Search arguments are used to search problem reporting databases. If the problem you are diagnosing was already reported and the symptoms are in the database, the search produces a match. Searching is an iterative process; you might need to gather additional data and continue your search.

Steps for searching the problem reporting databases

Use the following steps to search the problem reporting databases and determine if the problem was previously reported:

1. Create a search argument using the symptoms applicable for the type of loop being diagnosed:
 - a. Disabled loop search argument - LOOP module/CSECT name
 - b. Enabled loop search argument - LOOP module/CSECT name
 - c. Spinloop search argument - ABEND071 module/CSECT name.

Tip: Use free-format search arguments. For more information, see [“Searching problem reporting databases”](#) on page 8.

2. If the search finds no match, remove some symptoms or add some symptoms. Search again. Continue searching for matches by adding and removing symptoms.

Loop analysis

If the search finds that the problem was previously reported, request the fix.

3. If the search does not produce a match, continue with the next step. Use problem data from the preceding steps to create more symptoms; use these symptoms in later searches.
4. If you still cannot find the cause of the loop or if the problem is new, report the problem to IBM using the procedure in Chapter 27, “Reporting problems to IBM,” on page 335. Provide the following problem data as listed in [Chapter 28, “Problem diagnostic worksheet,”](#) on page 339.
 - Problem type: type of loop
 - Search argument
 - Dump, formatted by IPCS, online or printed
 - Range of loop
 - SMF records, if obtained
 - Accompanying messages: identifiers and texts
 - Hard-copy log, beginning 15 to 30 minutes before the problem, or master trace, if not wrapped between the problem and dump
 - All printed output and output data sets related to the problem
 - Data on any related problems
 - Module name and level
 - Name and level of the operating system(s) with a list of program temporary fixes (PTF) applied at the time of the problem and all installation modifications, exits, and products with other than Class A service
 - Other problem data developed while using the procedures in this document or other diagnosis books for the component, subsystem, or program

Related information

- See [“Searching problem reporting databases”](#) on page 8 for more information.
- See [z/OS MVS IPCS Commands](#) for the IPCS subcommands.

Chapter 14. Diagnosing an output problem

Overview of analyzing output problems

Most output problems occur during installation and testing of new functions or applications.

This chapter explains what information you will need to properly diagnose an output problem.

Symptoms of output problems: Your output is incorrect, incomplete, or missing, but messages indicate successful processing.

- **Incorrect output:** The processing produced all the expected output, but some output is incorrect. For example:
 - Some values in a report are wrong
 - The text in a message is incorrect
 - The return or reason code is not valid
 - The records in are not producing expected records out.
- **Incomplete output:** The processing did not produce all the expected output. For example, a column is missing from a report.
- **Missing output:** Some or all of the expected output is missing. For example, a report is missing.

Steps for diagnosing output problems

Before you begin: You need to have IPCS installed and have access to the following information:

- Job log
- Input and output data sets for the program
- The JCL for the program
- System log
- Logrec data set

Use the following flowchart to guide you through diagnosing an output problem:

Output problem analysis




Figure 59. Flowchart for output problem analysis

The following steps will guide you through diagnosing output problems:

1. [“Collecting problem data for an output problem” on page 238](#)
2. [“Analyzing data set allocation for an output problem” on page 239](#)
3. [“Analyzing the inputs and outputs” on page 239](#)
4. [“Analyzing installation exits for an output problem” on page 240](#)
5. [“Identifying the program or component” on page 241](#)
6. [“Searching the problem reporting databases for an output problem” on page 242](#)
7. [“Gathering additional data for output problems” on page 243](#)
 - [“Messages and logrec for output problems” on page 243](#)
 - [“Determine path for output problems” on page 243](#)
 - [“Teleprocessing for output problems” on page 243](#)
8. [“Reporting output problems to IBM” on page 244](#)

Collecting problem data for an output problem

To properly diagnose an output problem you need to collect as much information as possible.

Step for collecting problem data

- If the output is a data set, collect the following:
 - All input data sets for the program
 - Input macros, commands, and statements that are used to request output from the program
 - All output data sets produced by the program
 - The job log
- If the output is a message, return code, or reason code, collect the incorrect or incomplete message or code.

Analyzing data set allocation for an output problem

For missing data set output, analyze data set allocation.

Steps for analyzing data set allocation

1. Look for messages in the job log indicating that all data sets used by the program were properly allocated and unallocated. Look for:
 - The data set that should have contained the output
 - The data set that contained the input used by the program to create the output
 - Indications that the output was sent to another data set or different system
2. If problems are found, correct the JCL.
3. If the problem is not found, continue diagnosis with [“Analyzing the inputs and outputs” on page 239](#).

Related information:

- See [z/OS MVS JCL Reference](#) for JCL coding.
- For explanations of the messages, see:
 - [z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#)
 - [z/OS MVS System Messages, Vol 2 \(ARC-ASA\)](#)
 - [z/OS MVS System Messages, Vol 3 \(ASB-BPX\)](#)
 - [z/OS MVS System Messages, Vol 4 \(CBD-DMO\)](#)
 - [z/OS MVS System Messages, Vol 5 \(EDG-GLZ\)](#)
 - [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#)
 - [z/OS MVS System Messages, Vol 7 \(IEB-IEE\)](#)
 - [z/OS MVS System Messages, Vol 8 \(IEF-IGD\)](#)
 - [z/OS MVS System Messages, Vol 9 \(IGF-IWM\)](#)
 - [z/OS MVS System Messages, Vol 10 \(IXC-IZP\)](#)

Analyzing the inputs and outputs

For incorrect or incomplete data set output, analyze the inputs and outputs.

Steps for analyzing the inputs and outputs

Perform the following steps:

1. Compare the input and the output. For example, if a device model number is wrong in an IOS report, compare it to the model number specified on the in the IODF.
2. Compare the output received to examples of the output shown in the user's guide for the request.

Output problem analysis

3. Check the call, command, macro, or statement used to request the output. Make sure that all fields contain desired values.
4. Match your findings to the following:
 - For missing data set output, check the macro, command, or statement used to request the output. Make sure that the missing output was supposed to be received.
 - For an incorrect or incomplete message, match the message received to the message in the book with its explanation.
 - For incorrect return and reason codes, match the codes to the expected codes. Make sure that the code received and the code in the book are the same type: both hexadecimal or both decimal.

For the message text or a return or reason code in a message, see:

- [*z/OS MVS System Messages, Vol 1 \(ABA-AOM\)*](#)
- [*z/OS MVS System Messages, Vol 2 \(ARC-ASA\)*](#)
- [*z/OS MVS System Messages, Vol 3 \(ASB-BPX\)*](#)
- [*z/OS MVS System Messages, Vol 4 \(CBD-DMO\)*](#)
- [*z/OS MVS System Messages, Vol 5 \(EDG-GLZ\)*](#)
- [*z/OS MVS System Messages, Vol 6 \(GOS-IEA\)*](#)
- [*z/OS MVS System Messages, Vol 7 \(IEB-IEE\)*](#)
- [*z/OS MVS System Messages, Vol 8 \(IEF-IGD\)*](#)
- [*z/OS MVS System Messages, Vol 9 \(IGF-IWM\)*](#)
- [*z/OS MVS System Messages, Vol 10 \(IXC-IZP\)*](#)
- [*z/OS MVS Dump Output Messages*](#)
- The message book for a subsystem or program.

For a return or reason code with an abend code, see

- [*z/OS MVS System Codes*](#)

For a return or reason code for a macro, see:

- [*z/OS MVS Programming: Authorized Assembler Services Reference ALE-DYN*](#)
- [*z/OS MVS Programming: Authorized Assembler Services Reference EDT-IXG*](#)
- [*z/OS MVS Programming: Authorized Assembler Services Reference LLA-SDU*](#)
- [*z/OS MVS Programming: Authorized Assembler Services Reference SET-WTO*](#)
- [*z/OS MVS Programming: Authorized Assembler Services Guide*](#)
- [*z/OS MVS Programming: Assembler Services Reference ABE-HSP*](#)
- [*z/OS MVS Programming: Assembler Services Reference IAR-XCT*](#)
- [*z/OS MVS Programming: Assembler Services Guide.*](#)

Analyzing installation exits for an output problem

Analyze all installation exit routines used in obtaining the output.

Steps for analyzing installation exits

1. Check for problems in the logic of each exit routine. Many installation exits are invoked using the dynamic exits facility, which allows an updated exit to be refreshed into the system.
2. If no logic problems are found, remove the options that cause each exit routine to be invoked.
3. Rerun the program.

4. If this action stops the problem, the problem is in the exit routine that was eliminated. Continue diagnosis with that routine.

Related information:

See *z/OS MVS Installation Exits* for more information on coding installation exit routines.

Identifying the program or component

Identifying the component that is involved with the output enables you to determine the source of the problem.

Steps for identifying the program or component

Identify the program or component involved with the output from one of the following steps:

1. For output from a batch job, obtain the program name from the PGM parameter on the JCL EXEC statement. In the following example the program name is obtained from the JCL EXEC statement:

In the following example, the name of a program involved with the output is IKJEFT01. It can be found on the PGM parameter of the EXEC statement:

```
//*
//*=====
//* Batch TSO job (PGM=IKJEFT01)
//*=====
//IKJEFT01 EXEC PGM=IKJEFT01,REGION=4096K,DYNAMNBR=50
//*
```

2. Identify the program or component involved with the output from one of the following:
 - For output from interactive work, use the command being processed to identify the program.
 - For an error message, use the message prefix to identify the program or component or look for the component listed in the message explanation. Look at the preface of any of the z/OS MVS system messages books to find the element or component that the message is associated with and the book where you can find the explanation of the message:

For example, if you have an incorrect message with message number AHL002A, you can tell from the preface of *z/OS MVS System Messages, Vol 1 (ABA-AOM)* that the message was issued by GTF, and the explanation can be found in the same book:

Table 23. Example of using the message prefix to identify the component

Prefix	Component	Book Title	Order Number
AHL	Generalized trace facility (GTF)	– <i>z/OS MVS System Messages, Vol 1 (ABA-AOM)</i> , – <i>z/OS MVS Dump Output Messages</i>	– SA38-0668 – SA23-1378

3. For a return or reason code accompanying an abend, see the component listed in the explanation of the code. The following example identifies the component issuing an abend:

00D

Explanation: An error occurred during processing of a CTRACE or CTRACEWR macro. Register 15 contains xnnnnxx where nnnn is a reason code that further describes the error.

Reason Code (hex)
Explanation

Output problem analysis

0001

For the CTRACE macro, the parameter list version number is incorrect.

0002

For the CTRACE macro, the component name either does not begin with an alphabetic or national character, or it contains one or more characters that are not alphanumeric or national characters.

Source:

Component trace

4. For a component of z/OS, continue with [“Searching the problem reporting databases for an output problem”](#) on page 242.
5. For an IBM subsystem or product, continue diagnosis with the subsystem or product.
6. For an installation-provided program, including an installation exit routine, continue diagnosis with that program.

Related information:

- See [z/OS MVS JCL Reference](#) for the EXEC statement.
- For message prefixes for IBM components, subsystems, and products, see the following books:
 - [z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#)
 - [z/OS MVS System Messages, Vol 2 \(ARC-ASA\)](#)
 - [z/OS MVS System Messages, Vol 3 \(ASB-BPX\)](#)
 - [z/OS MVS System Messages, Vol 4 \(CBD-DMO\)](#)
 - [z/OS MVS System Messages, Vol 5 \(EDG-GLZ\)](#)
 - [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#)
 - [z/OS MVS System Messages, Vol 7 \(IEB-IEE\)](#)
 - [z/OS MVS System Messages, Vol 8 \(IEF-IGD\)](#)
 - [z/OS MVS System Messages, Vol 9 \(IGF-IWM\)](#)
 - [z/OS MVS System Messages, Vol 10 \(IXC-IZP\)](#)
- See [z/OS MVS System Codes](#) for explanations of the abend codes.
- See [z/OS MVS Diagnosis: Reference](#) to relate an IBM component, subsystem or product to a program identifier or module prefix.

Searching the problem reporting databases for an output problem

Use free-format search arguments to search the problem reporting databases. If the problem you are diagnosing was already reported and the symptoms are in the database, the search produces a match. Searching is an iterative process; you might need to gather additional data and continue your search.

Step for searching the problem reporting database

Use the following steps to search a problem reporting database to determine if this is a known problem.

1. Use a free-format search argument developed from the symptoms. For more information, see [“Searching problem reporting databases”](#) on page 8.
2. If the search finds no match, remove some symptoms or change the symptoms. Search again. Continue searching for matches by adding, removing, and changing symptoms.
3. If the search finds that the problem was previously reported, request the problem fix. If not, continue with the next step. Use problem data from following steps to create more symptoms; use these symptoms in later searches.

Related information:

See [“Searching problem reporting databases”](#) on page 8 for more information on developing a search argument.

Gathering additional data for output problems

Gathering additional data will increase your chances of finding a match in the problem reporting databases. Use the procedures outlined in this section to gather additional data and continue searching the problem reporting databases.

Steps for gathering additional information for output problems

Use the following steps to collect additional information about:

- [“Messages and logrec for output problems”](#) on page 243
- [“Determine path for output problems”](#) on page 243
- [“Teleprocessing for output problems”](#) on page 243

Messages and logrec for output problems

1. Collect and analyze messages and logrec records about the problem. Look at any messages or software, symptom, and hardware records for logrec around the time of the problem.
2. Look in the following:
 - Job log
 - TSO/E user's ISPF transaction log or session manager log
 - System log (SYSLOG) for the console with master authority or the alternate console
 - Logrec data set, formatted by EREP.
3. Check for:
 - I/O errors that could affect the output
 - Operator interactions that could affect the output
 - Problems with the access method or function involved: For example, VSAM, BTAM, JES, or WTO.

Determine path for output problems

1. Analyze the path the data should take.
2. For the request being processed, determine the correct path for the data from input to output.
3. Determine each program and component involved.
4. Check for messages about problems in these programs and components.
5. If an installation-provided exit routine receives control during the processing, check the routine.
6. Look at the environment. Specifically, look for recent hardware and software changes to the system and to any applications. A change in one program can affect others; for example, a change to an application that updates a database affects all other users of the database.
7. If needed, recreate the problem, using a SLIP trap and traces to obtain the data needed to isolate the problem.

Teleprocessing for output problems

1. Analyze the path the data should take.
2. Determine how the data flows through the programs and components that process it and through the systems and hardware. Use this knowledge to recreate the problem, using traces to checkpoint the data at certain spots along the path.
3. Track the data from a point where it was correct to a point where the data stopped or became incorrect.

Output problem analysis

Related information:

- For formatting of logrec records, see [Recording logrec error records in z/OS MVS Diagnosis: Tools and Service Aids](#).
- See [z/OS MVS Diagnosis: Tools and Service Aids](#) for the logrec records.
- For explanations of the messages, see:
 - [z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#)
 - [z/OS MVS System Messages, Vol 2 \(ARC-ASA\)](#)
 - [z/OS MVS System Messages, Vol 3 \(ASB-BPX\)](#)
 - [z/OS MVS System Messages, Vol 4 \(CBD-DMO\)](#)
 - [z/OS MVS System Messages, Vol 5 \(EDG-GLZ\)](#)
 - [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#)
 - [z/OS MVS System Messages, Vol 7 \(IEB-IEE\)](#)
 - [z/OS MVS System Messages, Vol 8 \(IEF-IGD\)](#)
 - [z/OS MVS System Messages, Vol 9 \(IGF-IWM\)](#)
 - [z/OS MVS System Messages, Vol 10 \(IXC-IZP\)](#)
 - [z/OS MVS Dump Output Messages](#)
 - The message book for a subsystem or program
- See [z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures](#) and [z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT](#) to diagnose VTAM® problems.
- See the SLIP chapter in [z/OS MVS System Commands](#) for information on setting a SLIP trap.
- See [z/OS MVS Diagnosis: Tools and Service Aids](#) for requesting dumps and traces.
- See [z/OS MVS Installation Exits](#) for exit routines.

Reporting output problems to IBM

If you have completed the procedures listed in this chapter, cannot find a match in the problem reporting databases, and you believe the failure was caused by a defect in IBM code, call the IBM Support Center.

Provide the following problem data:

- Problem type: INCORROUT, that is, incorrect, incomplete, or missing output in a data set, message, return code, or reason code
- Search argument
- All input associated with the problem, including all data sets, commands, macros, and statements
- All output associated with the problem, including data sets, reports, and records
- JCL for all data sets involved
- Source code for any exit routine involved
- Accompanying messages: identifiers and texts
- Hard copy log, beginning 30 to 60 minutes before the problem, or master trace, if not wrapped between the problem and dump
- Name and level of the operating system(s) with a list of program temporary fixes (PTF) applied at the time of the problem and all installation modifications, exits, and products with other than Class A service

Related information:

- See [Chapter 27, “Reporting problems to IBM,”](#) on page 335 for more information on reporting a problem.

- For formatting of logrec records, see the topic on Recording logrec error records in *z/OS MVS Diagnosis: Tools and Service Aids*.

Chapter 15. Diagnosing a performance problem

Overview of a performance problem

Most performance problems appear as unacceptable response times or resource usage. For example, system processing is slow because a program is using an excessive amount of system resources.

Symptoms of a performance problem:

- Jobs take more time to execute than normal
- Poor or erratic response time exists
- Service level objectives are being exceeded
- Users complain about slow response time
- Unexpected changes occur in response times or resource utilizations
- Other indicators show stress
- Monitor III Workflow/Exceptions occur.
- System resource indicators occur (for example, paging rates, DASD response)
- Expected throughput on the system is not being attained
- CP utilization for a job or address space is higher than normal.

Steps for diagnosing a performance problem

Before you begin: You need access to following types of information:

- System messages
- SVC dump
- Job log
- TSO/E user's ISPF transaction log or session manager log
- System log
- Logrec data set
- z/OS Resource Measurement Facility (RMF), SMF, or other system monitoring programs, such as IBM OMEGAMON z/OS Management Console.

When diagnosing a performance problems, using Predictive Failure Analysis (PFA) and Runtime Diagnostics can help you eliminate whether the problem is a damaged job or system. See [Chapter 5, “Runtime Diagnostics,”](#) on page 35 and [Chapter 7, “Predictive Failure Analysis overview and installation,”](#) on page 73.

RMF also provides comprehensive guidance for diagnosing performance problems. [z/OS RMF User's Guide](#) offers a practical, task-oriented approach to analyzing performance issues using RMF.

Use the following flowchart to guide you through diagnosing a performance problem:

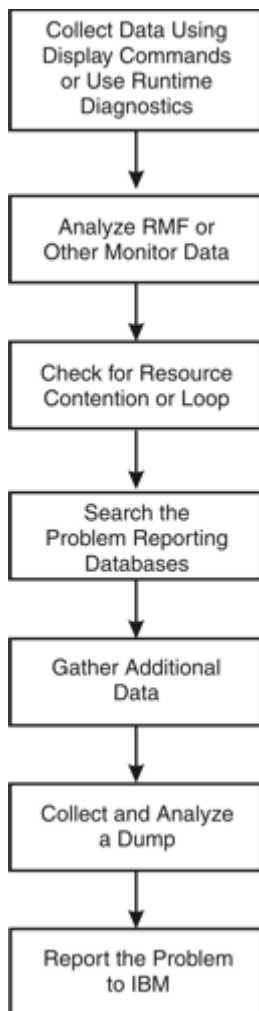


Figure 60. Flowchart for performance problem analysis

Use the following steps to guide you through diagnosing a performance problem:

1. [“Collecting data using commands” on page 248](#)
2. [“Checking for resource contention or loop” on page 251](#)
3. [“Searching the problem reporting database” on page 252](#)
4. [“Gathering additional data for performance problems” on page 253](#)
5. [“Analyzing a dump for performance problems” on page 253](#)
6. [“Reporting performance problems to IBM” on page 254](#)

Collecting data using commands

Collecting information with commands will give you a better understanding of the source of the problem. In a JES2 system, use JES2 commands to determine why JES2 is not able to schedule work. This topic is divided into two sections:

- [“Steps for collecting data using DISPLAY” on page 248](#)
- [“Steps for using JES2 commands to collect data” on page 250](#)

Steps for collecting data using DISPLAY

Use the following DISPLAY commands in the order listed to find the source of your problem:

1. Determine if the system is waiting for an operator action, for example, mounting of a volume. Other jobs might have to wait until the action is completed. In this case, the operator should perform the action. The following command displays outstanding messages requiring operator action:

```
DISPLAY R,LIST
```

In the following example of DISPLAY R,LIST output, there is one message, IEF434D, requiring immediate operator action. The system waits until a valid reply is entered. The operator should enter a valid reply to this message.

```
IEE112I 13.39.37 PENDING REQUESTS      FRAME LAST  F      E      SYS=SY1
RM=0     IM=0     CEM=2     EM=0     RU=0     IR=0     AMRF
ID:R/K   T  JOB  ID  MESSAGE TEXT
          3  C           *ILR005E PLPA PAGE DATA SET FULL, OVERFLOWING TO
          5  R           COMMON DATA SET
          *IEF434D CRITJOB - INVALID REPLY. REPLY 'HOLD'
          OR ' OHOLD'
```

2. Look for the generalized trace facility (GTF) in the started tasks. GTF may slow performance. A job step name of STARTING indicates that the system has not yet successfully completed initiation of the first step. If unsuccessful initiation of the job continues, diagnose this problem. The following command displays detailed information for active jobs and started tasks.

```
DISPLAY A,LIST
```

In the following example of DISPLAY A,LIST output, all of the steps have been successfully initiated. However, the performance problem could be due to GTF being active:

```
CNZ4105I 16.47.46 DISPLAY ACTIVITY      FRAME LAST  F      E      SYS=SY1
JOBS      M/S      TS  USERS      SYSAS      INITS      ACTIVE/MAX VTAM      OAS
00000     00005     00000     00016     00008     00000/00300     00000
LLA       LLA       LLA       NSW  S      VLF       VLF       VLF       NSW  S
JES2     JES2     IEFPROC   NSW  S      SDSF      SDSF      SDSF      NSW  S
GTF       GTF       IEFPROC   NSW  S
```

3. Look for the loss of a hardware component indicated by a message in the hard copy log. This loss might be causing jobs to wait. In this case, correct the hardware problem. The DISPLAY M command displays the hardware configuration.

```
DISPLAY M
```

4. Look for the name and status of current SLIP traps.

- a. Enter the DISPLAY SLIP command for a summary of SLIPs that are running on the system:

```
DISPLAY SLIP
DISPLAY SLIP=xxxx
```

Output using the DISPLAY SLIP command shows the name and status of the current SLIP traps. For more information see the [Figure 61](#) on page 249.

```
IEE735I 13.42.00 SLIP DISPLAY      FRAME LAST  F      E      SYS=SY1
ID  STATE  ID  STATE  ID  STATE  ID  STATE  ID  STATE
X013 ENABLED X028 ENABLED X0E7 ENABLED X0F3 ENABLED X13E ENABLED
X222 ENABLED X322 ENABLED X33E ENABLED X622 ENABLED X804 ENABLED
X806 ENABLED X80A ENABLED X9FB ENABLED XB37 ENABLED XD37 ENABLED
XE37 ENABLED
```

Figure 61. Output from the DISPLAY SLIP command

- b. Pick a SLIP ID that is ENABLED and enter DISPLAY SLIP=xxxx to check for enabled PER traps. In this case, disable the traps.

In [Figure 61](#) on page 249, all the SLIPs appear enabled. Entering DISPLAY SLIP=X013 yields the following enabled PER trap with an action of STRACE, which can slow performance:

```
IEE735I 09.14.03 SLIP DISPLAY      FRAME LAST  F      E      SYS=SY1
ID=X013,PER-SB,ENABLED(ACTIVE),ACTION=STRACE,SET BY TSO KLOGAN
```

Related information:

- See *z/OS MVS System Commands* for the SLIP and DISPLAY commands.
- See *z/OS MVS Diagnosis: Tools and Service Aids* for information about GTF.

Steps for using JES2 commands to collect data

Use the following procedures to collect problem data with JES2 commands:

1. Use the job entry subsystem display commands to find the status of jobs, queues, printer setups, requirements of SYSOUT data sets, and other problem data.

```
$D J1-9999
```

\$D J1-99 output displays the status of jobs. If the display shows that a range of jobs has been held, use the JES2 \$A J command to release specific jobs. Or, use the JES2 \$A A command to release all jobs in the system.

In the following \$D J1-99 output, all of the displayed jobs are being held:

```
$D J1-99
JOB00005 $HASP608 IEBGENER AWAITING HARDCOPY      Prio HELD 15 ANY
JOB00006 $HASP608 XEQN1   AWAITING XMITTER POK     Prio HELD 9  ANY
JOB00007 $HASP608 IEBGENER AWAITING HARDCOPY      Prio HELD 15 ANY
JOB00008 $HASP608 IEBGENER AWAITING EXECUTION A    Prio HELD 9  ANY
```

2. Use the following command to display the status of tasks started under JES2:

```
$D S1-9999
```

The following \$D S1-99 output shows tasks started under JES2:

```
$D S1-99
STC00001 $HASP608 SYSLOG   EXECUTING $           Prio 15 IBM2
STC00002 $HASP608 $MASCMM AWAITING HARDCOPY     Prio 15 ANY
STC00003 $HASP608 INIT     EXECUTING $           Prio 15 IBM2
INITASID=0015
STC00004 $HASP608 IRRDPTAB ON PRT1              Prio 1  IBM2
```

3. Use the following command to display the status of time-sharing users:

```
$D T1-9999
```

4. Use the following command to display the status of data set groups queued for output and the status of JES2-controlled local printers.

```
$D F
$D U, PRTS
```

If these displays show that no printers are set up with the needed forms, use the JES2 \$T PRTnnnn command to change a printer's setup to the needs of the output forms queue.

In the following \$D F and \$D U, PRTS output, there is only one item queued for the printers. All of the printers, however, are either drained or halted, which means that none of the printers are started. No printing can occur.

```
$D F
$HASP621 OUT R=LOCAL          F=STD      C=**** T=**** W= (NONE)
$HASP621 PRMODE=LINE        CLASS A=3

$D U, PRTS

$HASP603 PRT1      UNIT=0017,STATUS=HALTED,(STC00004 IRRDPTAB)
$HASP603 PRT2      UNIT=0002,STATUS=DRAINED
```

```
$HASP603 PRT3      UNIT=0017,STATUS=DRAINED
$HASP603 PRT4      UNIT=000E,STATUS=DRAINED
$HASP603 PRT5      UNIT=000F,STATUS=DRAINED
$HASP603 PRT6      UNIT=000F,STATUS=DRAINED
```

5. Use the following command to display the number of queued jobs.

```
$D Q
```

If the problem is that jobs are not running because they are held, you can use the JES2 \$A command with appropriate parameters to release the jobs.

Related information:

See [z/OS JES2 Commands](#) for the JES2 commands.

Checking for resource contention or loop

You can analyze output from RMF, SMF or another system monitoring program to look for resource contention and loops.

Steps for checking resource contention

Before you begin:

Perform the following steps to check for resource contention:

1. Use output from RMF, SMF or another system monitoring program to look for problems. Find someone in your installation who is familiar with the program and can interpret the output. The following lists some of the potential problems to look for:
 - A program using a lot of storage, whether it is real, virtual, auxiliary or extended storage.
 - Data set contention
 - ENQ contention
 - Tuning problems
 - System running over capacity
2. Identify the program, job or function that is involved in the performance problem. If using more resources than normal, gather RMF, trace output, or both from the time frame of when the slowdown occurred.

For a problem caused by resource contention, use dump output from [“Analyzing a dump for performance problems”](#) on page 253 or the following two choices:

- Use the ANALYZE output to identify the problem causing the contention:

In the following output, resource #0002, which is device 687, shows an intercept condition and is not running. There is one unit of work waiting for this device.

```

                                CONTENTION EXCEPTION REPORT
JOBNAME=IOS.      ASID=0001  UCBTAPE=00FC72C0
JOBNAME=IOS.      HOLDS THE FOLLOWING RESOURCE(S):
  RESOURCE #0002:  There are 0001 units of work waiting
                  for this resource
                  NAME=I/O Device 687 (TAPE    ) VOLSER=.....
                  DATA=(IOS) Active I/O with ASSIGN held.
                  (IOS) Device not ready.
                  (IOS) Intercept condition.

STATUS FOR THIS UNIT OF WORK:
  IRA10102I This address space is on the SRM IN queue.
```

Performance problem analysis

- For a problem caused by a batch program and identified through JES2 commands, obtain the program name from the PGM parameter on the JCL EXEC statement.

In the following example for obtaining the program name, the name of the program involved with the output is IKJEFT01. It can be found on the PGM parameter of the EXEC statement:

```
//*  
//*-----  
//* Batch TSO job (PGM=IKJEFT01)  
//*-----  
//IKJEFT01 EXEC PGM=IKJEFT01,REGION=4096K,DYNAMNBR=50  
//*
```

3. Continue diagnosis depending on what program caused the problem.
 - For a component of z/OS, continue with [“Searching the problem reporting database”](#) on page 252.
 - For an IBM subsystem or product, continue diagnosis with the subsystem or product.
 - For an installation-provided program, including an installation exit routine, continue diagnosis with that program.

Related information:

- See [z/OS MVS System Management Facilities \(SMF\)](#) for the SMF reports.
- See [z/OS RMF User's Guide](#) for information.
- See [z/OS MVS IPCS Commands](#) for the ANALYZE subcommand.
- See [z/OS MVS JCL Reference](#) for the EXEC statement.
- See [z/OS MVS Diagnosis: Reference](#) to find the IBM component, subsystem, or product for a program identifier or module prefix.

Searching the problem reporting database

Search arguments are used to search problem reporting databases. If the problem you are diagnosing was already reported and the symptoms are in the database, the search produces a match. Searching is an iterative process; you might need to gather additional data and continue your search.

Steps for searching the problem reporting databases

1. Search a problem reporting database to determine if the problem was previously reported.

Include any of the following free-format search arguments that apply to the performance problem, and the word PERFORMANCE. For example:

 - System address space name
 - Function name
 - Product involved
 - Module or CSECT name
 - Resource names
 - Message ID
 - Symptoms created from information in STATUS CPU output

For more information, see [“Searching problem reporting databases”](#) on page 8.
2. If the search finds no match, remove some symptoms or add some symptoms. Search again. Continue searching for matches by adding or removing symptoms.
3. If the search finds that the problem was previously reported, request the problem fix. If not, continue with [“Gathering additional data for performance problems”](#) on page 253. Use the problem data gathered there to create more symptoms; use these symptoms in later searches.

Related information:

- See [“Searching problem reporting databases”](#) on page 8.
- See [z/OS MVS IPCS Commands](#) for the STATUS and VERBEXIT SYMPTOM subcommands.

Gathering additional data for performance problems

Gathering additional data will increase your chances of finding a match in the problem reporting databases. Use the procedures outlined in this section to gather additional data and continue searching the problem reporting databases.

Steps for gathering additional information for performance problems

1. Collect and analyze messages and logrec records about the problem. Use time stamps to select messages and software, symptom, and hardware records for logrec. Look at any messages or records that occurred before and during the time the problem was first reported. Look in the following:
 - The job log
 - A TSO/E user's ISPF transaction log or session manager log
 - The system log (SYSLOG) for the console with master authority or the alternate console
 - VERBEXIT MTRACE dump output, which shows the buffer for system messages
 - VERBEXIT LOGDATA dump output, which formats the logrec buffer
 - Logrec data set, formatted by EREP
2. Use IPCS to look at the dump.

Related information:

- See [z/OS MVS IPCS Commands](#) for the VERBEXIT LOGDATA and VERBEXIT MTRACE subcommands.
- For formatting of logrec records, see [Recording logrec error records in z/OS MVS Diagnosis: Tools and Service Aids](#).
- For explanations of the messages, see:
 - [z/OS MVS System Messages, Vol 1 \(ABA-AOM\)](#)
 - [z/OS MVS System Messages, Vol 2 \(ARC-ASA\)](#)
 - [z/OS MVS System Messages, Vol 3 \(ASB-BPX\)](#)
 - [z/OS MVS System Messages, Vol 4 \(CBD-DMO\)](#)
 - [z/OS MVS System Messages, Vol 5 \(EDG-GLZ\)](#)
 - [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#)
 - [z/OS MVS System Messages, Vol 7 \(IEB-IEE\)](#)
 - [z/OS MVS System Messages, Vol 8 \(IEF-IGD\)](#)
 - [z/OS MVS System Messages, Vol 9 \(IGF-IWM\)](#)
 - [z/OS MVS System Messages, Vol 10 \(IXC-IZP\)](#)
 - See the message book for a subsystem or program.
 - See [z/OS MVS Dump Output Messages](#) for dump output messages.

Analyzing a dump for performance problems

When all else fails, analyzing a dump might yield some additional information about performance problems. A dump captures a moment in time during the system execution and potentially a small amount of SYSTRACE data. This is typically not enough data to diagnose performance problems, but you might find a loop or a resource contention for the jobs that are performing poorly.

Steps for collecting and analyzing a dump for performance problems

Before you begin: Request a dump. If you did not obtain a dump, recreate the problem and request the dump using the following steps:

1. Enter a DUMP command to request an SVC dump and reply with SDATA options to dump global resource serialization control blocks and the nucleus.
2. Specify the address spaces that are experiencing performance reduction. For example, if the problem appears to be in JES2 or JES3, specify the JES address space in the reply.

The following example shows a request for an SVC dump with a sample of the parameters you might use:

```
DUMP COMM=(text)
REPLY id,ASID=1,SDATA=(GRSQ,NUC,CSA,SQA,TRT),END
```

Tip: Be sure to give your dumps meaningful names.

3. Use IPCS to format the dump.
4. If your dump contains the JES address space and your problem appears to be in JES2 or JES3, format the dump using IPCS as follows:
 - For a JES2 system, select JES2 in the Component Analysis panel. Then select JES2 Control Blocks from the JES2 Component Data Analysis panel. From the JES2 Control Block List panel, select the control blocks you wish to format.
 - For a JES3 system, select JES3 in the Component Analysis panel. Then select JES3 Control Block Information from the IPCS JES3 - Primary Options panel. You can also use JMF to analyze performance problems for JES3.

Related information:

- See *z/OS MVS Diagnosis: Tools and Service Aids* for requesting an SVC dump.
- See *z/OS MVS System Commands* for the DUMP operator command.
- See *z/OS JES2 Diagnosis* for information on diagnosing JES2 problems and using IPCS for JES2 diagnosis.
- See *z/OS JES3 Diagnosis* for information on diagnosing JES3 problems, using IPCS for JES3 diagnosis, and JMF.

Reporting performance problems to IBM

If you have completed the procedures listed in this chapter, cannot find a match for your performance problem in the problem reporting databases, and believe the problem is a defect in IBM code, call the IBM Support Center.

Provide the following problem data:

- Problem type: performance
- Search argument
- Dump, formatted by IPCS, online or printed
- System responses to DISPLAY and JES commands
- Parmlib members analyzed
- SMF records, if obtained in steps from the hang or wait procedure
- Hard copy log, beginning 30 to 60 minutes before the problem, or master trace, if not wrapped between the problem and dump
- Logrec records, beginning 30 to 60 minutes before the problem
- All printed output and output data sets related to the problem

- Name and level of the operating system(s) with a list of program temporary fixes (PTF) applied at the time of the problem and all installation modifications, exits, and products with other than Class A service

Related information:

- See [Chapter 27, “Reporting problems to IBM,”](#) on page 335 for more information.
- For formatting of logrec records, see [Recording logrec error records](#) in *z/OS MVS Diagnosis: Tools and Service Aids*.

Part 5. Diagnosing component-specific problems

After you identify a component-specific problem, use this information as a guide for operational problem determination of the problem. Each component section contains basic commands and function you can use for determining problems and collecting and sending documentation to IBM support. Additionally, where possible each component identifies component-specific problems, and supplies:

- Problem symptoms
- Investigation techniques
- Recovery actions
- Best practices

See:

- [Chapter 16, “Catalog component operational problem determination,” on page 259](#)
- [Chapter 19, “PDSE operational problem determination,” on page 277](#)
- [Chapter 20, “RRS operational problem determination,” on page 285](#)
- [Chapter 21, “System Data Mover \(SDM\) operational problem determination,” on page 301](#)
- [Chapter 22, “System logger operational problem determination,” on page 307](#)
- [Chapter 23, “VSAM component operational problem determination,” on page 311](#)
- [Chapter 24, “VSAM record-level sharing \(RLS\) operational problem determination,” on page 317](#)
- [Chapter 25, “XES and XCF operational problem determination,” on page 321](#)

Chapter 16. Catalog component operational problem determination

The major commands used to help in problem determination are: F CATALOG,LIST – Which shows the active tasks in the catalog address space D GRS,CONTENTION – displays resource contention F CATALOG,TAKEDUMP – Takes a dump of the Catalog address space F CATALOG,RESTART – Restarts the catalog address space

Catalog messages in general start with IEC3

Basic Catalog problem determination functions

Catalog provides the following functions to help with problem determination, data collection, and recovery:

- Use the DISPLAY GRS,CONTENTION command to display resource contention.
- Use the MODIFY CATALOG,LIST command to show the active tasks in the catalog address space.
- Use the MODIFY CATALOG,RESTART command to restart the catalog address space.
- Use the MODIFY CATALOG,TAKEDUMP command to take a dump of the Catalog address space.
- Make sure that all catalog checks are activated and running to warn you of impending catalog problems. For details, see Catalog checks (IBMCATALOG) in *IBM Health Checker for z/OS User's Guide*.

Catalog component-specific problems and recovery

- [“Hang in the Catalog address space or in the user address waiting on a request to the Catalog address space” on page 259](#)
- [“Damaged or broken catalogs” on page 260](#)
- [“Slow performance in various address spaces due to requests to the catalog address space taking excessive time ” on page 261](#)

Hang in the Catalog address space or in the user address waiting on a request to the Catalog address space

Symptoms

- Jobs hung in open or close processing
- TSO users hang when doing a simple ISPF 3.4 data set list
- Specific requests for a data set hang
- ISV monitoring products may indicate a problem accessing catalogs

How to investigate

1. Issue the command F CATALOG,LIST, wait 30 seconds to a minute, and reissue the command. The list of active tasks should have changed. Any tasks waiting over 5 seconds should be investigated. The output from the command will give the requesting address space name, type of request, how long the request has been waiting and in some cases the reason the request is waiting.
2. If any of the requests indicate waiting for an enqueue or reserve or latch issue the command D GRS,C. The resources that are likely to be in contention are SYSIGGV2, SYSZVVDS, SYSVTOC, SYS.IGG.CATALOGLATCHSET. If any of these resources are in contention you should obtain a dump of the catalog address space on each system that could be sharing catalogs. Issuing the command F CATALOG,TAKEDUMP and routing that to all systems in the sysplex will enable in depth problem

Catalog problem determination

determination. If there are systems outside the scope of a sysplex that may share some catalogs, data sets or volumes, then dumps should be obtained from those systems also.

Recovery actions

1. Issue the command F CATALOG,RESTART on the systems that appear to be hung.
2. F CATALOG,RESTART will be followed by an IEC363D if the F CATALOG,TAKEDUMP was not used to obtain a dump prior, answer 'y' to IEC363D and 'n' to the following IEC364D query to obtain a dump. Otherwise, reply 'n' to IEC363D to proceed with the restart. For more information, see "Restarting the Catalog Address Space" in *DFSMS Managing Catalogs*.
3. If a restart of catalog is not desirable, use instructions listed in "Ending a Catalog Request Task" in *DFSMS Managing Catalogs* to specifically end a service task believed to be hanging while holding a resource. Ensure that F CATALOG,TAKEDUMP is issued prior to taking action to obtain a dump

Best practices

Issue the D GRS,C on a periodic basis (perhaps 15 minute intervals). Concern yourself with any of the resources listed previously if the queuing increases and the same TCBs are holding resources.

Damaged or broken catalogs

Symptoms

- Jobs hung in open or close processing or other hang symptoms
- Unexplainable catalog errors during locate errors

How to investigate

1. Run the following commands in an IDCAMS batch job to identify errors:

```
//STEP1 EXEC PGM=IDCAMS,REGION=0M
//SYSPRINT DD SYSOUT=A //SYSIN DD *
DIAGNOSE ICFCATALOG INDATASET(catalog.name)
EXAMINE NAME(catalog.name) ITEST NODTEST
EXAMINE NAME(catalog.name) NOITEST DTEST
LISTCAT ENT(catalog.name) CAT(catalog.name) ALL
LISTCAT CAT(catalog.name) ENTRIES
```

2. If DIAGNOSE or EXAMINE indicate a condition code 8 or greater, run the following DSS PRINT jobs to capture a picture of the catalog:

```
//STEP2 EXEC PGM=ADRDSU,REGION=0M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
PRINT DATASET('catalog.name') -
  INDYNAM(volser)
PRINT DATASET('catalog.name.CATINDEX') -
  INDYNAM(volser)
```

Recovery actions

Catalog breakages can range from minor logical DIAGNOSE errors to major structural EXAMINE errors.

- For only DIAGNOSE errors, follow recovery actions listed under IDC21364I in *MVS System Messages, Vol 6*.
- If errors include EXAMINE errors, attempt the following actions:
 1. Determine if a cc=0 backup can be taken by using IDCAMS EXPORT. Ensure the number of records exported resembles a number close to the number of records you expect in the catalog.

```
//STEP1 EXEC PGM=IDCAMS,REGION=0M
//CATBKUP DD DSN=CAT.BACKUP,UNIT=SYSDA,
```

```
//          DISP=(NEW,CATLG),SPACE=(CYL,(10,10))
//SYSPRINT DD      SYSOUT=A
//SYSIN   DD      *
EXPORT catalog.name OUTFILE(CATBKUP) TEMPORARY
```

2. If EXPORT runs fine, refer to instructions listed in Chapter 6 "Backing Up and Recovering Catalogs" in *DFSMS Managing Catalogs* section "Recovering a BCS" to IMPORT this back-up, which will rebuild the catalog index.
3. If EXPORT fails, recovery from back-up will be necessary. With IBM products, this can be performed by locating the last good back-up taken of the catalog and then forward recovering the back-up to the current time. This will require a back-up produced by EXPORT, SMF 61,65, and 66 records, and use of the IBM supplied forward recovery utility ICFRU. Refer the chapter "Integrated Catalog Forward Recovery Utility" in *DFSMS Managing Catalogs* for instruction in using ICFRU.
4. If EXPORT fails and a good catalog back-up cannot be located, salvaging entries will be necessary. Refer to "Merging Catalogs" in *DFSMS Managing Catalogs* for direction on how to use REPRO MERGECAT to move entries from the broken catalog to a new catalog.

Best practices

To help avoid breakages, ensure that catalogs being shared between systems reside on DASD generated as shared and that the catalog is defined with share options (3 4). If the catalog is accessed from outside of the sysplex, ensure sharing is performed with RESERVEs rather than through a convert RNL for SYSIGGV2.

To ensure catalog recovery can be performed in a timely manner, take regular back-ups via IDCAMS EXPORT. Ensure DIAGNOSE and EXAMINE are ran prior to each EXPORT to ensure a valid back-up is obtained. If back-ups are taken with another product like DSS, an additional step of restoring the back-up and then exporting that copy will be required in order to use the ICFRU product since that product only accepts EXPORT output for input.

Slow performance in various address spaces due to requests to the catalog address space taking excessive time

Symptoms

Timeout events such as Abend522 or application specific errors.

How to investigate

1. Issue the command F CATALOG,LIST and wait 30 seconds to a minute before reissuing the command. The list of active tasks should have changed. Any tasks waiting over 5 seconds should be investigated. The output from the command will give the requesting address space name, type of request, how long the request has been waiting and in some cases the reason the request is waiting.
2. Issue D GRS,C to identify any contention on a particular catalog (SYSIGGV2) or volume (SYSZVVDS)
3. Issue F CATALOG,TAKEDUMP during time of slow performance to obtain a dump in case support must be engaged

Recovery actions

1. Cancel non-critical jobs that are involved in the contention
2. F CATALOG,RESTART. This will be followed by an IEC363D if the F CATALOG,TAKEDUMP was not used to obtain a dump prior, answer 'y' to IEC363D and 'n' to the following IEC364D query to obtain a dump. Otherwise, reply 'n' to IEC363D to proceed with the restart. For more information, see "Restarting the Catalog Address Space" in *DFSMS Managing Catalogs*.

Best practices

1. Review "Diagnosing a Catalog Performance Problem" in *DFSMS Managing Catalogs* for possible solutions and best practices
2. Investigate any jobs or users that may be issuing generic requests such as locates for *. These processes may hold a shared SYSIGGV2 ENQ for a longer time than typical locate requests.
3. If performance is concerning a particular catalog, research whether certain aliases directed to that catalog can be split off into a separate catalog as this may reduce contention during peak periods
4. If performance is concerning a particular volume, research whether there are other volume processing is occurring at that time such as defragmentation jobs or back-up processing.

Chapter 17. Global resource serialization problem determination

Global resource serialization (GRS) is a component of z/OS used to provide services for resource serialization in a multi-tasking, multi-processing environment. These services include:

- ENQ/DEQ/RESERVE and ISGENQ to obtain and release resources
- QQSCAN/ISGQUERY to query the status of a resource(s), as well as, owners/waiters.
- Latch services to create, obtain, release, and purge latches.

A latch is a serialization mechanism similar to ENQ/DEQ resource but designed for speed. The scope is limited to a single system, with limited RAS, and is only available to authorized programs.

- ISGECA provides enhanced contention analysis services

Exploiters of GRS services include other z/OS components, subsystems and applications. For more information about GRS see [z/OS MVS Planning: Global Resource Serialization](#).

Basic GRS problem determination functions

GRS provides the following functions to help with problem determination, data collections, and recovery:

- **Displaying GRS information:** The DISPLAY GRS command provides the following:
 - D GRS, S provides the status of systems in the GRS complex.
 - D GRS, ANALYZE and D GRS, C provide resource contention information by resource name, jobname, or device name.

The ANALYZE function is recommended because it detects dependencies (including deadlocks) between requestors, provides a complex wide view of system and systems ENQ resources, and takes time into consideration to highlight the possible root cause of the contention.

Contention or C provides an alphabetical list of resources that are in contention for both ENQ resources and latches.

- D GRS, C, LATCHID and D GRS, ANALYZE, LATCH provide latch contention information.

See the DISPLAY GRS in [z/OS MVS System Commands](#).

- **Dumping GRS information:** SVC dump option SDATA=(GRSQ,...) specifies that GRS ENQ/DEQ resource information be collected from the GRS address space when a dump is requested.

You can also use the SETGRS command to control how much information is dumped by the GRS dump exit. Options include: GRSQ=LOCAL | CONTENTION | ALL . CONTENTION is the default.

See [z/OS MVS System Commands](#).

- **ENQ/RESERVE/DEQ monitor tool:** The GRS monitor tool monitors ENQ, DEQ, and RESERVE requests, collecting data about the resources and their requestors. It is intended for short-term use only as it adds additional overhead to GRS processing. See the section on ENQ/RESERVE/DEQ monitor tool in [z/OS MVS Planning: Global Resource Serialization](#).
- **RMF ENQ Activity report;** The RMF ENQ Activity report provides information about resources that are in contention. See [z/OS RMF Report Analysis](#) for further information.
- **SMF87 subtype 1 and 2 records:** SMF 87 record subtype 1 helps you identify global generic queue scan issuers. SMF 87 record subtype 2 helps you diagnose application problems related to ENQ and DEQ and monitor the usage of ENQ/DEQ/ISGENQ/RESERVE requests. See the section on SMF record type 87 in [z/OS MVS Planning: Global Resource Serialization](#).

GRS component-specific problems

- “Job hang/delay due to GRS resource contention” on page 264
- “Job hang due to latch contention” on page 265
- “GRS related delays across the sysplex” on page 266
- “High CPU usage in GRS” on page 267
- “Wait state X'0A3' for GRS” on page 268

Job hang/delay due to GRS resource contention

When a workload is hung or running with degraded performance, often the first symptom to check for is GRS ENQ contention. ENQ resource contention occurs when, for example, job A-Z cannot run because it needs a resource held by Job B. See the topic on checking for ENQ contention problems in [z/OS MVS Planning: Global Resource Serialization](#).

Symptoms

Symptoms of a job hang/delay due to GRS resource contention include:

- Output from the D GRS ,C or D GRS ,ANALYZE commands show consistent ENQ resource contention:

```
D GRS,C
ISG343I 08.15.49 GRS STATUS 251
S=SYSTEMS SYSZAPPC APPC_PPLU_LOG
SYSNAME  JOBNAME      ASID      TCBADDR  EXC/SHR
STATUS
JB0      APPC              0252      006D9578 EXCLUSIVE
OWN
Z0       APPC              0157      007D9508 EXCLUSIVE
WAIT
JC0      APPC              015A      007D9508 EXCLUSIVE
WAIT
JA0      APPC              0153      007D9508 EXCLUSIVE
WAIT
JE0      APPC              0157      007D94A0 EXCLUSIVE
WAIT
J80      APPC              015B      007D9508 EXCLUSIVE
WAIT
J90      APPC              0155      007D94B8 EXCLUSIVE
WAIT
JF0      APPC              0151      007D9508 EXCLUSIVE
WAIT
NO REQUESTS PENDING FOR ISGLOCK

D GRS,ANALYZE,WAITER
ISG349I 03.25.55 GRS ANALYSIS 273
LONG WAITER ANALYSIS:  ENTIRE SYSPLEX
WAITTIME  SYSTEM  JOBNAME E/S  SCOPE QNAME  RNAME
00:45:27 SYSG   RRS    *E*  SYSS SYSZATR SYSPLEX-RESTART
BLOCKER  SYSJ   RRS    E    OTHER BLOCKERS: 0 WAITERS: 1
00:39:35 SYSG   RRS    *E*  SYSS SYSZATR SYSPLEX-RESTART
BLOCKER  SYSJ   RRS    E    OTHER BLOCKERS: 0 WAITERS: 1
00:01:43 SYSG   DCOTCMPM*E* SYSS SYSDSN  CICS.PROD.APPLLIB
BLOCKER  SYSG   CICS1  S    OTHER BLOCKERS: 1 WAITERS: 2
00:01:43 SYSG   DCOTCMPM*E* SYSS SYSDSN  SYSCICS.PARMLIB
BLOCKER  SYSG   CICS1  S    OTHER BLOCKERS: 1 WAITERS: 2
00:01:43 SYSG   DCOTCMPM*E* SYSS SYSDSN  CICS.TRAN.APPLLIB
BLOCKER  SYSG   CICS1  S
00:01:43 SYSG   DCOTCMPM*E* SYSS SYSDSN  CICS.ACF2.PARMLIB
BLOCKER  SYSG   CICS1  S    OTHER BLOCKERS: 1 WAITERS: 2
00:01:37 SYSG   CICS4  *S*  SYSS SYSDSN  CICS.ACF2.PARMLIB
```

- Messages IOS071I and IOS431I, which the system issues when there are delays due to hardware reserves held. IOS431I identifies the system and resource being delayed and IOS issues the D GRS,DEV=xxxx command on the owning system to identify the jobname that is delayed:

```
IOS071I 4341,**,*MASTER*, START PENDING
IOS431I DEVICE 4341 RESERVED TO CPU=9999992097,LPAR ID=01
SYSTEM=SYS2
```

How to investigate

- Issue the D GRS command repeatedly to observe whether the owner/blocker is consistently the same job or changes.
- Request an SVC dump of GRS and the job or jobs identified as the holder or blocker, specifying the following parameters on the DUMP command:

```
JOBLIST=(GRS, jobname, SDATA=(RGN,GRSQ, COUPLE, XESDATA, SQA, TRT, CSA, ALLNUC)
```

where *jobname* is the job or jobs identified in message ISG343I as the holder/blocker of the resource. You can use this dump to identify the root cause of the enqueue lockout.

- Check SYS1.LOGREC for software errors reported for the holding or blocking job *jobname*.

Recovery actions

If necessary, issue the CANCEL command to cancel the holding or blocking job *jobname*. This prompts the GRS termination resource manager to cleanup resources held. If the output from the D GRS, C command shows that the job *jobname* still owns the resource even though it is no longer active, try purging the initiator with a \$PI JES2 command or using the FORCE command to terminate the address space where job *jobname* was last running.

Job hang due to latch contention

Latches are used as a serialization mechanism similar to ENQ resources, but allow more granularity. Latches are used to serialize resources (like control blocks or data sets) within a single address space or across several address spaces in a single MVS system with minimal overhead. Latch manager callable services are for authorized programs only - see [z/OS MVS Programming: Assembler Services Guide](#).

Symptoms

The symptom of a job hang due to latch contention is when output from the D GRS, C, LATCH or D GRS, ANALYZE, LATCH commands show consistent contention for a latch:

```
ISG343I 08.56.53 GRS STATUS 165
LATCH SET NAME: SYS.BPX.A000.FSLIT.FILESYS.LSN.01
CREATOR JOBNAME: OMVS          CREATOR ASID: 0010
LATCH NUMBER: 120
REQUESTOR ASID  EXC/SHR  OWN/WAIT  WORKUNIT  TCB  ELAPSED TIME
SRVEMUG  0120  EXCLUSIVE  OWN       007FF1D0  Y   05:43:06.453
SRVEMUG  0168  SHARED    WAIT      007FF1D0  Y   05:42:03.258
SRVEMUG  0131  SHARED    WAIT      007FF1D0  Y   05:40:11.452
SRVEMUG  0181  EXCLUSIVE  WAIT      007FF1D0  Y   05:35:52.075
SRVEMUG  014F  SHARED    WAIT      007FF1D0  Y   04:24:45.355
SRVEMUG  0151  EXCLUSIVE  WAIT      007FF1D0  Y   03:57:33.404
SRVEMUG  0167  EXCLUSIVE  WAIT      007FF1D0  Y   03:48:56.081
SRVEMUG  0110  EXCLUSIVE  WAIT      007FF1D0  Y   00:05:43.140
```

How to investigate

- Request an SVC dump of the owner and creator job name or names specifying the following parameters on the DUMP command:

```
JOBNAME=(jobx, jobz, ...), SDATA=(RGN, SQA, GRSQ, CSA, TRT, COUPLE, ALLNUC)
```

where *jobx* is the **creators** jobname and *joby* is the **owners** jobname, as displayed in message ISG343I.

- Check for software entries in SYS1.LOGREC for the owner jobname
- If the latch in contention has a name beginning with SYS.BPX, indicating that z/OS UNIX System Services (z/OS UNIX) owns it, issue a D OMVS command for additional information regarding the latch contention.

Recovery actions

- If you can, determine the latch creator *jobx* component and refer to the problem determination section for that product or component, if one exists, for recovery actions.
- Try canceling the owner *joby*, using the system CANCEL command to drive the GRS termination resource manager to cleanup latches held.
- If the output from the D GRS, C command shows that the job *joby* still owns the resource even though it is no longer active, try purging the initiator with a \$PI JES2 command or using the FORCE command to terminate the address space where job *joby* was last running to force cleanup.

GRS related delays across the sysplex

There is an instance of GRS on every system in a sysplex or ring. If one instance of GRS is not healthy, it will affect the other instances of GRS adversely. The following is a list of ways to investigate the most common symptoms and recommended actions to take:

1. **Check for sysplex partitioning in progress** for any system in sysplex by checking for an outstanding WTOR for IXC102A or IXC402D using the D R,L command:

```
*3574 IXC102A XCF IS WAITING FOR SYSTEM XYZ
DEACTIVATION. REPLY DOWN WHEN MVS ON SQ0 HAS BEEN
SYSTEM RESET
*2170 IXC402D E003 LAST OPERATIVE AT 11:17:03. REPLY DOWN AFTER
SYSTEM RESET, OR INTERVAL=SSSSS TO SET A REPROMPT TIME.
```

Recovery action:

- Reply 'DOWN' as quickly as possible to prevent sysplex-wide sympathy sickness. See [“Sysplex sympathy sickness”](#) on page 322.
 - Leverage sysplex failure management (SFM) and base control program internal interface (BCPii). See [z/OS MVS Setting Up a Sysplex](#).
2. **Check for contention on the GRS list lock** by looking for message ISG361A on every system in the sysplex. Any system that does not display message ISG361A is probably the one owning the list lock in contention. The system issues message ISG362I when the contention for the list lock is relieved.

```
*ISG361A GLOBAL REQUEST PROCESSING DELAYED ON
SYSTEM SYS1. THE SYSTEM HAS BEEN WAITING 174 SECONDS FOR THE LISTLOCK.
ISG362I GLOBAL REQUEST PROCESSING RESUMED ON SYSTEM
SYS1.

ISG362I GLOBAL REQUEST PROCESSING RESUMED ON SYSTEM SYS1
THE SYSTEM OBTAINED THE LIST LOCK.
```

Recovery actions: Rebuild the ISGLOCK structure to try freeing up contention on the GRS list lock, if it exists:

```
SETXCF START, RB, STRNM=ISGLOCK
```

3. **Check to see if XCF has detected that GRS is impaired.** GRS joins the sysplex as a critical XCF member and XCF monitors its health, which can impact the entire sysplex. If XCF detects GRS as hung or not functioning on one system, it issues messages IXC633I, IXC636I and IXC635E:

```
IXC633I GROUP SYSGRS MEMBER SYSX JOB GRS ASID 0007
DEEMED IMPAIRED AT 07/06/2014 04:04:58.009741 ID: 1.1
LAST MSGX: 07/06/2014 04:04:19.062379 17 STALLED 78 PENDINGQ
LAST GRPX: 06/30/2014 05:21:54.526031 1 STALLED 1 PENDINGQ
LAST STAX: 07/06/2014 04:03:21.867614 1 STALLED

IXC636I GROUP SYSGRS MEMBER SYSX JOB GRS ASID 0007
IMPAIRED, IMPACTING FUNCTION Not Specified

*IXC635E SYSTEM SYSX HAS IMPAIRED XCF GROUP MEMBERS
```

Recovery actions:

- Request dumps of GRS on all systems in the sysplex as documentation for the problem using the following command:

```
DUMP Title=(GRS is Hung)
R xx, JOBNAME=GRS, SDATA=(RGN, SQA, TRT, XESDATA, COUPLE,
SQA, NUC), REMOTE=(SYSLIST=(*'GRS')), SDATA), END
```

- Rebuild the ISGLOCK structure to try to free up contention on the GRS list lock, if it exists:

```
SETXCF START, RB, STRNM=ISGLOCK
```

4. Look for a non-responsive instance of GRS:

- Issue a D GRS, C command and look for message ISG378I to determine which instance of GRS is not responding:

```
ISG378I GRS QSCAN ERROR COMMUNICATING WITH SYSTEM sysname, DIAG=diag
```

- Look for dumps accompanying abend X'30D' for module ISGGTRM0 (GRS termination resource manager) on an active system that is timing out while waiting for a non-responsive system.

```
END OF MEMORY RESOURCE MANAGER HANG DETECTED: TCB =
007E2828, NAME = ISGGTRM0 - SCSDS
```

Recovery actions:

- Use the RMF III Job Report Menu (or equivalent) to verify that GRS is not contending for resources on any of systems in the sysplex. See the section on the monitor III primary menu in [z/OS RMF User's Guide](#).
- Use Runtime Diagnostics to check to see if GRS is contending for any suspend locks on any of the systems in the sysplex. See [Part 2, "Runtime Diagnostics," on page 33](#) and ["LOCK - Local lock suspension" on page 44](#).
- If a system or instance of GRS is hung or non-responsive, either wait for XCF to partition the critical member from the sysplex or consider partitioning it to minimize impact to other systems in the sysplex (see ["Sysplex sympathy sickness" on page 322](#)). Use the VARY XCF, *sysname* command to partition the system. See [z/OS MVS System Commands](#).

Note that GRS is not restartable.

High CPU usage in GRS

GRS runs in cross memory mode under an application, as well as in SRB and TCB mode in the GRS address space. If GRS is incurring higher than normal CPU use, it is typically due to an increase in the rate and type of requests being driven to GRS. To address the problem, you must identify the application or applications driving the increased GRS workload.

You may also notice this problem when an application encounters a delay on a GRS request.

How to investigate and recover

Using an online monitor, identify the GRS unit of work that is consuming a higher than normal amount of CPU. Then take one of the following actions:

- If the unit of work or TCB for the entry point is identified as ISGWDRVR, the increased CPU usage might be the result of higher number of global ENQ requests or reduced performance on operations to the ISGLOCK structure. Gather a before and after picture of ENQ rates using:
 - The RMF ENQ Activity report, see [z/OS RMF User's Guide](#).
 - The ENQ/RESERVE/DEQ monitor tool, see sample, job A-Z cannot run because it needs a resource held by Job B. See the topic on checking for ENQ contention problems in [z/OS MVS Planning: Global Resource Serialization](#).
 - Start collecting SMF 87 subtype 2 records which monitor ENQ/DEQ/RESERVE activity. See the section on diagnosing global resource serialization in [z/OS MVS Planning: Global Resource Serialization](#).

Global resource serialization problem determination

- Activate monitoring of GQSCAN requests using SMF 87 subtype 1 records. See the section on SMF record type 87 in *z/OS MVS Planning: Global Resource Serialization*.

Gather a before and after picture of service times for requests going to the ISGLOCK structure using the Review RMF CF Activity Reports. Make sure that the CFs have adequate central processor (CP) resource, see *z/OS MVS Setting Up a Sysplex*.

- If the TCB for the entry point is identified as ISGQDRVR, this indicates an increase in processing of GQSCAN or ISGQUERY requests. Non-specific requests are more costly in GRS CPU time. To resolve this problem, activate SYSGRS CTRACE with options FLOW9 and REQUEST2 or request an SVC dump that includes GRS to analyze request types and rates:

```
TRACE CT,8M,COMP=SYSGRS
R xx,OPTIONS=(FLOW9,REQUEST2),END

or

DUMP COMM=(Dump of GRS)
R xx,JOBNAME=GRS,SDATA=(RGN,TRT,GRSQ,COUPLE,XESDATA,
SQA,ALLNUC),END
```

- If the TCB for the entry point is identified as ISGWCCNF, this instance of the GRS is the contention notification server (CNS) for the sysplex, and it might be processing a higher amount of contention in the system or sysplex than it has the capacity for. Ensure that the CNS system has enough processing capacity to handle contention notification requests efficiently or else use the SETGRS CNS=*sysname* command to assign the CNS function to another system in the sysplex with greater capacity.
- Check false lock contention rates for the ISGLOCK structure in a RMF CF Activity Report - see *z/OS RMF Report Analysis*. If false contention exceeds 2% then consider increasing the size of the ISGLOCK structure. A high false contention rate implies that GRS is experiencing contention for lock entries on the ISGLOCK structure as a result of many different resources hashing to the same lock entry. Use the Coupling Facility Structure Sizer (CFSizer) tool to estimate ISGLOCK structure size, then increase the size of the ISGLOCK structure by updating and activating a new CFRM policy. A larger physical lock structure will provide more locks for GRS to hash the resource names into.

You can find the CFSizer tool at [Coupling Facility sizer \(www.ibm.com/support/docview.wss?uid=isg3T1027062\)](http://www.ibm.com/support/docview.wss?uid=isg3T1027062).

Wait state X'0A3' for GRS

The following sections show how to diagnose and recover from wait state X'0A3' for GRS:

- [“Wait state X'0A3' with reason code X'09C' or X'0BC' after a CF upgrade” on page 268](#)
- [“Wait state X'0A3' with reason code X'064', X'068', or X'06C' indicating a GRS RNL mismatch” on page 269](#)

Wait state X'0A3' with reason code X'09C' or X'0BC' after a CF upgrade

The GRS STAR environment uses sysplex resources and must be able to allocate the ISGLOCK structure in the coupling facility and access the sysplex and CFRM couple data sets during initialization. If either of these resources are unavailable, GRS issues wait state X'0A3' during system initialization with one of the following reason codes:

X'09C'

GRS was unable to query the CFRM couple data set for information about the ISGLOCK structure.

X'0BC'

GRS was unable to allocate the ISGLOCK structure.

If this is the first system being IPLed into a GRS STAR complex, the following messages will be issued and you will be prompted to take action:

```
ISG379E GRS UNABLE TO CONNECT TO THE ISGLOCK STRUCTURE.
VALIDATE THAT THERE IS A COUPLING FACILITY DEFINED IN THE CFRM
POLICY AND THAT IT IS PHYSICALLY CONNECTED TO THE SYSTEM.
```

ENSURE THAT THE CF IS IN THE PREFLIST FOR THE ISGLOCK STRUCTURE.

ISG380E GRS UNABLE TO USE THE ISGLOCK STRUCTURE DUE TO INSUFFICIENT NUMBER OF LOCKS. VALIDATE THAT THE SIZE AND INITSIZE IN THE CFPM POLICY ARE SUFFICIENT FOR ISGLOCK AND THAT THE CF HAS THAT STORAGE AVAILABLE.

ISG381D GRS UNABLE TO INITIALIZE IN STAR MODE. RELOAD THE SYSTEM OR REPLY NONE TO CONTINUE THE IPL IN NONE MODE.

Otherwise, see [“Unable to bring new or upgraded coupling facility into use”](#) on page 326 to solve this problem.

Wait state X'0A3' with reason code X'064', X'068', or X'06C' indicating a GRS RNL mismatch

The entries in the GRSRNLxx member used by an IPLing system must match the RNLs in use by the rest of the sysplex exactly. Not only must each entry be the same, they must also be in exactly the same order. If an IPLing system fails with a wait state X'0A3' with reason code X'064', X'068', or X'06C', then there is an RNL mismatch.

To solve the mismatch, on one of the systems that is already up and running, issue a SET GRSRNL=xx command, where xx identifies the GRSRNLxx parmlib member.

- If the SET GRSRNL=xx command is accepted, then something has changed in the GRSRNLxx member. Retry the IPL.
- If the SET GRSRNL=xx command in step one is not accepted, make a copy of the GRSRNLxx member and try to IPL using the new copy. For example, if you copy the GRSRNLxx member into a member called GRSRNLyy, IPL with GRSRNL=yy. If the IPL'ing system cannot find GRSRNLyy, then it is probably not using a correctly cataloged SYS1.PARMLIB member for GRSRNLyy. Use the correct SYS1.PARMLIB member and retry the IPL.

Note: This procedure works only in an environment where the systems in the sysplex are the same as the systems in the GRS complex.

Chapter 18. JES2 operational problem determination

In a JES2 multi-access spool (MAS), sometimes a problem can occur on one system that can affect the whole MAS. Some of the problems include checkpoint locks, high CPU, and long dispatching times for processor control elements (PCE). The information in the following section covers the steps that you can take to solve the problems as quickly as possible.

Many JES2 problems require you to collect additional documentation to diagnose the problem. In general, the three essential documentation types are:

JES2 DUMP
SYSLOG
LOGREC

Many JES2 errors produce dumps, and alert the operator with the message sequence: \$HASP095/\$HASP096 and \$HASP088. The message information for each event is captured in the SYSLOG. The SYSLOG information is often enough to find an existing APAR or fix, but sometimes a full dump is required.

The JES2 problem determination information covers the following topics:

- [“High CPU usage in the JES2 address space” on page 271](#)
- [“JES2 contention” on page 272](#)
- [“JES2 checkpoint contention” on page 274](#)

In addition, the considerations and the resources needed for a checkpoint reconfiguration are covered in [“JES2 checkpoint reconfiguration” on page 274](#).

For complete JES2 diagnosis information including symptoms, how to recover, and information about setting dumps and traces, as well as using IPCS, see [z/OS JES2 Diagnosis](#)

High CPU usage in the JES2 address space

Symptoms

JES2 is taking a longer time to process work. You notice slow throughput of jobs through the conversion or the execution phase of job processing. There are occasions when you might notice that print is backed up. Print backup is mainly with vendor products that are using the SYSOUT application programming interface (SAPI).

Note: To record status in SYSLOG consider installing the **\$DPERFDATA** command sequence as an automatic command.

Diagnosing the problem

The best way to monitor for this condition is to use PERFDATA that measures the amount of CPU being used by each processor control element over a designated time period. Using PERFDATA can also give you a better indication of where to focus your attention.

Note: To record status in SYSLOG, consider installing the **\$DPERFDATA** command sequence as an automatic command.

Resolving the problem

1. Enter the following command sequence:

```
$D PERFDATA          <Display historical data>
D A,JES2
```

```
$T PERFDATA(*),RESET <Reset counts>
```

2. Wait 15 minutes.
3. Reissue **\$DPERFDATA** for time multiple time periods. Review SYSLOG for one or more process control elements (PCE) that might have a high CPU percent. The following example shows two PCEs with CPU data:

```
INTERVAL=31:23.194512, Time since last $TPERFDATA  
$HASP660 PCENAME=CKPT, TIME=7.672396, CPU=7.022071, CPU%=18.10,  
$HASP660 PCENAME=JQRP, TIME=3.357481, CPU=2.259784, CPU%=5.82,  
$HASP660 QSUSE_TIME=0.989827, IOCOUNT=6626, CKPT_COUNT=111476,
```

If the analysis of the initial data warrants it, use the following commands to refine the PERFDATA to a specific PCE:

- \$D PERFDATA(INITSTAT)** – JES2 initialization stats.
- \$D PERFDATA(QSUSE)** – PCE \$QSUSE summary.
- \$D PERFDATA(PESTAT)** – detailed PCE stats.
- \$D PERFDATA(SAMPDATA)** – WLM Sampling data.
- \$D PERFDATA(CPUSTAT)** – PCE CPU usage.
- \$D PERFDATA(CKPTSTAT)** – CKPT read/write stats.
- \$D PERFDATA(SUBTSTAT)** – JES2 subtask.
- \$D PERFDATA(EVENT)** – \$EVENTS captured.
- \$D PERFDATA(*)** – All of the above.
- \$T PERFDATA(*), RESET** – Resets performance data.

When you are reviewing the SYSLOG, you can use the following commands on the results:

- \$JDHISTORY**
- \$JDDetails**
- \$JDSTATUS**
- \$JDMONITOR**

4. If the SYSLOG investigation detects a looping job or task, cancel it. If a task is causing the issue, it might need to be recycled.
5. If the causing job or task is not identifiable, before you can restart JES2, you must take a console dump of JES2. The console dump is more useful than the dump created by the **\$PJES2,abend** command. The console dump command is as follows:

```
DUMP COMM=(dump title)  
Rxx, JOBNAME=JES2, SDATA=(PSA, SQA, LSQA, RGN, TRT, LPA, CSA, NUC, SUM, SWA), END
```

Best practice: Install the **\$DPERFDATA** command sequence as an automatic command. Periodically issue **PERFDATA** and **MONITOR** commands to obtain historical data. Doing so gives you a snapshot of statistics for normal JES2 performance to compare against JES2 error conditions. Scan PSP buckets for all HIPER APAR activity and especially for any that are associated with SYSOUT application programming interface (SAPI) processing.

JES2 contention

In a JES2 multi-access spool (MAS), one member is dominating the processing cycles.

Symptoms

JES2 monitors performance and alerts the operator in \$WTO messages that identify the potential issue. The HASP92xxx message indicates a delay with JES2 processing. The main messages are HASP9201 through HASP9213. The following messages are issued the most frequently:

- \$HASP9203 - Long PCE dispatch time.

- \$HASP9207 - JES2 checkpoint lock held.
- \$HASP9301 - Alert cleared.

Causes

Messages HASP9201 through HASP9213 can be an indication of a tuning issues within the JES2 multi-access spool (MAS).

Diagnosing the problem

- \$HASP9203 indicates that a particular PCE is dispatched from a JES2 dispatcher for the duration indicated in the message. A long PCE dispatch might indicate a problem or complex processing that is taking a long time to complete. It can also be caused by a loop that spans a large area of memory, or involves calling a service in a separate module.
- \$HASP9207 is issued when the checkpoint data set is held for significantly longer than was specified on the MASDEF HOLD= statement. The alert is issued with other alerts to inform you of the multisystem aspect of this problem. If this member is holding the checkpoint lock, other members of the JES2 MAS cannot process JES2 requests. If this message appears without any other alerts, it might indicate a checkpoint problem.

Resolving the problem

For \$HASP9203, long PCE dispatches can be an indication of a tuning problem, a complex process that takes a long time to complete, or an error situation. If the JES2 main task is not getting enough CPU resource, it can appear to the monitor as a loop.

1. First, enter the **D A, JES2** command, or use a performance monitor to ensure that JES2 is indeed consuming CPU.
2. After you eliminate CPU performance problems, check if the PCE displayed is the command processor. Examine the command that was entered. Some commands can take a long time to complete, such as a command to alter all 200,000 jobs or 500,000 JOEs in the system. If progress is being made (the job number that is displayed is being updated) for a complex command, then it might be normal processing.

Message \$HASP9207 indicates the multi-member nature of another problem that exists. However, if this message appears without any other monitor alerts, it might indicate one of the following problems:

- A potential problem with the JES2 checkpoint PCE. Issue the **\$DPCE (CKPT) , DETAILS** command to determine where the checkpoint PCE is waiting. If the PCE is not being dispatched (for example, the wait time not advancing) this is the cause of the problem. If the wait is in exit code, review the exit code. If the wait is in IBM-distributed code, request a dump of the JES2 member that is issuing the message and contact IBM support.
- A problem with the checkpoint data set. Consider issuing the **\$TCKPTDEF, RECONFIG=YES** command to reconfigure the checkpoint and use a different checkpoint data set.

Remember: Before calling IBM support, it's important to capture the required problem determination documentation.

If there are repeated HASP92xx messages, it's time set the following SLIP and capture a dump:

```
SLIP SET,MSGID=HASP9207,ACTION=SVCD,ID=HASP9207,
SDATA=PSA,SQA,LSQA,SWA,RGN,TRT,LPA,CSA,GRSQ,SUM,ALLNUC,COUPLE,XESDATA),
END
```

- For complete SLIP syntax and usage information, see the [SLIP command](#) in *z/OS MVS System Commands*.
- For help with the **\$TCKPTDEF, RECONFIG=YES** command, see: [z/OS JES2 Commands](#).

JES2 checkpoint contention

Checkpoint contention reported by \$HASP263. The system that issues the message is the victim. Quick response to the message is urged. First-time data capture of the correct system provides for the most timely problem determination and solution. Message \$HASP263 typically means that there is a problem with the JES2 multi-access spool (MAS). There can be other reasons why the system issues this message ranging from tuning of the hold and dormancy values on CKPTDEF to workload spikes and changes on a different member. The system can also issue this message for JES2 checkpoint contention. When a checkpoint is on a DASD volume that is not a dedicated volume and other data sets are being accessed, it delays the checkpoint.

Symptoms

Checkpoint contention that is reported by the message \$HASP263:

```
WAITING FOR ACCESS TO__JES2 CHECKPOINT__VOLUME
volser_____>

>__ _LOCK HELD BY MEMBER member_name
```

Diagnosing the problem

The member of the MAS that is trying to acquire the checkpoint lock issues message \$HASP263. SYSLOG displays the complete message text for the HASP263 message. The text identifies which member is holding the checkpoint lock. When you gather the problem documentation, it is important to understand that the system that issues the message is the victim. When reporting the problem to IBM, you need to collect the SYSLOG and the console dump from the member that is identified in message \$HASP263.

Resolving the problem

A console dump of the member is more timely than a dump of the abended JES2 member (by using \$P JES2, ABEND).

1. First, collect a console dump of the member that is identified in \$HASP263.

```
DUMP COMM=(dump title)
Rxx,JOBNAME=JES2,SDATA=(PSA,SQA,LSQA,RGN,TRT,LPA,CSA,NUC,SUM,SWA),END
```

2. Review the SYSLOG. The SYSLOG can reveal a cause, such as a vital task abend. If so, you can try to recycle the task to solve the lock contention.
3. If the member that is holding the checkpoint lock failed, enter the following command to try to reset the lock (and to avoid a restart of JES2):

```
$eckptlock,heldby=failingsystem
```

4. After all the problem documentation is collected, contact IBM service. A hot-start of JES2 can typically clear the problem, but be sure to first collect all the problem documentation.

JES2 checkpoint reconfiguration

Considerations and resources that are needed when JES2 checkpoint reconfiguration is necessary.

Checkpoint reconfiguration can be needed as a result of a hardware or a software issue. Issues can stem from a failed coupling facility structure or a DASD volume. Checkpoint reconfiguration can also be needed to enlarge a checkpoint data set to handle an increase in workload such as higher number of job numbers and all the associated control blocks. Keep in mind that JES2 does not support moving the checkpoint data set from one volume serial to another unless the checkpoint dialog is used.

You can use the checkpoint reconfiguration facility to dynamically redefine the checkpoint data set definitions for your JES2 multi-access spool (MAS) configuration. Whether the data sets is on a DASD

volume or on a coupling facility structure, the reconfiguration uses its operator dialog to perform the following tasks:

- Replace a checkpoint data set.
- Discontinue the use of a checkpoint data set.
- Resume by using a previously suspended checkpoint data set.
- End JES2 during the dialog.
- Start JES2 by using a new checkpoint data set.

JES2 does not support moving the checkpoint data set from one volume serial to another unless the checkpoint dialog is used.

Drive a checkpoint reconfiguration by using the information in *z/OS JES2 Initialization and Tuning Guide*.

It is critical to examine and reply to each message that the system issues during the reconfiguration process. JES2 initiates the checkpoint reconfiguration when:

- During initialization processing:

Either JES2 might not determine the checkpoint data set specifications, or JES2 requires operator verification of the data set specifications.

You specified **PARM=RECONFIG** as a JES2 start option.

You specified **RECONFIG=YES** on the CKPTDEF initialization statement.

- Because of an I/O error. JES2 automatically initiates the reconfiguration when an I/O error occurs during a read or write of the checkpoint data set. For the JES2-initiated checkpoint reconfiguration, see the topic about "I/O error" in *z/OS JES2 Initialization and Tuning Guide*
- The coupling facility structure becomes volatile. As result of processing based on your **CKPTDEF VOLATILE=** specification when a coupling facility on which a checkpoint is becomes volatile. A coupling facility structure is considered volatile if when power to the coupling facility is lost, all data on the facility is lost. A volatile structure might not be an immediate problem, but is serious enough that JES2 prompts you to take precautions to prevent potential data loss.

Note: A coupling facility is always considered volatile unless you take measures to protect it against the loss of power. For a description of how to monitor the volatility of coupling facilities, see IBM Systems Processor Resource/Systems Manager (PR/SM) Planning Guide on IBM Resource Link. For information about the JES2-initiated checkpoint reconfiguration when the coupling facility structure becomes volatile, see the topic about "Volatile coupling facility" in *z/OS JES2 Initialization and Tuning Guide*

- Operator initiated. Depending on the way your installation uses DASD, coupling facility structures, or both to maintain your checkpoint data, you might need to occasionally change your checkpoint definitions. For example, you might choose to initiate a checkpoint reconfiguration because the device on which the checkpoint data set is located is scheduled for service. However, if your checkpoint is on a coupling facility, it might be simpler to use a system-managed rebuild to move your checkpoint to another coupling facility. See the topics about "Moving a JES2 checkpoint to a coupling facility using system-managed rebuild" and "Operator-initiated entrance into a checkpoint reconfiguration dialog" in *z/OS JES2 Initialization and Tuning Guide*

The reconfiguration allows minimal disruption across all checkpoint data set failures, which prevents the need for a JES2 restart. Review the reconfiguration process by using the reference material that is listed here and practice the reconfiguration on a test system.

Chapter 19. PDSE operational problem determination

Partitioned Data Set Extended (PDSE) is an advanced implementation of the MVS partitioned (PO) data set organization. PDSE files:

- Have a directory structure which does not require pre-allocation and which will dynamically expand to fit the number of members stored within.
- Store members in a manner where periodic data set compression is not needed to reclaim "dead" space.
- Automatically reclaim space when a member is deleted or replaced.

PDSE can be used in place of a PDS to store data, or to store executable programs in the form of program objects. PDSE data sets are processed using most of the same access methods and macros as PDS files.

PDSE data set access is accomplished by means of a client server relationship between the PDSE user (batch jobs, started tasks, or TSO users) and the SMSPDSE and, optionally, the SMSPDSE1 address space.

Basic PDSE problem determination functions

PDSE provides the following functions to assist with first failure data capture and problem determination.

- System generated ABEND0F4 dumps
- EREP reports
- SYSLOG messages
- The SMSPDSE and the SMSPDSE1 monitor task
- PDSE system commands
- The IEBCPDSE PDSE Verification Utility program.

Collecting documentation for PDSE

The support center routinely requests a standard set of diagnostic data for PDSE issues, included is:

- The first system generated ABEND0F4 dump.
- Detailed EREP report or raw LOGREC data that brackets the dump by two hours. Raw LOGREC is the preferred format.
- SYSLOG or OPERLOG for the same time frame as LOGREC/EREP.
- In the event of PDSE data set corruption, an ADRDSSU physical dump of the affected PDSE data set.

In event that a system generated dump was not taken or if the dump was DAE suppressed then an operator initiated dump of the PDSE address space or spaces may be required.

Pre-allocating an IEADMCxx PARMLIB member, to dump the PDSE processing environment, will save time. Create the member as follows:

```
JOBNAME=(*MASTER*, SMSPDSE*),
SDATA=(PSA, CSA, SQA, GRSQ, LPA, LSQA, RGN, SUM, SWA, TRT, COUPLE, XESDATA)
```

Issue the following command when you need to obtain a dump:

```
DUMP COMM=(title),PARMLIB=xx
```

Where, 'xx', is the suffix of the IEADMCxx parmlib member that contains the SMSPDSE dump options.

PDSE specific problems

Common PDSE issues that might require Support Center assistance:

- [“ABEND0F4 failures” on page 278](#)
- [“MSGIGW038A possible PDSE problems” on page 279](#)
- [“PDSE data set corruption” on page 281](#)
- [“Failure of the SMSPDSE or SMSPDSE1 address space” on page 283](#)

ABEND0F4 failures

PDSE code performs extensive validity checking and will issue a system 0F4 abnormal termination (ABEND0F4) completion code for internally detected errors. These ABEND0F4 are accompanied by unique return and reason codes that identify the module where the failure occurred and the reason for the abnormal termination.

By default the ABEND0F4 dump will include the primary, home and secondary address spaces insuring that both the SMSPDSE/SMSPDSE1 server and client address spaces are captured in the dump.

If more than one dump is generated then the first dump in the series is usually the most pertinent to the cause of the failure, as subsequent dumps may be recording issues that were encountered during recovery processing.

Note: The first dump in a series can be something other than a S0F4 ABEND, such as ABEND0C1, ABEND0C4, or ABEND30D.

Symptoms

ABEND0F4 failures in batch jobs, TSO users, or started task address spaces.

How to investigate

- Note the DUMP TITLE from the system IEA794I (msgIEA794I) message and record the following data:
 - The failing CSECT name
 - Offset within the CSECT
 - Maintenance level
 - Return code
 - Reason code.

Note: The reason code is recorded in general purpose register 0 and the return code is recorded in general purpose register 15. Check these two registers if the reason and return codes are not readily apparent.

In most instances this information is sufficient enough to identify a matching APAR.

The following is a sample S0F4 DUMP TITLE, for an ABEND0F4 RC24 RSN1467A02B failure out of IGWISRCH+0E92.

```
DUMP TITLE=COMPID=DF115,CSECT=IGWISRCH+0E92,DATE=03/18/11,MAINTID=
NONE,ABND=0F4,RC=00000024,RSN=1467A02B
```

- Search the Technical Support database or IBMLINK for:
 - Available APARs to resolve known issues
 - Suggested recovery actions.

Recovery actions

- Implement the suggested recovery action or actions and schedule the install of all maintenance that was identified in your problem search. Note, that the implementing all PDSE maintenance requires an IPL with CLPA.
- Gather the standard set of PDSE diagnostic data and contact the Support Center if additional assistance is required.

Best practices

MSGIGW038A possible PDSE problems

PDSE employs an optional monitor that checks PDSE resources on regular intervals for possible problems. The monitor reports potential PDSE problems in an IGW038A messages.

Note: The IGW038A message only warns of a possible problem. Further actions must be taken to determine if a problem actually exists.

A PDSE ANALYSIS command should be issued as soon as possible after the receipt of an IGW038A message. The IGW038A message is issued from either the SMSPDSE or SMSPDSE1 address space. The issuing address space is identified in the first line of the message:

```
IGW038A POSSIBLE PDSE PROBLEM(S). (SMSPDSE|SMSPDSE1)
RECOMMEND ISSUING V SMS,PDSE,ANALYSIS
```

The recommended action is to issue a PDSE ANALYSIS command. IBM strongly suggests that clients employ their system automation package to respond to IGW038A messages.

Symptoms

IGW038A messages

How to investigate

1. The target address space for the PDSE ANALYSIS command depends on the address space that issued the IGW038A message:

- ANALYSIS command for IGW038A issued from the SMSPDSE address space:

```
VARY SMS,PDSE,ANALYSIS
```

- ANALYSIS command for IGW038A issued from the SMSPDSE1 address space:

```
VARY SMS,PDSE1,ANALYSIS
```

Route the command to all LPARs in the SYSPLEX.

2. The VARY ANALYSIS command output is displayed in an IGW031I message. This message identifies any latch/lock contention problems, latch/lock hangs, and outstanding cross system messages. (PDSE uses XCF messaging for serialization between systems in a SYSPLEX.) Refer to the manual *z/OS DFSMSdfp Diagnosis*, Chapter "PDSE Diagnostic Aids", section "Analysis Command", subheading "Recommended Usage", for a description of the IGW031I command output.

Note that a single IGW038A message is not an indication of a problem, especially if the IGW031I message states:

```
++ No exceptional data set conditions detected
```

In all likelihood, this message indicates that the detected exception condition was temporary in nature. Proceed to the next step if IGW038A messages continue to occur.

PDSE problem determination

3. Parse the IGW031I message for an indication of cross system issues. Make note of any remote subsystem IDs (ssssssss)

```
++ Message to ssssssss pending for nnnnnnnn seconds
```

4. Obtain a dump of the identified PDSE operating environment. The recommended method for obtaining a dump, in this situation, is to use the Monitor's DUMPNEXT processing option. This dump will automatically include both the SMSPDSE or SMSPDSE1 server and the client address space.

The DUMPNEXT processing option should be enabled on the LPAR where the IGW038A message was issued and any remote systems identified in the IGW031I message.

Issue the following command to enabling DUMPNEXT processing for the SMSPDSE address space:

```
VARY SMS,PDSE,MONITOR,DUMPNEXT
```

Issue the following command to enabling DUMPNEXT processing for the SMSPDSE1 address space:

```
VARY SMS,PDSE1,MONITOR,DUMPNEXT
```

Recovery actions

1. After obtaining the dump or dumps CANCEL or CANCEL and FORCE, if necessary, any TSO users or jobs identified in the IGW031I message. In most instances the termination of the address space that is holding the contested PDSE resource will clear the latch.

Note: In most instances, ASRBULCH latches are secondary issues and not the primarily contested PDSE resource.

2. Reissue the ANALYSIS command after terminating job or jobs. If the IGW031I shows the same contested latch then the FREELATCH command should be issued. Again the format of the FREELATCH command is dependent on the address space where the resource issue was detected.

```
V SMS,PDSE(1),FREELATCH(latchaddr,ASID,tcbaddr)
```

where latchaddr, ASID, and tcbaddr are from the IGW031I output.

Note: The FREELATCH command can only be used to free latches. This command will have no effect against PDSE data set locks.

Important: If this command is used to release a latch held by a process (on the same system or on another system in a multisystem environment) that is still running, it could result in breaking the PDSE and it could cause system abends.

3. Run the ANALYSIS command again after issuing the FREELATCH command. If problems still exist, there is a good possibility that an SMSPDSE1 address space restart or an IPL will be required to clear the problem.

Consideration should be given to engaging the service center at this point.

Important: All support releases of PDSE have the option of running with a re-startable SMSPDSE1 address space. A restart of the SMSPDSE1 address space should be attempted prior to an IPL.

Refer to the manual *DFSMSdfp Diagnosis Reference*, chapter "PDSE Diagnostics Aids", for a discussion of consideration for restarting the SMSPDSE1 address space.

Issue the following command if it been determined that the PDSE1 address can and should be restarted:

```
V SMS,PDSE1,RESTART
```

4. An IPL will be required if:

- The restart of the SMSPDSE1 fails to free the contested PDSE resource
- If the SMSPDSE1 address space is not activated.

Consideration should be given to engaging the support center before an IPL is attempted as the support center may be able to offer other recovery options.

5. Open an incident with the support center, COMPID 5695DF115, after all recovery actions have been completed or if attempted recovery actions fail to resolve the PDSE resource contention issue. Submit the following diagnostic data to the support center:
 - All dumps that were captured
 - Detailed EREP report or raw LOGREC data that brackets the dumps by two hours. Raw LOGREC data is preferred format.
 - SYSLOG or OPERLOG for the same time frame as LOGREC/EREP.

Best practices

PDSE data set corruption

Symptoms

PDSE corruption is normally discovered through an operator message or an ABEND0F4 failure.

Common operator messages which indicate possible PDSE data set corruption include:

- IEC036I 002-A4 IGC0005E
- IEC143I 213-50
- IGW01330T DIAG 271B0409, 281C014A, 270C03F7,
- CSV031I LIBRARY SEARCH FAILED FOR MODULE xxxxxxx, 27080409
- 'I/O error reading data' when using ISPF Browse

Common ABEND0F4 reason codes that may indicate PDSE corruption are:

Reason Code	Module	Explanation
010E5AB8	IGWDARD1	
IGWDADC0	JCDM_NOT_ALL_SARS_RETURNED	
050A0046	IGWBEXT2	BDS_INVALID_DSCB_EXTENTS
06215910	IGWBVLP1	JCDM_BAD_VDF
070A0021	IGWBVLP1	CDS_INVALID_LSCB
130B0203	IGWLH3AB	CLM_FIND_HL3B_HASH_FAILED
1451A030	IGWDDCR2	IMF_DUPLICATE
145AA033	IGWISRCH	IMF_INDEX_PAGE_DAMAGED
1465A03E	IGWISUBS	IMF_ABEND_DURING_PROC
150A001E	IGWBIX1	INVALID_DIRECTORY_SUFFIX
150BC008	IGWBIEX1	RSNS_IO_ERROR_DIRECTORY
150BC009	IGWBIEX1	RSNS_IO_ERROR_MEMBER
1C0752EE	IGWDRRRC	JCDM_RECORD_PREFIX_ERROR
270C03F7	IGWC DGTA	DESR_S_IEWBX_ILO_ERROR
271B0409	IGWC DGTL	DESR_S_UNEXPECTED_JCDM_ERROR

PDSE corruption can be caused by improper sharing or serialization of the PDSE resources in the user's processing environment. Most PDSE sharing problems violate one of the following rules:

- GRS must be active on any system that shares PDSEs with another system.
- In a PDSESHARING(NORMAL) environment all sharing PDSE must be in the same GRSPLEX.

PDSE problem determination

- In PDSESHARING(EXTENDED) environment all sharing systems PDSE must be in the same SYSPLEX. (Extended PDSE sharing requires XCF signaling between the sharing systems and XCF signaling is restricted to systems within a single SYSPLEX.)

How to investigate

Identify the problem PDSE data set.

1. The data set name will appear in the following operator messages:

- IEC036I
- IEC143I
- CSV031I – will contain the failing DDNAME

Proceed to the next step if PDSE data set name cannot be identified from an operator message.

2. Format the dump and issue the following command:

```
IPCS STATUS
```

Some of the PDSE subcomponents will record the PDSE data set name in the SDWA Variable Recording Area of this report. Proceed to the next step if the data set name was not recorded in the SDWA Variable Recording Area.

3. A more detailed analysis of the ABEND0F4 dump will be needed if a PDSE data set is not identified by either of the previous two steps. Open an incident with the service center and transfer the ABEND0F4 dump in for assistance in identifying the failing PDSE.

Recovery actions

1. Skip this step if the LPAR where the failure is occurring is a monoplex. Otherwise, attempt an ISPF browse against the identified PDSE data set from another LPAR in the SYSPLEX. Proceed to step 4 if the browse fails. If the ISPF browse from another LPAR in the SYSPLEX succeeds then the DASD copy of the file is intact and the most likely problem is that the cached directory pages for the PDSE data set are out of sync with the DASD copy of the file. Proceed to step 3 if the failing LPAR is not release at z/OS 1.13 or above.
2. z/OS 1.13 and above. Issue the following command to discard the cached directory pages and reread the directory from DASD VARY SMS,PDSE|PDSE1,REFRESH,DSN(dsname)[,VOL(volser)] Attempt to browse the identified PDSE data set on the failing LPAR after the command completes. Proceed to the next step if the failure still occurs.
3. If implemented, restart the SMSPDSE1 address space.

```
VARY SMS,PDSE1,RESTART
```

Restarting the PDSE address space will rebuild the connections to all opened PDSE data sets. Proceed to the next step if the SMSPDSE1 address space is not implemented or if the file still cannot be browsed after the restart completes.

4. Open an incident with the service center. Gather and transfer in for review the previously defined basic PDSE diagnostic data.
5. The PDSE data set is corrupted, begin data set recovery efforts.
6. IEBCOPY the corrupted data set into a newly allocated PDSE data set. Proceed to the next step if the IEBCOPY job fails.
7. If the IEBCOPY fails, then the data set will have to be recovered from a viable back up. Open an incident with the service center and gather and transfer in for review the previously defined basic PDSE diagnostic data. In addition to the basic PDSE diagnostic data, collect and transfer in the following additional data:
 - A DFSS physical dump copy of the corrupted PDSE data set.

- An IEHLIST LISTVTOC FORMAT report or a DFDSS PRINT of the VTOC from the volume where the PDSE resides.
- SMF TYPE14/15 records from all sharing LPARs back to the last successful access.

Best practices

Failure of the SMSPDSE or SMSPDSE1 address space

Symptoms

The deactivation of the SMSPDSE is indicated by an IGW007E message.

```
IGW007E SMS HAS BEEN DEACTIVATED DUE TO A FATAL ERROR RETURN CODE (IN
HEX): return-code REASON CODE (IN HEX): reason-code {A DUMP HAS BEEN
TAKEN|NO DUMP HAS BEEN TAKEN}
```

The deactivation of the SMSPDSE1 address space is indicated by an IGW077E message.

```
IGW077E SMSPDSE1 DEACTIVATED DUE TO A FATAL ERROR RETURN CODE:
return-code REASON CODE: reason-code A RESTART OF SMSPDSE1 IS
RECOMMENDED. TO RESTART ENTER: VARY SMS,PDSE1,RESTART
```

External symptoms include, but are not limited to, delays in job initialization and job termination, delays in job processing, and S30D abends.

How to investigate

- The deactivation of the SMSPDSE is indicated by an IGW007E message.
- The deactivation of the SMSPDSE1 is indicated by an IGW077E message.

These messages are issued as a highlighted operator action message that requires immediate attention. Both messages will remain on the console until recovery action is taken.

In the case of the deactivation of the SMSPDSE1 address space PDSE processing will be held or corralled until such point that the SMSPDSE1 address space is restarted. This corraling of PDSE activity will have an adverse affect on LPAR where the error occurred.

Recovery actions

1. Proceed to step 2 if the SMSPDSE address space was deactivated. Continue with this step if the SMSPDSE1 address space was deactivated. Attempt to restart the SMSPDSE1 address space by issuing the following command

```
VARY SMS,PDSE1,RESTART
```

A restart of the SMSPDSE1 address space is almost always successful. Proceed to the next step in the event that the address space fails to restart.

2. Obtain a stand alone dump.
3. IPL the LPAR.
4. Open an incident with the service center and gather and transfer in for review the SADUMP and previously defined basic PDSE diagnostic data.

Best practices

- Enable the SMSPDSE1 address space when executing in a PDSESHARING(EXTENDED) environment.
- Run with the SMSPDSE and SMSPDSE1 Monitor task enabled.
- Automate the response to the IGW038A message.
- Conform to PDSE sharing rules. See [z/OS DFSMS Using Data Sets](#) for a summary of PDSE sharing rules.

Chapter 20. RRS operational problem determination

Resource Recovery Service (RRS) is a syncpoint manager that consists of protocols and program interfaces that, when requested, can coordinate consistent changes to one or more protected resources. These resources can be accessed through different resource managers and can be on different systems within a sysplex. For more information, see [z/OS MVS Programming: Resource Recovery](#).

Basic RRS problem determination functions

RRS provides the following functions to help with problem determination, data collection, and recovery:

- Use the DISPLAY RRS command to query status of resource managers and units of recovery (UR). Units of recovery represent transactions.
- Use CTRACE options specific to SYSRRS to gather detailed information about RRS requests and events.
- Use the SETRRS command to cancel or shut down RRS, and enable and disable archive logs. RRS issues messages with the ATR prefix.
- Make sure that all RRS checks are activated and running to warn you of impending storage usage and other RRS problems. For details, see [RRS checks \(IBMRRS\)](#) in [IBM Health Checker for z/OS User's Guide](#).

Collecting documentation for RRS

This section contains the following topics:

- [“Dumping RRS information” on page 285](#)
- [“Important RRS CTRACE information” on page 286](#)

Dumping RRS information

When the IBM Support Center requests an RRS dump, they typically ask for the same data every time. To save time, create two IEADMCxx parmlib members.

Note: If your installation setup RRS with a job name other than 'RRS', replace that value in both the JOBNAME and DSPNAME parameters in the following examples:

- To dump RRS: Create an IEADMCxx parmlib member with the requested RRS dump parameters as follows:

```
JOBNAME=(RRS), DSPNAME=('RRS'.*),
SDATA=(ALLNUC, LPA, LSQA, PSA, RGN, SQA, TRT, CSA, GRSQ, COUPLE)
```

When you need a dump of RRS, use the following command:

```
DUMP COMM=(title), PARMLIB=xx
```

xx is the suffix of the IEADMCxx parmlib member that contains the RRS dump options.

- **To dump RRS and system logger:** Create an IEADMCxx parmlib member with the requested RRS and system logger dump parameters as follows:

```
JOBNAME=(RRS, IXGLOGR), DSPNAME=('RRS'.*, 'IXGLOGR'.*),
SDATA=(ALLNUC, LPA, LSQA, PSA, RGN, SQA, TRT, CSA, GRSQ, COUPLE)
```

Then, when you need a dump of RRS and system logger, use the following command:

```
DUMP COMM=(title), PARMLIB=xx
```

RRS problem determination

xx is the suffix of the IEADMCxx parmlib member that contains the RRS and system logger dump options.

If exploiters are involved, you can include them in the JOBNAME list in your IEADMCxx parmlib members.

Important RRS CTRACE information

RRS strongly suggests that you have the following CTRACE options, as found in the RRS CTRACE samplib member ATRTCRRS, set at all times, . This ensures you have the most trace information available in the RRS dumps. Create a CTIRRSxx parmlib member that contains the following options:

```
TRACEOPTS
ON
BUFSIZE(500M)
OPTIONS('EVENTS(URSERVS,LOGGING,CONTEXT,EXITS,STATECHG,RRSAPI,RESTART)')
```

To ensure that these CTRACE settings are always active:

- Issue the S RRS command, specifying the CTIRRSxx parmlib member.
- Update the RRS procedure with the CTMEM value to ensure that RRS CTRACE is running on each startup.

If RRS is already started, to set these options, issue the following z/OS command:

```
TRACE CT,ON,COMP=SYSRRS,PARM=CTIRRSxx
```

RRS CTRACE in a sysplex cascaded transaction environment and problem determination mode: If you are running a sysplex cascaded transaction environment and are in problem determination mode, use the external writer to trace more data. In addition, update the RRS trace to use the following option:

```
OPTIONS('EVENTS(ALL)')
```

You are in problem determination mode when you are collecting documentation for a recurring problem, such as a transaction hang or delay.

RRS recovery options

This section contains the following topics:

- [“RRS warm start” on page 286](#)
- [“RRS cold start” on page 287](#)

RRS warm start

An RRS warm start is also named recycling RRS. For an RRS warm start:

1. Whenever doing a warm start of RRS, first shut down its exploiters whenever possible, as documented in *z/OS MVS Programming: Resource Recovery*. For example, WebSphere® does not handle RRS termination and terminates itself when it detects RRS is not available. Additional exploiters might encounter abends when RRS terminates if they had outstanding processing activity at that time. IMS can terminate for some abend conditions.
2. Terminate RRS with one of the following commands:
 - SETRRS CANCEL, which terminates RRS with an abend code X'222'.
 - SETRRS SHUTDOWN, which first stops each resource manager and then terminates RRS without an abend (normal termination).
3. Restart RRS with the Automatic Restart Manager (ARM) or by issuing the START RRS command.

RRS cold start

A cold start of RRS requires deletion and redefinition of the RRS RM.DATA log stream with the IXCMIAPU utility. You can find sample JCL for an RRS cold start in the ATRCOLD member of SYS1.SAMPLIB.

When unrecoverable corruption is detected in either the RRS RM.DATA or RESTART log streams, process an RRS cold start. For instructions about RRS cold start processing, see [Internal cold start in z/OS MVS Programming: Resource Recovery](#).

RRS component-specific problems and recovery

- [“RRS resource contention” on page 287](#)
- [“RRS suspended, waiting for signal from system logger” on page 288](#)
- [“RRS log stream gap condition” on page 289](#)
- [“RRS log stream data loss condition” on page 290](#)
- [“RRS high processor usage” on page 293](#)
- [“RRS address space hang” on page 293](#)
- [“RRS high storage usage” on page 293](#)
- [“Resource manager is unable to start with RRS” on page 295](#)
- [“Resource manager termination delay” on page 295](#)
- [“RRS transaction hang” on page 296](#)
- [“RRS severe error on RRS RM.DATA log stream, message ATR250E” on page 297](#)
- [“Resolving RRS problems in a sysplex cascaded transaction environment” on page 297](#)
 - [“Sysplex cascaded transaction hang” on page 298](#)
 - [“Sysplex cascaded transaction hang messages ATR246I and ATR247E” on page 299](#)

RRS resource contention

Symptoms

RRS resource contention shows up in DISPLAY GRS command output as contention for resource SYSZATR *sysplex_name*-RESTART. RRS obtains this global resource when manipulating the ATR.*logging_group_name*.RM.DATA log stream. Operations such as RRS start/stop, resource manager start/unset, and internal log stream management obtain this resource exclusively. If any of these operations hang while holding this resource, all other operations in the sysplex that require this resource also hang.

How to investigate

Look for these symptoms:

- Message ATR248E or ATR249E to identify a potential hang in log stream processing. See [“RRS suspended, waiting for signal from system logger” on page 288](#).
- ATR* and IXG* messages that involve the resource manager data log stream.

Recovery actions

Periodically monitor for SYSZATR RESTART resource contention with the DISPLAY GRS,CONTENTION command. If you detect contention, Follow the action specified in the system logger service return and reason codes shown in message ATR248E or ATR249E. See the IXGWRITE return and reason codes in [z/OS MVS Programming: Authorized Assembler Services Reference EDT-IXG](#).

If the action in the return and reason code does not resolve the contention, or if you cannot detect the cause of the contention:

1. Capture a dump of RRS, system logger, and GRS on the system that is holding the resource.

RRS problem determination

2. Recycle RRS on the holding system to free the resource.

Note: If the issue that causes the holder to hang persists, the resource contention can continue until the issue is resolved.

Search problem reporting databases for a fix for the problem. If no fix exists, report the problem to the IBM Support Center and provide the dump, system logs, and formatted LOGREC data set. (A symptom record is created on the system experiencing the failure and stored in the LOGREC data set.)

Best practices

Periodically monitor for SYSZATR RESTART resource contention with DISPLAY GRS,CONTENTION.

RRS suspended, waiting for signal from system logger

RRS uses system logger services to manage resource manager and unit of recovery (UR) status information. Sometimes RRS ends up suspended while waiting for a signal from system logger.

Symptoms

The system issues one of the following messages when an RRS function is suspended in Logger processing:

- ATR248E RRS IS WAITING FOR SIGNAL FROM LOGGER TO RESUME PROCESSING
RETURN CODE: *returncode* REASON CODE: *reasoncode* DIAGNOSTIC
INFORMATION: *diag1 diag2 diag3 diag4*
- ATR249E RRS IS WAITING FOR SIGNAL FROM LOGGER TO RESUME PROCESSING
LOGSTREAM NAME: *logstreamname* RETURN: *returncode* REASON: *reasoncode*
DIAGNOSTIC INFORMATION: *diag1 diag2 diag3 diag4*

The messages indicate that RRS has a suspended TCB waiting for a signal from system logger before it can resume processing. The TCB is suspended for at least 30 seconds. The problem occurs when RRS is in the process of updating an RRS log stream and encounters a temporary error condition. The error condition is identified by the return and reason codes from the IXGWRITE service shown in the messages.

Generally, RRS is holding the SYSZATR RESTART global resource when suspended. The system obtains this resource when RRS updates the ATR.logging_group_name.RM.DATA log stream, which it can do during:

- Internal RRS housekeeping
- Restart or unset processing

All processing that involves the RM.DATA log stream throughout the RRS logging group (generally the sysplex) is suspended while waiting for the signal from system logger.

How to investigate

To investigate:

- Review the meaning of the return and reason codes found in message ATR248E or ATR249E, which are found in the topic about IXGWRITE - Write log data to a log stream in [z/OS MVS Programming: Authorized Assembler Services Reference EDT-IXG](#).
- Review SYSLOG for prior messages that show an issue with system logger processing. For example, look for IXG messages that involve the RM.DATA log stream.

Recovery actions

Check for and correct problems with system logger by reviewing the reported diagnostic information. If the problem cannot be identified or resolved, capture a dump of RRS and Logger to provide to the IBM Support Center. Provide the dump along with SYSLOG and formatted LOGREC data set.

If additional RRS work is suspended, as indicated by SYSZATR RESTART global resource contention, you might need to recycle RRS. Recycle on the system that is receiving the ATR248E message or ATR249E

message. When the temporary error condition reported by system logger remains persistent, the problem reappears until you resolve the log stream issue.

RRS log stream gap condition

RRS uses system logger to manage resource manager and UR status information. RRS encountered a gap condition when browsing the log stream, which can lead to a loss of data.

Symptoms

The following message is issued when the system detects a log stream gap condition.

```
ATR202D GAP FOUND IN logstreamname. REPLY RETRY TO RETRY OR ACCEPT TO
ACCEPT THE DATA LOSS
```

When the gap condition is against the RRS RM.DATA log stream, the following message may also be issued.

```
ATR212I RRS DETECTED LOG DATA LOSS ON LOGSTREAM logstreamname
DUE TO INACCESSIBLE LOG DATA. LOG DATA FROM lowGMT TO
highGMT ARE AFFECTED.
```

See [“RRS severe error on RRS RM.DATA log stream, message ATR250E”](#) on page 297 for more information.

How to investigate

To investigate:

- Review SYSLOG for prior messages that report the gap condition.
- Sometimes, when an HSM backup occurs against the volume that contains the RRS log stream data set, it prevents system logger from opening the data set. This condition can result in a gap condition for the log stream. To determine whether the gap condition is the cause, look for the following accompanying messages:

```
ARC0722I BACKUP STARTING ON VOLUME vvvvv(SMS) AT
ARC0722I (CONT.) hh:mm:ss ON yyyy/mm/dd SYSTEM ssss

IEC161I 052(015)-084,jjj,sss,,
IEC161I hlq.ATR.logging_group_name.logstream.Axxxxxxx,
IEC161I hlq.ATR.logging_group_name.logstream.Axxxxxxx,.DATA,ICFCAT.SYSCAT03

IXG268I LOGSTREAM DATASET hlq.ATR.logging_group_name.logstream.Axxxxxxx
CAN NOT BE OPENED FOR JOB RRS DUE TO
INCORRECT VSAM SHAREOPTIONS OR OTHER ERROR,
REQUESTED DATA MAY NOT BE AVAILABLE.

ATR202D GAP FOUND IN ATR.logging_group_name.logstream.
REPLY RETRY OR ACCEPT TO ACCEPT THE DATA LOSS

ARC0723I BACKUP ENDING ON VOLUME vvvvv AT hh:mm:ss
```

Recovery actions

- If you can ascertain the cause of the gap, correct the problem and reply RETRY to the message. If the gap is caused by HSM backup activity, reply RETRY to message ATR202D after you get accompanying message ARC0723I.

Note: Replying ACCEPT to message ATR202D hardens the gap as a data loss to RRS. Internal flags are set that cannot be updated without a warm start.

When the gap condition is against the RRS RM.DATA or RESTART log stream and:

1. Can be corrected, use an RRS warm start to reset the internal flag settings.
 2. Cannot be corrected, use an RRS cold start to recover.
- If you cannot find the cause of the gap condition, capture a dump of RRS and system logger to provide to the IBM Support Center. Provide the dump along with SYSLOG and formatted LOGREC data.

RRS log stream data loss condition

RRS uses system logger to manage resource manager (RM) and unit of recovery (UR) status information. An RRS log stream data loss condition can occur when RRS does not disconnect cleanly from system logger before a power-on-reset (POR) and IPL. Any exploiter of system logger could be exposed if the system is brought down hard before the connections to the logstream are broken. It is important to disconnect cleanly when the exploiter is brought down so that system logger knows to offload the data to DASD.

Symptoms

- A power-on-reset (POR) and IPL of the MVS images using a log stream result in a data loss condition, RC08 RSN84D.
- The system issues the ATR210E message when it detects a log stream data loss condition when RRS attempted to browse its RM.DATA log stream.

```
ATR210E INACCESSIBLE LOG DATA DETECTED ON THE RRS RM DATA LOGSTREAM logstreamname
```

The system can also issue accompanying message ATR218I to describe an RRS process that failed because of the inaccessible log data.

When the data loss condition is against the RRS RM.DATA log stream, the following message may also be issued.

```
ATR212I RRS DETECTED LOG DATA LOSS ON LOGSTREAM logstreamname  
DUE TO INACCESSIBLE LOG DATA. LOG DATA FROM lowGMT TO  
highGMT ARE AFFECTED.
```

See [“RRS severe error on RRS RM.DATA log stream, message ATR250E” on page 297](#) for more information.

How to investigate

Review SYSLOG for prior messages from system logger that reports the data loss condition.

Recovery actions

Action to take when RRS suffers loss of data:

- If a loss of data was detected against the RM.DATA log stream, a cold start is required for RRS to successfully initialize.
- If a loss of data was detected against the RESTART log stream, RRS has already successfully initialized but any resource manager attempting to RESTART with RRS will fail with messages ATR212I and ATR209I . In order to allow for these resource managers to restart with RRS, you must perform an RRS cold start.

Note: It is strongly recommended that you use unconditional logstream duplexing for the RM.DATA and RESTART logs since any loss of data, unresolved gap, or any permanent error will force an RRS cold start.

If you cannot determine the cause of the data loss condition, capture a dump of RRS and system logger to provide to the IBM Support Center. Provide the dump along with SYSLOG and formatted LOGREC data.

Best practices

To avoid data loss conditions, do the following:

- Use logstream duplexing for both the RM.DATA and RESTART log streams to protect against data loss.
- Do the following before issuing a POR of a coupling facility to avoid data loss:
 1. To make sure that RRS has disconnected from its log streams, issue the following command:

```
DISPLAY LOGGER,L
```

This command displays the sysplex view of the system system logger connections, such as which systems are connected to which log streams, including the log stream name, structure name, number of connections, current status, and the name of the connected systems that are connected and whether they are using staging data sets.

The system issues message IXG601I to display the results of the DISPLAY LOGGER,L command. If an RRS log stream is still connected, the output shows INUSE for that log stream, for example:

```

SY2 IXG601I 09.00.21 LOGGER DISPLAY 662
INVENTORY INFORMATION BY log stream
LOGSTREAM STRUCTURE #CONN STATUS
-----
ATR.SYSPLEX.DELAYED.UR SYSPLEX_DELAY 000005 LOSS OF DATA
SYSNAME: SY1
STG DS: NO
SYSNAME: SY2
STG DS: NO
ATR.SYSPLEX.MAIN.UR SYSPLEX_MAIN 000005 IN USE
SYSNAME: SY1
STG DS: NO
SYSNAME: SY2
STG DS: NO
ATR.SYSPLEX.RETART SYSPLEX_RESTRT 000005 LOSS OF DATA
SYSNAME: SY1
STG DS: NO
SYSNAME: SY2
STG DS: NO
ATR.SYSPLEX.RM.DATA SYSPLEX_RMDATA 000005 IN USE
SYSNAME: SY1
STG DS: NO
SYSNAME: SY2
STG DS: NO

```

- For log streams that are displayed as INUSE in the output from DISPLAY LOGGER,L, issue the following command to stop RRS and have it break its connection from all its log streams, signaling system logger to move all the data from the CF to DASD:

```
SETRRS CANCEL
```

If a SETRRS CANCEL is not effective, use a FORCE command.

- If after issuing SETRRS CANCEL, the log stream still shows connections when you think there should not be any, use the following commands to help determine why system system logger is still connected:
 - DISPLAY LOGGER,C,LSN=LOGSTREAMNAME,DETAIL - Issue this command from the system that is still connected to show any application that is still connected. If no application is showing as connected, the log stream is most likely in a failed state and will need recovery. Reconnecting to the log stream will drive this recovery. Note that if you are not using staging data sets to duplex data, this reconnection must be done prior to POR of the CF or IPLing the MVS image.
 - DISPLAY XCF,STR,STRNAME=STRUCTURENAME - Issue this command to display a XES view of who is connected to a given structure, but will tell nothing of the log streams contained within that structure. If the output of the command shows STATUS: NOT ALLOCATED, there are no log streams connected in that structure, but note there could still be problems with any of those log streams.

The following shows output from a D XCF,STR,STRNAME=STRUCTURENAME command - note that in this instance it shows a log stream which is still allocated:

```

D XCF,STR,STRNAME=LIST14
.
IXC360I 09.12.15 DISPLAY XCF
STRNAME: LIST14
STATUS: ALLOCATED
POLICY SIZE : 4000 K
POLICY INITSIZE: N/A
REBUILD PERCENT: N/A
DUPLEX : DISABLED
PREFERENCE LIST: CF01 CF01N CF02 CF02N
ENFORCEORDER : NO
EXCLUSION LIST : LIST01 LIST02 LIST03

```

```
LIST04
ACTIVE STRUCTURE
-----
ALLOCATION TIME: 02/24/2000 09:07:45
CFNAME : CF01N
COUPLING FACILITY: SIMDEV.IBM.EN.CF0100000000
PARTITION: 0 CPCID: 00
ACTUAL SIZE : 4096 K
STORAGE INCREMENT SIZE: 256 K
PHYSICAL VERSION: B3A5BA73 B0687909
LOGICAL VERSION: B3A5BA73 B0687909
SYSTEM-MANAGED PROCESS LEVEL: NOT APPLICABLE
DISPOSITION : DELETE
ACCESS TIME : 0
MAX CONNECTIONS: 16
# CONNECTIONS : 2
CONNECTION NAME ID VERSION SYSNAME JOBNAME ASID STATE
-----
IXGLOGR_SY1 01 00010001 SY1 IXGLOGR 0014 ACTIVE
IXGLOGR_SY2 02 00020001 SY2 IXGLOGR 0014 ACTIVE
```

To reiterate, it is only safe to POR a CF when:

1. The output of the D LOGGER, L command shows zero connections and status available for all log streams, which means that the application (RRS, IMS, or CICS, for example) has disconnected cleanly.
2. The output of D XCF, STR, STRNAME issued for **every existing structure** shows STATUS NOT ALLOCATED for the log stream.

If these conditions are not satisfied, you are guaranteed to lose data.

Factors decreasing the chances of a RRS log stream data loss condition

The possibility of getting a data loss condition (RC08 RSN84D) is related to a number of different factors and the interrelationship of these factors to each other. The following factors decrease your chances of an RRS log stream data loss condition:

- A non-volatile CF. Your CF is non-volatile if has a back up battery or alternate power source. A volatile CF, which has no alternate power source, is more vulnerable to data loss conditions.
- Failure independent systems. A failure independent system is one that does not share a power source with the CF. A failure dependent system, on the other hand, does share a power source with the CF, which means that a power failure effects both system and CF, making it more vulnerable to a data lost condition.
- Log streams duplexed with staging data sets, as follows:
 - Unconditional duplexing: When you specify STG_DUPLEX(YES) with DUPLEXMODE(UNCOND) in the log stream definition in the LOGR policy, system logger writes each transaction to a staging data set on DASD every time a transaction is written to the CF.
Note: It is strongly recommended that you use unconditional duplexing for the RM.DATA and RESTART logs since any loss of data, unresolved gap, or any permanent error will force an RRS cold start.
 - Conditional duplexing: When you specify STG_DUPLEX(YES) with DUPLEXMODE(COND), system logger will check for either condition, VOLATILITY or FAILURE DEPENDENCE. If either is true, then system logger writes each transaction to a staging data set on DASD, as in unconditional duplexing.

Without duplexing (STG_DUPLEX(NO) specified on the log stream definition), system logger writes each transaction to interim storage located in a system logger dataspace every time a transaction is written to the CF.

How factors interact:

- A combination of failure independence, non-volatile CF, and log streams duplexed to staging data sets decreases your chances of a data loss condition dramatically.

Note that if your system or systems are failure independent and the CF is non-volatile, then you do not actually need to duplex to staging data sets to limit exposure to the data loss condition.

- If you do not duplex to staging data sets and have one other factor making you more vulnerable to data loss, you could be exposed to data loss if both system and CF go down or are shut down **and** any exploiters of LOGGER still have connections to their log streams. In other words, you are exposed to data loss when you do not have duplexing and either:
 - Your system or systems are failure dependent and the CF is non-volatile.
 - Your system or systems are failure independent but the CF is volatile.

For more information, see the chapter on system logger in [z/OS MVS Setting Up a Sysplex](#).

RRS high processor usage

Symptoms

You detect High processor usage for RRS.

Recovery actions

High processor usage detected for RRS processing suggests a loop. Collect a dump of RRS. If this condition is a sysplex cascaded transaction processing environment, collect dumps on each system in that environment. Recycle RRS on the system with high processor usages. Search the problem reporting data bases for a fix for the problem. If no fix exists, report the problem to the IBM Support Center, providing the dump or dumps along with SYSLOG and formatted LOGREC data.

RRS address space hang

Symptoms

RRS is not processing work.

How to investigate

Look for exception messages for RRS check RRS_Storage_NumServerReqs. For information about the check, see RRS_Storage_NumServerReqs in *IBM Health Checker for z/OS User's Guide*. This check monitors the number of server requests within RRS. The system issues exception message ATRH016E if the number of server requests exceeds the threshold monitored:

```
ATRH016E The current number of server task requests in RRS is
         curreqs which exceeds the threshold
```

If this number continues to grow, RRS might have a hang. RRS has a limited number of TCBS to process each server request.

You can also observe transaction hangs by looking at the output of the following DISPLAY command:

```
D RRS,UREX
```

Recovery actions

If RRS_Storage_NumServerReqs check exception message ATRH016E shows the number of server requests growing, check for persistent RRS resource contention. Capture a dump of RRS and the holder of RRS resources, if applicable, and provide them to IBM Support Center. Recycle RRS to address the suspected hang processing server requests.

RRS high storage usage

Symptoms: High storage usage detected in the RRS address space.

How to investigate: Look for RRS IBM Health Checker for z/OS check exception messages that warn of growth in RRS storage use. See [RRS checks \(IBMRRS\)](#) in *IBM Health Checker for z/OS User's Guide* for a description of RRS checks:

RRS problem determination

- RRS_Storage_NumLargeLOGBlks - Monitors the count of large log buffer blocks in RRS. The system issues exception message ATRH020E if the count exceeds the threshold monitored, which can affect RRS private storage. (The default threshold is 1,000.)
- RRS_Storage_NumLargeMSGBlks - Monitors the count of large message buffer blocks in RRS. The system issues exception message ATRH018E if the count exceeds the threshold monitored, which can affect RRS private storage. (The default threshold is 1,000.)
- RRS_Storage_NumTransBlks - Monitors the count of active URs (transactions) with RRS. The system issues exception message ATRH014E if the count exceeds the threshold monitored, which can impact common storage usage. (The default threshold is 10,000.) In addition to the UR control block created in private storage, RRS creates additional control block data in common storage.

How to investigate

Look for RRS IBM Health Checker for z/OS check exception messages that warn of growth in RRS storage use. See [RRS checks \(IBMRRS\)](#) in *IBM Health Checker for z/OS User's Guide* for a description of RRS checks:

- RRS_Storage_NumLargeLOGBlks - Monitors the count of large log buffer blocks in RRS. The system issues exception message ATRH020E if the count exceeds the threshold monitored, which can affect RRS private storage. (The default threshold is 1,000.)
- RRS_Storage_NumLargeMSGBlks - Monitors the count of large message buffer blocks in RRS. The system issues exception message ATRH018E if the count exceeds the threshold monitored, which can affect RRS private storage. (The default threshold is 1,000.)
- RRS_Storage_NumTransBlks - Monitors the count of active URs (transactions) with RRS. The system issues exception message ATRH014E if the count exceeds the threshold monitored, which can impact common storage usage. (The default threshold is 10,000.) In addition to the UR control block created in private storage, RRS creates additional control block data in common storage.

Recovery actions

If the amount of storage that is being monitored continues to grow, determine whether it is expected behavior. Use the available RRS data collection techniques such as panels, console display command, or batch program. These techniques can help you assess the level of transaction activity in RRS and determine if it is unusual or unexpected. The threshold values can be updated for each IBM Health Checker for z/OS check to values more appropriate to the environment. Follow the action based on:

- RRS_Storage_NumLargeMSGBlks exception messages:
 - Look for hangs with sysplex cascaded transactions and base actions on the hang reported
 - Use the DISPLAY RRS,UREX command to list all active units of recovery (UR) waiting for an event.
 - If possible, resolve the hang. Otherwise, capture a dump of RRS and the associated resource manager to provide to the IBM Support Center for further assistance.
- RRS_Storage_NumTransBlks exception messages:
 - Enter the DISPLAY RRS, UR command to list all active UR.
 - Investigate which application is creating these URs and check for a problem with that application, such as looping, not cleaning up transactions, or recursive errors.
 - If you do not understand the increase in transaction activity:
 - Collect a dump of RRS and any resource manager or work manager identified as the source of the increased transaction activity.
 - Provide the dump to the IBM Support Center for further assistance.
- For any high storage count:
 - Look for the RRS_Storage_NumServerReqs check exception message, ATRH016E. The message suggests a possible hang in RRS processing, which can cause storage usage to grow.
 - Capture a dump of RRS and provide it to the IBM Support Center for further assistance.

If you cannot find the cause of storage growth, collect a dump of RRS and any involved exploiter for analysis. Provide the dump along with SYSLOG and formatted LOGREC data set to the IBM Support Center. If the growth continues, recycle RRS to clean up the storage usage after collecting documentation to avoid reaching a critical storage shortage.

Best practices

Make sure that all RRS checks are activated and running to warn you of impending storage usage and other RRS problems.

Resource manager is unable to start with RRS

Symptoms

In a known example of this problem, when DB2 fails to restart with RRS, it issues one of the following messages:

- DSNX982I db2regionWLM attempt to perform RRS attach function failed with
RRS RC8 RSN 00F30091
- DSN3031I *DBVC DSN3RRSI RRS ATTACH INITIALIZATION FAILED, RRS CALL=CRGSEIF, RRS RETURN CODE='X'00008004'
*DSNV086E *DBVC DB2 ABNORMAL TERMINATION REASON=00F30095

How to investigate

Enter the DISPLAY RRS, RM command on each system in the sysplex to determine the state of the resource manager.

Recovery actions

Use one of the following recovery actions:

- When the resource manager is in SET state, use the RRS ISPF panels to "Unregister RM". You can enter "n" next to the RM in SET state in the resource manager list. You can also use the ATRSRV utility with REQUEST(UNREGRM) to unregister the specific RMNAME. This state might be observed when the DSN* DB2 messages are observed.
- When the resource manager is in UNSETINPROG state, check for RRS global resource contention. If observed, collect a dump of the holding address space, RRS, system logger, and GRS. Provide the dump along with SYSLOG and formatted LOGREC data to the IBM Support Center. Recycle RRS on the holding system.
- When the resource manager is in RUN, UNSET or RESET state, work with the exploiter to determine why it is failing to restart. Also determine if it involves RRS.

Resource manager termination delay

Symptoms

Resource manager termination delay can be indicated by the following WTOR messages:

- These messages can delay the termination of resource manager until they receive a reply.

```

ATR225D  CANCEL DELAYED. REPLY WAIT, BACKOUT, OR COMMIT TO RESOLVE
          INDOUBT UR. URID=uridentifier
ATR226D  MEMTERM DELAYED. REPLY WAIT, BACKOUT, OR COMMIT TO RESOLVE
          INDOUBT UR. URID=uridentifier
ATR227D  CANCEL DELAYED. REPLY WAIT, BACKOUT, OR COMMIT TO RESOLVE
          INDOUBT UR. URID=uridentifier
ATR228D  MEMTERM DELAYED. REPLY WAIT, BACKOUT, OR COMMIT TO RESOLVE
          INDOUBT UR. URID=uridentifier
ATR229D  CANCEL DELAYED. REPLY WAIT, BACKOUT, OR COMMIT TO RESOLVE
          INDOUBT UR. URID=uridentifier
ATR230D  MEMTERM DELAYED. REPLY WAIT, BACKOUT, OR COMMIT TO RESOLVE
          INDOUBT UR. URID=uridentifier
ATR231D  CANCEL DELAYED. REPLY WAIT, BACKOUT, OR COMMIT TO RESOLVE
          INDOUBT UR. URID=uridentifier
ATR232D  MEMTERM DELAYED. REPLY WAIT, BACKOUT, OR COMMIT TO RESOLVE

```

RRS problem determination

```
                INDOUBT UR. URID=uridentifier
ATR233D  CANCEL DELAYED. REPLY BACKOUT, OR COMMIT TO RESOLVE INDOUBT
                UR. URID=uridentifier
ATR234D  MEMTERM DELAYED. REPLY BACKOUT, OR COMMIT TO RESOLVE INDOUBT
                UR. URID=uridentifier
```

- Task or memory termination of the resource manager address space is delayed because of an outstanding transaction in the "in-doubt" state. RRS forces the operator to manually resolve the "in-doubt" transaction before allowing termination to continue.

How to investigate

If a resource manager address space termination is not completing, review SYSLOG for messages ATR225D - ATR234D.

Recovery actions

Reply to the outstanding WTOR to commit or back out the "in-doubt" units of recovery.

If the "in-doubt" units of recovery is not expected, enter `D RRS,UR,DETAILED,URID=uridentifier` to determine the resource managers with an interest in the units of recovery. Provide a dump of RRS and the resource managers involved in the units of recovery. Include SYSLOG and formatted LOGREC data to the IBM Support Center

Best practices

Replying to these WTORs immediately can prevent RRS caused delay to address space termination. Consider setting up automation to ensure that the messages are responded to when they are issued.

RRS transaction hang

Symptoms

The RRS exploiter reports one or more outstanding transactions exist.

How to investigate

Enter the following display command, which issues message ATR624I showing RRS units of recovery (UR) exceptions:

```
DISPLAY RRS,UREXceptions
ATR624I  hh.mm.ss
        RRS UR EXCEPTION id | SYSTEM URID | WAIT FOR |
        sysname urid waitfortext
```

Message ATR624I displays all the units of recovery that are waiting for completion of other tasks on the specified system. A UR is the RRS representation of a transaction. Examine each exception to determine whether it represents a problem.

Identify persistent exceptions in the display by looking for units of recovery that show the same status in consecutive display output. This condition suggests a possible problem with the transaction. For example, look at the following `D RRS, UREX` example output.

```
D RRS,UREX
ATR624I 14.57.34 RRS UR EXCEPTION
SYSTEM URID                                WAIT FOR
SYS1    C4D2B56B7EB296E80000432201010000
        SYS2    C4D2B56B7EB3CA5C0000728801020000 RMGR2 PREPARE Exit
SYS1    C4D2B56B7EB29A5C000003E0601010000
        SYS2    C4D2B56B7EB3D4B800006B0A01020000 RMGR2 PREPARE Exit
```

This output from a `DISPLAY` command issued from system `SYS1` lists the exception URID followed by the `WAIT FOR` text indented on the next line.

- The two `SYS1` exception units of recovery are waiting for cascaded, subordinate units of recovery on `SYS2`.

- The units of recovery on SYS2 are waiting for resource manager RMGR2 to complete its PREPARE Exit.

Recovery actions

1. Try to resolve the transaction hang.
2. If you cannot get the resource manager that is using RRS to resolve the transaction hang, capture a dump of RRS and the exploiter or exploiters involved in the hang on the systems involved.
3. Provide the dumps with SYSLOG and formatted LOGREC data on the systems involved to the IBM Support Center.

RRS severe error on RRS RM.DATA log stream, message ATR250E

RRS uses system logger to manage resource manager information. A gap or data loss condition against the RRS RM.DATA log stream could lead to a log stream error.

Symptoms

A gap or data loss condition against the RRS RM.DATA log may result in the following message being issued.

```
ATR212I RRS DETECTED LOG DATA LOSS ON LOGSTREAM logstreamname
DUE TO INACCESSIBLE LOG DATA. LOG DATA FROM lowgmt TO
highgmt ARE AFFECTED.
```

When the condition is against the RRS RM.DATA log, the following message will also be issued when all RRS systems in the sysplex are at version z/OS V2R1 or higher.

```
ATR250E RRS LOGSTREAM ERROR FOUND. CORRECT THE ERROR OR
OPTIONALLY REPLY COLDSTART TO BEGIN A RRS INTERNAL COLD
START.
```

How to investigate

Review SYSLOG for prior messages from system logger that reports the data loss condition.

Recovery actions

RRS needs to be terminated on all systems in the RRS group in order to resolve this problem. Once terminated, request a cold start of RRS using the ATRCOLD procedure and then restart RRS on each system in the RRS group. This can be done manually and will result in all outstanding transactions being lost and not recoverable. Optionally, a reply of COLDSTART may be given to instruct RRS to attempt an Internal Cold Start. RRS will remain active, but new work will not be accepted until the cold start is complete. An attempt will be made to save the in storage transactions which will be relogged as part of the Internal Cold Start procedure.

Best practices

Create an automation routine that looks for message ATR250E and take the COLDSTART action as described in [“Recovery actions” on page 297](#).

Resolving RRS problems in a sysplex cascaded transaction environment

The sysplex cascaded transactions environment is complex. When problems occur, determining what actions to take can also be complex. This section is intended to help you:

- Identify the problem
- Identify the systems involved
- Gather the relevant documentation
- Recover transaction processing

This section also covers the following sysplex-specific cascaded transaction environment problems:

RRS problem determination

- “[Sysplex cascaded transaction hang](#)” on page 298
- “[Sysplex cascaded transaction hang messages ATR246I and ATR247E](#)” on page 299

RRS sysplex cascaded transaction processing exploiters include IMS TM, OTMA, and APPC synchronous, shared queue, cross queue transaction processing.

Collecting documentation for a sysplex cascaded transaction environment

- **RRS documentation:** Collect an RRS dump and CTRACE as covered in “[Collecting documentation for RRS](#)” on page 285. Use a IEADMCxx parmlib member for the RRS dump options and, if possible, use an external writer to trace to a data set as documented in “[Important RRS CTRACE information](#)” on page 286.
- **Dumping IMS:** Prepare to dump IMS by creating an IEADMCyy parmlib member and include the following statements:

```
JOBNAME=(IMScontrolregion),  
SDATA=(ALLNUC,LPA,LSQA,PSA,RGN,SQA,TRT,CSA,GRSQ,SUM)
```

Then, when a dump of IMS is necessary, enter the following z/OS command, referencing your IEADMCyy parmlib member:

```
DUMP COMM=(whatever title you choose),PARMLIB=yy
```

When you want to dump both RRS and IMS, use the following z/OS command, where xx and yy are the suffix of the IEADMCxx members created for RRS and IMS:

```
DUMP COMM=(whatever title you choose),PARMLIB=(xx,yy)
```

- **Tracing IMS:** To trace IMS, either run the following traces or specify them in the DFSVSMxx member of SYS1.PROCLIB. If possible, use option log and external trace data sets to externalize the trace data.
 - For IMS data, : run or specify IMS RRST trace:

```
/TRA SET ON TABLE RRST
```

- For OTMA users, run or specify the OTMT trace:

```
/TRA SET ON TABLE OTMT
```

- For APPC users, run or specify IMS LUMI trace:

```
/TRA SET ON TABLE LUMI
```

If you are using external trace data sets, the data can wrap and that IMS does not provide archive support. It is the users responsibility to ensure that the data is not overwritten. Consider automation that copies the inactive data set when IMS switches from one data set to another, or preserve the data manually in problem situations.

Sysplex cascaded transaction hang

Symptoms

Any of the following can alert you to a sysplex cascaded transaction hang:

- IMS Display Region command shows hung threads:

```
/DIS ACT REG
```

- IMS shared message queue growth
- RRS display command output shows UR exceptions:

```
DISPLAY RRS,UREXceptions
```

How to investigate

- Use the IMS /DIS ACT REG command to display status of regions that seem to be hung. If a thread is active in the same state for two consecutive displays, that suggests that the thread is hung. You can also use the DISPLAY RRS,UREXceptions command to display status of transaction exceptions.
- When a transaction is hung, obtain a dump of RRS and IMS on both the frontend (FE) and back-end (BE) systems. The FE and BE system names can be determined by the /DIS ACTIVE REGION output. The following z/OS command can be used to obtain the dumps:

```
RO (fesysname,besysname),DUMP COMM=(title you choose),PARMLIB=(xx,yy)
```

- If you cannot determine the systems that are involved in the hang, request dumps on each system in the sysplex. Also collect the IMS SLDS and IMS control region job logs for all systems dumped, OPERLOG or SYSLOG, and LOGREC. Use the following command to dump all systems:

```
RO *ALL,DUMP COMM=(title you choose),PARMLIB=(xx,yy)
```

- If you find IMS regions that are hung with a status other than what is in this list, dump all the IMS and RRS pairs that are participating in the shared queue group:
 - WAIT-SYNCPOINT
 - WAIT-RRS/PC
 - TERM-WAIT SYNCPT
 - TERM-WAIT RRS

Recovery actions

If the IMS /DIS ACTIVE REGION display shows a hung dependent region:

- Collect the following documentation from the back-end and frontend systems:

```
RO (fesysname,besysname),DUMP COMM=(title),PARMLIB=(xx,yy)
```

- Try to terminate the hung dependent region with the following command:

```
/STO REG ABDUMP
```

- If you cannot terminate the hung thread, you might recycle the IMS control region that owns the hung dependent region.
- If the problem persists, check for RRS high processor usage spikes, which suggest that RRS must be recycled.

Sysplex cascaded transaction hang messages ATR246I and ATR247E

Symptoms

Look for the following messages:

- ATR246I RRS HAS DETECTED A *controlblockname* CONTROL BLOCK ERROR - UNEXPECTED ERROR DUMP REQUESTED
- ATR247E RRS HAS DETECTED A SEVERE ERROR – TERMINATE RMS AND OPTIONALLY REPLY SHUTDOWN TO SHUTDOWN RRS.

How to investigate

- Message ATR246I indicates that RRS detected a potential problem. RRS takes an unexpected error dump and continues processing without impact. Collect the dump and report this issue to the IBM Support Center.
- Message ATR247E indicates that RRS detects a more severe error, and is expecting that action be taken to terminate RRS. This error might be due to a corrupted control block chain and RRS cannot continue processing.

RRS problem determination

Recovery actions

For ATR246I, no recovery action is necessary.

For ATR247E, take the following actions:

1. Collect dumps of RRS and IMS on all systems:

```
ROUTE *ALL,DUMP COMM=(title),PARMLIB=(xx,yy)
```

2. Cancel IMS on this system:

```
F IMS,STOP
```

3. After IMS is stopped, terminate RRS replying to ATR247E with SHUTDOWN. After RRS is terminated, the message is DOMed.

4. Start RRS with the following command:

```
S RRS
```

5. After RRS finishes initializing, start IMS with the following command:

```
S IMS
```

Best practices

Create an automation routine that looks for the message ATR247E message and take actions in [“Recovery actions”](#) on page 300.

Chapter 21. System Data Mover (SDM) operational problem determination

The system data mover (SDM) is a DFSMS/MVS component that interacts with data storage subsystems and with various advanced copy services functions to efficiently move large amounts of data. As updates occur to primary volumes, the SDM manages the process of copying those updates to secondary volumes.

The SDM ensures that updates to secondary volumes are made in the same order in which they were made to the primary volumes, maintaining sequence consistency.

Basic SDM problem determination functions

For problem determination with SDM related issues, the following items are used:

- CQUERY DEVN('XXXX') PATHS
- F ANTAS000,LISTSESS ALL (to list XRC sessions on the HW)
- XQUERY ssid SC DETAIL (reports all primary storage controls with volumes in XRC session)
- SYSLOG, LOGREC, JOBLLOG, System dumps (General documentation to collect, if available)

SDM specific problems

- [“ANTP0095I Unable to determine PPRC paths” on page 301](#)
- [“ANTX5104E RC=0901 \(XRC\)” on page 302](#)
- [“ANTX5104E RC=0647 REASON=0053 \(XRC\)” on page 303](#)
- [“ANTX5104E RC=0647 REASON=0002 \(XRC\)” on page 303](#)
- [“ANTAS00* ASIDs consuming excessive storage below 2GB” on page 304](#)
- [“Converting to IR, RC=1017” on page 304](#)
- [“Microcode issue impacting concurrent copy” on page 305](#)

ANTP0095I Unable to determine PPRC paths

While querying paths using TSO CQUERY DEVN(XXXX) PATHS to PPRC secondary devices, the path status reports as "UNABLE TO DETERMINE".

Symptoms

```

CQUERY DEVN(X'dddd') PATHS.
ANTP0095I CQUERY FORMATTED LVL 3 .
PATHS REPORT.
***** PPRC REMOTE COPY CQUERY - PATHS *****
* PRIMARY UNIT: SERIAL#= 000000serial SSID= ssid SS= 2107 LSS= ls *
*          FIRST          SECOND          THIRD          FOURTH          *
*          SECONDARY      SECONDARY      SECONDARY      SECONDARY      *
*SERIAL NO: 000000serial ..... *
* SSID LSS:  ssid ls ..... *
*  PATHS:      0          0          0          0          *
*          SAID DEST S*  SAID DEST S*  SAID DEST S*  SAID DEST S*  *
*          ----- *
*          1: XXXX XXXX FF  ----  ----  00  ----  ----  00  ----  ----  00  *
*          2: XXXX XXXX FF  ----  ----  00  ----  ----  00  ----  ----  00  *
*          3: XXXX XXXX FF  ----  ----  00  ----  ----  00  ----  ----  00  *
*          4: XXXX XXXX FF  ----  ----  00  ----  ----  00  ----  ----  00  *
* SUBSYSTEM          WWNN          LIC LEVEL          *
* ----- *
* PRIMARY... (primary wwnn )          x.x.xx.xxx          *
* SECONDARY.1 (secondary wwnn)          *
*

```

SDM problem determination

```
* S* = PATH STATUS: *.  
* 00=NO PATH          01=ESTABLISHED ESCON  02=INIT FAILED *.  
* 03=TIME OUT         04=NO RESOURCES AT PRI 05=NO RESOURCES AT SEC*.  
* 06=SERIAL# MISMATCH 07=SEC SSID MISMATCH 08=ESCON LINK OFFLINE *.  
* 09=ESTABLISH RETRY  0A=PATH ACTIVE TO HOST 0B=PATH TO SAME CLUSTR*.  
* 10=CONFIG ERROR    FF=UNABLE TO DETERMINE *.  
*****  
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE dddd. COMPLETION CODE: 00 .
```

How to investigate

This query output is a normal response from the PPRC secondary devices. Although the paths are defined, they are not recognized until they are actually in use.

Recovery actions

No applicable recovery.

Best practices

None.

ANTX5104E RC=0901 (XRC)

This error is issued if the data mover detects either a No Record Found or Invalid Track Format error on a volume. The data mover will automatically reinitialize the track on which the error occurred. If a subsequent error occurs during resynchronization for the same volume, the volume pair will be suspended.

Symptoms

This error is only reporting the "out of synch" condition when XRC detects it. Post analysis is needed to determine how the volume pair got out of synch.

How to investigate

Run the Dynamic Volume Compare Utility to scan for any discrepancies between volumes. See [XRC Dynamic Volume Compare \(www.ibm.com/support/docview.wss?uid=isg3T7000248\)](http://www.ibm.com/support/docview.wss?uid=isg3T7000248) for instructions on using the utility.

FTP the following documentation to the IBM support center for further diagnostics:

- DUMP (if available)

If not, ensure the following diagnostic flags are enabled to collect a dump on any future hits:

```
F ANTAS00n,CTFLG NRFITF ON  
F ANTAS00n,CTFLG TIF ON  
F ANTAS00n,CTFLG TIF_ERROR ON  
F ANTAS00n,CTFLG ABEND_LIC ON
```

- SYSLOG
- LOGREC

Recovery actions

The user should XDELPAIR, then XADDPAIR the volume pair that received the RC901 error. This will perform a full copy of the data, ensuring the volume is 100% mirrored. After the mirror is up,

Technote: See: <http://www-01.ibm.com/support/docview.wss?uid=isg3T1018957>

Best practices

ANTX5104E RC=0647 REASON=0053 (XRC)

XRC has detected this condition during normal data mover processing. The scope of this condition is for a single storage control session. After you determine the reason for the error and correct the condition, you can issue an XADPAIR command to add the suspended volumes back to the session. This is for the 5695DF117 System Data Mover (SDM) component.

Symptoms

Error message reported:

```
ANTX5119E XRC SESSION(session) ENCOUNTERED AN ERROR PROCESSING STORAGE CONTROL
ssid SESSION ##, RC=0647 REAS=0053 SRVC=0106
```

If a SUSSESS ALL command was entered either directly by the user, or by automation, then this is an expected condition for every session in the data mover.

If unexpected, look for other ANTXxxxxE messages containing a different return and reason code in the SYSLOG around the time of the suspend. This will be the original reason for the suspend. The REAS53 messages are symptoms of the original suspend

How to investigate

This message is a result of the suspension and a preceding message should be searched for.

FTP the following documentation to the IBM support center for further diagnostics: SYSLOG

Recovery actions

No applicable recovery.

Best practices**ANTX5104E RC=0647 REASON=0002 (XRC)**

An attempt to issue I/O from the SDM host to the primary storage subsystem has timed out. The timeout value is stored in the user defined StorageControlTimeout value. This is for component 5695DF117 System Data Mover (SDM).

Symptoms

Error messages:

- ANTX5119E XRC SESSION(session) ENCOUNTERED AN ERROR PROCESSING STORAGECONTROL ssid SESSION ##, RC=0647 REAS=0002 SRVC=0106
- Possible IOS messages

How to investigate

Typically, the presence of IOS messages indicate a networking issue. The I/O timeout is justified by the loss of connectivity from the SDM hosts to the primary storage subsystem. Check the SYSLOG for IOS messages. If present, users should follow up with their networking vendor. Once the network is stabilized, the user can resume the mirror.

FTP the following documentation to the IBM support center for further diagnostics: SYSLOG

Recovery actions

Correct any connectivity issues in the hardware network.

Best practices

ANTAS00* ASIDs consuming excessive storage below 2GB

When running close to the real storage installation, you may experience messages indicating storage shortages on SDM LPARs. Because SDM address spaces use a large amount of fixed storage, the SDM address spaces will always report first in the IRA404I messages as users of large amounts of storage.

Symptoms

Error messages reported:

- IRA400E 04,PAGEABLE STORAGE SHORTAGE
- IRA404I ANTAS00x ASID 00xx OWNS 0000##### PAGES, 0000##### FIXED, 0000##### FIXED IN SHORTAGE AREA
- IRA401E 04,CRITICAL PAGEABLE STORAGE SHORTAGE (sometimes)

How to investigate

If you aren't witnessing a *IEA602I ADDRESS SPACE CREATE FAILED message, then this can be considered only a warning until storage is adjusted.

F ANTAS000,DUMP (generates dump, gathering storage usage from all address spaces)

Note: If no impact to environment, and only warnings are displayed, follow-up during next business day.

What to send:

- DUMPs (Generated by previous command, from all address spaces)
- SYSLOG (Including timestamps leading up to and from reported error)

Recovery actions

Reduce storage demand by terminating job tasks until shortage is relieved.

Best practices

Converting to IR, RC=1017

While converting from normal XRC to XRC-IR (PPRC and XRC), clients received error messages indicating there was a mismatch of storage control sessions on the SWAP (PPRC secondary) storage controls. This is for component 5695DF117 System Data Mover (SDM).

Symptoms

Error messages reported:

- ANTX5129E XRC SESSION(session) ERROR ASSOCIATING SWAP VOLUME(####) SCSESSION(aa) WITH PRIMARY VOLUME(volume) SCSESSION(aa), ANTX5129E (CONT) RC=1017 REAS=0
- ANTA5107E XADDPAIR FAILED FOR VOLUME PAIR(volume,XRCUTL) FOR SESSION(session), RC=1017 REAS=0

How to investigate

FTP the following documentation to the IBM support center for further diagnostics:

- F ANTAS000,LISTSESS ##### (for both primary and swap controllers, where ##### is the device number)
- F ANTAS000,LISTDVCS ##### ss (for all devices in affected session, ##### is the device number, and ss is the storage control session number)
- F ANTAS00n,DUMP (Dump of the impacted data mover address space)
- SYSLOG

- LOGREC

Recovery actions

Consult IBM Software Technical Support.

Best practices**Microcode issue impacting concurrent copy**

A known open microcode defect causes a counter on the box to reach its maximum limit without clearing itself. This counter keeps track of active sessions on a control unit (in hardware). There is a limit to how many active sessions (XRC and CC) you can have on a control unit, so when this counter is maxed out, the client is unable to execute more Concurrent Copy sessions. This issue is regarding the 5695DF117 System Data Mover (SDM) component.

Symptoms**Error messages reported:**

- ANTPC2 messages flood the LOGREC, one for every CC request, which can be hundreds.

How to investigate

- This issue requires the user to take a hardware statesave on the box. A statesave is like taking a dump, but it also performs "checks and balances".
- The user should open a Hardware PMR to request the statesave.

Recovery actions

Engage DASD Hardware Support for recovery assistance.

Best practices

Chapter 22. System logger operational problem determination

System logger is an MVS component that provides a set of services allowing an application to log data from a sysplex. Applications use these services to read or write log data to and from the log stream, which is simply a collection of data spanning interim storage (such as the coupling facility and DASD) and permanent DASD storage managed by system logger. As needed, data is offloaded to permanent DASD storage. System logger externalizes errors in allocating DASD data sets and monitors for delays in data movement in order to provide early alerts of filled log stream resource conditions that prevent applications from writing data.

For related information on resolving system logger allocation error, see the system logger chapter in the component reference of *z/OS MVS Diagnosis: Reference*.

Basic system logger problem determination functions

System Logger provides the following functions to help with problem determination, data collection, and recovery:

- Use the DISPLAY LOGGER command to query status of log streams and their connectors.
- Use CTRACE options specific to SYSLOGR to gather detailed information about system logger requests and events.
- Use the SETLOGR command to force a log stream to be disconnected from a system, delete a log stream from the LOGR policy, or stop system logger from waiting for a particular DASD data set to be recalled.
- Make sure that all System Logger health checks are activated and running to warn you of imminent or occurring storage usage full conditions. For details, see system logger checks (IBMIXGLOGR) in *IBM Health Checker for z/OS User's Guide*.

Collecting documentation for system logger

This section contains the following topics:

- [“Dumping system logger information” on page 307](#)
- [“Important system logger CTRACE information” on page 308](#)

Dumping system logger information

System logger provides four DUMP command parmlib samples in SYS1.SAMPLIB:

- IEADMCLC – Dumps system logger, Resource Recovery Service (RRS), and CICS data
- IEADMCLG – Dumps system logger and global resource serialization data
- IEADMCLS – Dumps system logger and all related data
- IEADMCLX – Dumps system logger and XCFdata

You can copy these SYS1.SAMPLIB members to SYS1.PARMLIB and run them using one of the following DUMP commands:

```
DUMP PARMLIB=xx
- OI -
DUMP TITLE=(dump_title_text),PARMLIB=xx
```

System logger problem determination

When requesting a system logger dump using the DUMP command, be sure to include the system logger component trace data by dumping the system logger dataspace. To do this, add the following parameter to the z/OS console DUMP command:

```
DSPNAME=( 'IXGLOGR' .*)
```

Important system logger CTRACE information

System logger suggests that you activate the following system logger CTRACE options so that you get the most trace information possible in logger dumps. You can either update the CTILOG00 parmlib member or create a CTILOGxx parmlib member with these trace options:

```
TRACEOPTS ON
          BUFSIZE(512M)
          OPTIONS('CONNECT','DATASET','SERIAL',
                 'LOGSTRM','MISC','RECOVERY','LOCBUFF')
```

Then use following z/OS command to activate the CTRACE settings:

```
TRACE CT,ON,COMP=SYSLOGR,PARM=CTILOGxx
```

Starting with z/OS V1R13, you can also use the SETLOGR CTRACE command to modify the system logger component trace options on the CTRACE parameter in the IXGCNFxx system logger parmlib member. You can use the SETLOGR CTRACE command as an alternative to the TRACE CT,ON,COMP=SYSLOGR,PARM=CTILOGxx command overriding the options in default CTRACE parmlib member CTILOG00. The SETLOGR CTRACE command syntax is:

```
SETLOGR CTRACE(parmlib_member_name)
```

System logger component-specific problems

- [“Delay in offload or log stream data set allocation” on page 308](#)
- [“System logger high CPU usage” on page 309](#)
- [“System logger allocation errors” on page 309](#)

Delay in offload or log stream data set allocation

System logger monitors both log stream offload activity and specific log stream data set allocation requests. If an offload or data set allocation request appears to be hung or is taking too long to complete, system logger issues messages after waiting set amounts of time for an offload or log stream data set request to make progress.

Symptoms

Symptoms of delay in offload or log stream data set allocation include the following messages:

```
IXG310I  SYSTEM LOGGER CURRENT OFFLOAD IS NOT PROGRESSING FOR
         LOGSTREAM logstream STRUCTURE: strname
         request DSN=dsnhlq.dsnlsn.dsnllq

IXG311I  SYSTEM LOGGER CURRENT OFFLOAD HAS NOT PROGRESSED DURING
         THE PAST seconds SECONDS FOR LOGSTREAM logstream
         STRUCTURE: strname
         request DSN=dsnhlq.dsnlsn.dsnllq

IXG312E  OFFLOAD DELAYED FOR logstream, REPLY "MONITOR", "IGNORE",
         "FAIL", "AUTOFAIL", OR "EXIT".

IXG271I  LOGGER DATA SET REQUEST IN group taskname SERVICE TASK
         DELAYED DURING THE PAST seconds SECONDS FOR
         LOGSTREAM logstream staging
```

```
DSN=dsname, DIAG=diag
```

```
IXG272E  LOGGER group taskname TASK DELAYED, REPLY "MONITOR",
         "IGNORE", "FAIL", "EXIT".
```

How to investigate

Look for inhibitors to offload processing by issuing the following DISPLAY commands:

```
D  LOGGER,C,LSN=logstreamname
D  LOGGER,L,LSN=logstreamname
D  LOGGER,STATUS,RECALLS
D  XCF,STRUCTURE,STRNAME=structname
D  GRS,C
```

Recovery actions

Before responding to messages IXG310I, IXG311I, IXG312E, IXG271I, or IXG272E, make sure you read the complete message descriptions in *z/OS MVS System Messages, Vol 10 (IXC-IJZP)*. Your response to these messages can inadvertently create further problems. For example, if system logger issues message IXG272E, only respond to this message after you have tried correcting any general allocation or recall related messages. If message IXG271I indicates that the request delay affects a log stream data set and you reply "FAIL" to accompanying message IXG272E, this might affect log stream duplexing and the log stream connection might fail. In general, replying "FAIL" to the IXG272E is not recommended and should only be done if you cannot determine why the request does not complete.

Take the following recovery actions:

1. Correct any problems observed from the DISPLAY commands. For example, if the DISPLAY LOGGER,STATUS,RECALLS command output includes message IXG601I stating that log stream data set recalls are not being satisfied, consider issuing the SETLOGR FORCE,NORECALL command to stop system logger from waiting for a particular data set to be recalled.
2. If the problem persists, respond to the IXG312E message to stop offload processing for the log stream referenced in the message to allow it to run on another system, if possible. Other work backed up by the delayed offload named in the IXG312E message might then begin to run.

Note that several of IXG312E messages might appear at the same time for different log streams that are offloading data, but often only one log stream is actually experiencing a delay - the others might simply be waiting for one log stream to finish its allocation processing. Usually the log stream causing the delay is the **first** log stream to be reported.

Another way to determine which log stream is causing the delay is to check to see if the data set name displayed in the message ends with an explicit sequence number, such as 'A0000001' for example. If so, this is usually the log stream experiencing a problem. On the other hand, if the data set name ends with '<SEQ#>' it's more likely that this log stream is just waiting on a delayed log stream.

System logger high CPU usage

If you detect high CPU usage by system logger, this suggests a loop. Make sure that system logger component trace is active and request a system logger dump. Typically, a system logger high CPU usage condition is caused by a log stream connector driving a workload - the component trace information can indicate which workload is causing problems and where it comes from.

See ["Collecting documentation for system logger" on page 307](#).

System logger allocation errors

Symptoms

When system logger detects allocation errors for offload or staging data sets, it issues first allocation message IXG251I and then other messages, such as IXG002E and IXG003I. For example:

```
IXG251I IKJ56893I
DATA SET
IXGLOGR.CICSTS13.CICSVR.DFHLGLOG.A0000000 NOT
  ALLOCATED+
IXG251I IGD17103I CATALOG ERROR WHILE DEFINING VSAM DATA
SET
  IXGLOGR.CICSTS13.CICSVR.DFHLGLOG.A0000000
IXG251I RETURN CODE IS 56 REASON CODE IS 6 IGG0CLFT
IXG251I IGD306I UNEXPECTED ERROR DURING IGG0CLFT PROCESSING
IXG251I RETURN CODE 56 REASON CODE 6
IXG251I THE MODULE THAT DETECTED THE ERROR IS IGDVTSCU
IXG251I SMS MODULE TRACE BACK - VTSCU VTSCT VTSCH VTSCD VTSCC

  VTSCR SIRT
IXG251I SYMPTOM RECORD CREATED, PROBLEM ID IS IGD00007
IXG251I IGD17219I UNABLE TO CONTINUE DEFINE OF DATA SET
  IXGLOGR.CICSTS13.CICSVR.DFHLGLOG.A0000000
IXG002E LOGR POLICY PROCESSING ENDED WITH RETCODE=00000008
  RSNCODE=00000805
IXG003I LOGR POLICY PROCESSING ENCOUNTERED AN UNEXPECTED ERROR.
DIAGNOSIS INFORMATION: 00000004 000042CF 0107001B 00000000
```

How to investigate

The diagnosis information displayed in message IXG003I contains the SMS return and reason codes from the component that returned the allocation error. You can use this information to search the system log for the IGD message identifying the component. For example, using the messages shown in [“Symptoms” on page 310](#), do the following:

1. Look at the second hexadecimal word of the diagnosis information in the message output for IXG003I
2. Convert the second word from hexadecimal to decimal. In message IXG003I in [“Symptoms” on page 310](#), the second word is X'42CF', which is 17103 decimal. This is the suffix for IGD message you are looking for, giving you message IGD17103.
3. Search the system log for message IGD17103I for more information related to the allocation error.

See message IDC3009I in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#) for an explanation of the catalog return and reason codes (such as return code 56, reason code 6) displayed in the IXG251I message issued immediately after the IGD17103I message.

Other common reasons for an allocation error may include:

- IXGLOGR ASID not defined with TRUSTED/PRIVILEGED authority
- Insufficient DASD space
- Data set not in the CATALOG or the CATALOG cannot be accessed.

Recovery actions

If the IXGLOGR ASID has not been defined with TRUSTED/PRIVILEGED authority, update the IXGLOGR address space to have TRUSTED/PRIVILEGED authority. For new authority to take effect, the IXGLOGR address space must be stopped and restarted.

If there is insufficient space on DASD to allocate the data set, free up space or allow SMS to use more volumes.

Chapter 23. VSAM component operational problem determination

VSAM is an access method that arranges records in data sets by an index key, relative record number or relative byte addressing. VSAM is used for direct or sequential access to either fixed-length or variable-length records on DASD. See the manual *DFSMS Using Data Sets* for more information on VSAM data sets.

Basic VSAM problem determination functions

For problem determination with VSAM data sets, the following items are used:

- IDCAMS LISTCAT ALL COMMAND
- IDCAMS EXAMINE DATATEST INDEXTEST for VSAM data sets with an index component
- SMF type 60 through 66 records contain information about VSAM data sets, including definition, deletion, altering of VSAM data sets
- System dumps produced by VSAM in certain conditions with no customer request
- System dumps produced as a result of issuing the command F CATALOG,VDUMPON to generate a dump when certain conditions occur
- EREP reports in some cases where a call to media manager produced an unexpected result.

VSAM specific problems

- [“VSAM Index Trap” on page 311](#)
- [“Hang in VSAM record management code” on page 313](#)
- [“Loop in VSAM record management code” on page 313](#)
- [“Unexpected return codes from VSAM record management” on page 314](#)
- [“Issues opening, closing, extending VSAM data sets” on page 314](#)

VSAM Index Trap

When writing a record to a data set with an index or an ICF catalog, VSAM record management will check to see if the adding or updating of the record will damage the index component. If the code detects that it will damage the index, the request is **not** done and any changes made during the request will be backed out.

Symptoms

- Message IDAI1001E for VSAM data sets
- Message IDAI1002E for ICF Catalogs
- RPL Feedback code of X'xx08006D' for the request that detected this
- RPL Feedback code of X'xx08006E' for subsequent accesses
- System dump produced VSAM DYNAMIC RPL DUMP - IDAM19R3 +0xxx FEEDBACK CODE: 0108006D

How to investigate

1. Close VSAM data set if possible.
2. Run IDCAMS LISTCAT ALL of the data set and any associated data sets, such as AIX's

```
//STEP1 EXEC PGM=IDCAMS,REGION=0M
//SYSPRINT DD SYSOUT=A
```

VSAM problem determination

```
//SYSIN DD *  
LISTCAT ENT(VSAM.DATA.SET) ALL
```

3. If the data set is closed, run the following IDCAMS commands:

```
EXAMINE INDEXTEST NODATATEST ... followed by  
EXAMINE NOINDEXTEST DATATEST
```

```
//STEP1 EXEC PGM=IDCAMS,REGION=0M  
//SYSPRINT DD SYSOUT=A  
//SYSIN DD *  
EXAMINE VSAM.DATA.SET INDEXTEST NODATATEST  
EXAMINE VSAM.DATA.SET NOINDEXTEST DATATEST
```

4. If the data set cannot be closed or is an ICF Catalog, run the following IDCAMS commands:

```
VERIFY  
EXAMINE IND EXTEST NODATATEST  
EXAMINE INDEXTEST NODATATEST,  
EXAMINE DATATEST NOINDEXTEST  
EXAMINE DATATEST NOINDEXTEST
```

```
//STEP1 EXEC PGM=IDCAMS,REGION=0M  
//SYSPRINT DD SYSOUT=A  
//SYSIN DD *  
VERIFY DATASET(VSAM.DATA.SET)  
EXAMINE VSAM.DATA.SET INDEXTEST NODATATEST  
EXAMINE VSAM.DATA.SET INDEXTEST NODATATEST  
EXAMINE VSAM.DATA.SET.NOINDEXTEST DATATEST  
EXAMINE VSAM.DATA.SET NOINDEXTEST DATATEST
```

5. Preserve the data set by renaming the cluster, data component and index component in case further information is needed from the data set. If the VSAM data set has associated components such as AIX's and PATHs these should also be renamed

```
//STEP1 EXEC PGM=IDCAMS,REGION=0M  
//SYSPRINT DD SYSOUT=A  
//SYSIN DD *  
ALTER VSAM.DATA.SET -  
NEWNAME(VSAM.DATA.SET.SAVE)  
ALTER VSAM.DATA.SET.DATA -  
NEWNAME(VSAM.DATA.SET.DATA.SAVE)  
ALTER VSAM.DATA.SET.INDEX -  
NEWNAME(VSAM.DATA.SET.INDEX.SAVE)
```

6. For ICF Catalogs, perform a DFDSS Physical Dump of the catalog, in case further information is needed.

```
//STEP001 EXEC PGM=ADDRSSU  
//SYSPRINT DD SYSOUT=*  
//DSYS004 DD DISP=(,CATLG),DSN=SYS1.DSSPHY.DUMP,  
// UNIT=SYSDA,VOL=SER=XXXXXX,  
// SPACE=(CYL,(1110,200),RLSE)  
//ISYS004 DD DISP=SHR,UNIT=SYSDA,VOL=SER=SYS004  
//SYSIN DD *  
DS(INCL(CATALOG.NAME)) -  
PHYSINDDNAME( -  
(ISYS004) -  
) -  
OUTDDNAME(DSYS004)
```

7. Gather SMF type 60 through 66 records from all systems that may have accessed the data set or catalog from a time when the data set or catalog was known to be good.
8. Submit the output from the above to the Support Center along with syslog and joblog (if available).

Recovery actions

- If the IDCAMS EXAMINE commands run with a return code of 4 or less, the data set or catalog should not need any further processing before being used.

- If the IDCAMS EXAMINE commands return a return code of 8 or higher, then the data set or catalog will need to be recovered using your normal recovery processing.

Best practices

S

Hang in VSAM record management code

When processing a VSAM data set, a hang may occur in VSAM record management code

Symptoms

- Job that is processing VSAM data sets stops doing I/O to the VSAM data sets and is not using CPU
- RMF or other Monitors may indicate job is waiting in a VSAM module named IDA019xx

How to investigate

1. Obtain system dump of address space with the following SDATA:
SDATA=(ALLNUC,CSA,GRSQ,LPA,LSQA,PSA,RGN,SQA,SUM,SWA,TRT)
2. If the monitor includes detail on the data set that the last request was for, obtain an IDCAMS LISTCAT ALL of the data set and any associated data sets such as AIX's

```
//STEP1    EXEC   PGM=IDCAMS,REGION=0M
//SYSPRINT DD    SYSOUT=A
//SYSIN    DD    *
LISTCAT ENT(VSAM.DATA.SET) ALL
```

3. Submit the dump, LISTCAT output (if available), joblog from the job that is hung, and syslog to the Support Center.

Recovery actions

Cancel the hung job and attempt to restart.

Best practices

Loop in VSAM record management code

When processing a VSAM data set, a loop may occur in VSAM record management code.

Symptoms

- A job that is processing VSAM data sets stops doing I/O to the VSAM data sets and is continuing to use CPU
- RMF or other Monitors may indicate that a job is using CPU in a VSAM module named IDA019xx
- Some applications may indicate that a loop is occurring. For example, CICS will issue message DFHFC0004 indicating a loop at offset X'FFFF' in module DFHFCVR.

How to investigate

1. Obtain system dump of address space with the following SDATA:
SDATA=(ALLNUC,CSA,GRSQ,LPA,LSQA,PSA,RGN,SQA,SUM,SWA,TRT)
2. If the monitor includes detail on the data set that the last request was for, obtain an IDCAMS LISTCAT ALL of the data set and any associated data sets such as AIX's

```
//STEP1    EXEC   PGM=IDCAMS,REGION=0M
//SYSPRINT DD    SYSOUT=A
//SYSIN    DD    *
LISTCAT ENT(VSAM.DATA.SET) ALL
```

VSAM problem determination

3. Submit the dump, LISTCAT output (if available), joblog from the job that is hung, and syslog to the Support Center.

Recovery actions

Cancel or shutdown the hung job and attempt to restart.

Best practices

Unexpected return codes from VSAM record management

Code to process VSAM data sets will normally handle expected return codes such as No Record Found, Duplicate Key, and Successful request. However at times, other return codes may be returned that are not handled by the code or indicate an error in VSAM record management processing.

Symptoms

The application or program issues a message or otherwise indicates it received an unusual return code from a request to VSAM record management.

How to investigate

1. Review the RPL Feedback code returned with the descriptions in the manual *DFSMS: Macro Instructions for Data Sets*. If the RPL feedback code is one that the job/application should handle, change the application to handle that feedback code.
2. If the RPL Feedback code indicates the problem may be in VSAM code a dump may be needed to determine the cause. Review the use of the F CATALOG,VDUMPON command and issue the command to capture a dump or contact the Support Center who can assist in providing the values for the F CATALOG,VDUMPON command.
3. Submit the dump, joblog from the job that is receiving the return code, and syslog to the Support Center.

Recovery actions

- If the job/application needs to be changed, change the application.
- The Support Center will provide guidance for any other cases.

Best practices

Issues opening, closing, extending VSAM data sets

At times there may be issues with opening, closing or extending VSAM data sets.

Symptoms

- Message IEC161I is issued for open problems (or in some cases information)
- Message IEC251I is issued for close problems (or in some cases information)
- Message IEC070I is issued for extend problems

How to investigate

1. Review the messages and codes manual for the specific return and reason codes. If they are informational decide if further action is necessary from your applications perspective. For example, some common messages such as

```
IEC161I 001(DW) -255,P0331406,STEP1,AURM,, ,VSAM.DATA.SET
```

are informational as are some other messages such as the following:

```
IEC161I 056-084,P0331406,STEP1,AURM,,,VSAM.DATA.SET  
IEC161I 062-086,P0331406,STEP1,AURM,,,VSAM.DATA.SET
```

The return and reason codes for messages IEC251I and IEC070 are documented under message IEC161I

2. In some cases VSAM may take a DUMP during this processing. These dumps will have a title similar to the following:

```
DUMP TITLE=VSAM O/C/E0V FFDC DUMP - IDA0192Y + 00000298 RC=020 CCC=054
```

The return and reason codes are the same as those documented for messages IEC161I. The reason for the dump being taken was to gather doc at the time of the issue.

3. If the reason for problem in open, close or extend processing is not clear from the message description, submit the joblog, syslog and any dumps produced to the support center .

Recovery actions

- If the job/application needs to be changed according the messages received, change the job/application.
- The Support Center will provide guidance for any other cases.

Best practices

Chapter 24. VSAM record-level sharing (RLS) operational problem determination

VSAM record-level sharing (RLS) is a data set access mode that allows multiple address spaces, CICS application owning regions (AORs) on multiple MVS systems, and jobs to access data at the same time. With VSAM RLS, multiple CICS systems can directly access a shared VSAM data set, eliminating the need for function shipping between application owning regions (AORs) and file owning regions (FORs). CICS provides the logging, commit, and rollback functions for VSAM recoverable files; VSAM provides record-level serialization and cross-system caching. CICS, not VSAM, provides the recoverable files function.

For additional info reference chapter 16 of *z/OS DFSMSdfp Storage Administration*.

Basic VSAM record-level sharing (RLS) problem determination functions

The following is a list of the commands that may be useful in diagnosing VSAM RLS problems. This is a referential list, and will be repeated in the subsequent sections about various possible problems.

Display commands:

- D GRS,C
 - Shows any outstanding GRS conflicts. Watch for SYSVSAM ENQs.
- D SMS,SMSVSAM,DIAG(C)
 - Lists any outstanding latch contention and the associated delays and the TCBs involved.
 - SCOPE=SYSTEM
- D SMS,CFLS
 - • Shows current lock structure information, such as false contention and lock rate.
- D XCF,STR,STRNM=strname (usually IGWLOCK00 or cache structure)
 - Lists XCF information related to the SMSVSAM coupling facility structures.
 - Can be used for lock structures or cache structures.
- D SMS,SMSVSAM,QUIESCE
 - Outlines any active quiesce activity within the SMSVSAM region on that system. This can often reveal potential hangs if a registered region is not responding.
 - SCOPE=SYSTEM
- D SMS,SMSVSAM,ALL
 - Shows the status of SMSVSAM around the sysplex.
 - SCOPE=SYSPLEX
- D SMS,TRANVSAM,ALL
 - Provides the status of Transactional VSAM (TVS) around the PLEX.
 - SCOPE=SYSPLEX

Bringing up/Taking down SMSVSAM:

- V SMS,SMSVSAM,ACTIVE
 - Starts SMSVSAM after a TERMINATESERVER
- V SMS,SMSVSAM,TERMINATESERVER
 - Brings down the SMSVSAM ASID

If TERMINATESERVER fails:

VSAM RLS problem determination

- FORCE SMSVSAM,ARM
 - Force down SMSVSAM while still allowing recovery routines to run
- FORCE SMSVSAM
 - Force SMSVSAM into end-of-memory

Dump commands:

- Console dump

```
DUMP COMM=(some meaningful dump title)
R xx,JOBNAME=(SMSVSAM,XCFAS),CONT
R yy,DSPNAME=('SMSVSAM'.*, 'XCFAS' .*),CONT
R nn,SDATA=(PSA,NUC,SQA,LSQA,SUM,RGN,GRSQ,LPA,
            TRT,CSA,XESDATA),CONT
R zz,REMOTE=(SYSLIST=(*('SMSVSAM')),DSPNAME,SDATA),END
```

- Dump using IEADMCxx PARMLIB member:

```
JOBNAME=(*MASTER*,SMSVSAM),DSPNAME=('SMSVSAM' .*),
SDATA=(COUPLE,PSA,NUC,SQA,LSQA,SUM,RGN,GRSQ,LPA,TRT,CSA,XESDATA),
REMOTE=(SYSLIST=(*('SMSVSAM')),DSPNAME,SDATA),END
```

VSAM record-level sharing (RLS) specific problems

- [“HANG/WAIT in RLS/TVS” on page 319](#)
- [“ABEND0F4 failures” on page 318](#)
- [“SMSVSAM will not start up” on page 320](#)
- [“Share Control Datasets not specified” on page 320](#)

ABEND0F4 failures

Symptoms

In the event of a logic error internal to the RLS code, an ABEND0F4 dump will be generated. In many cases, DUMPSRV will also generate dumps on the SMSVSAM images across the plex. The symptom string will include the 5695DF122 (COMPID=DF122) component ID. No additional console commands are typically necessary to diagnose this type of error.

How to investigate

Search IBM support portal for any APARs that match the symptoms of the ABEND

Recovery actions

For most ABEND issues, there is no immediate solution. However, recovery may involve restarting SMSVSAM, if it has not automatically done so. For many issues, no actions are required. To identify any pertinent recovery actions, check the IBM support portal for any APARs that match the symptoms of the ABEND.

Best practices

Listed in APAR if available.

HANG/WAIT in RLS/TVS

Symptoms

If transactions are not processing, quiesces not quiescing, or data sets failing to open or close, you may have a form of hang within the SMSVSAM address space.

How to investigate

In order to troubleshoot the vast majority of RLS hang / wait / slowdown / loop situations, the following diagnostic commands will need to be issued:

```
- D GRS,C                - ENQ contention (system level)
- D SMS,SMSVSAM,DIAG(C) - RLS latch contention (system level)
- D SMS,SMSVSAM,QUIESCE - Quiesce event status (system level)
- IDCAMS SHCDS LISTALL  - lists registered subsystems & lock info
- D SMS,CFLS(lock_structure) - displays lock structure information
- D XCF,STR,STRNM=[IGWLOCK00 | secondary_lock_structure]
                          - another display of the lock str
```

For all "system level" commands, ensure that they are issued on every system in the plex. The best way to accomplish this is to use the route command. For example, RO *ALL,D GRS,C

Once the commands have been issued, dump SMSVSAM around the plex. Be sure to including the DATASPACEs for RLS as well as a minimum of SDATA parms GRSQ & XESDATA. If any CICS regions are affected, ensure that they are added to the dump specification as well.

Here is an example command to dump RLS, XCF and a CICS region on one system:

```
DUMP COMM=(some meaningful dump title)
R xx,JOBNAME=(SMSVSAM,XCFAS,CICS1),CONT
R yy,DSPNAME=('SMSVSAM'.*, 'XCFAS'.*),CONT
R nn,SD=(COUPLE,PSA,NUC,SQA,LSQA,SUM,RGN,GRSQ,LPA,TRT,CSA,XESDATA),END
```

Adding the REMOTE keyword will issue the same dump command on each member in the plex:

```
DUMP COMM=(some meaningful dump title)
R xx,JOBNAME=(SMSVSAM,XCFAS,CICS1),CONT
R yy,DSPNAME=('SMSVSAM'.*, 'XCFAS'.*),CONT
R nn,SD=(COUPLE,PSA,NUC,SQA,LSQA,SUM,RGN,GRSQ,LPA,TRT,CSA,XESDATA),CONT
R zz,REMOTE=(SYSLIST=(*('SMSVSAM')),DSPNAME,SDATA),END
```

The dumping process can be simplified by including an entry similar to the following example in the IEADMCxx PAMRLIB member:

```
JOBNAME=(*MASTER*,SMSVSAM,CICS1),DSPNAME=('SMSVSAM'.*),
SDATA=(COUPLE,PSA,NUC,SQA,LSQA,SUM,RGN,GRSQ,LPA,TRT,CSA,XESDATA),
REMOTE=(SYSLIST=(*('SMSVSAM')),DSPNAME,SDATA)
```

Once the member is created, issuing DUMP COMM=(title),PARMLIB=xx will dump RLS all around the plex.

Recovery actions

As with all RLS problems, please be sure to collect appropriate documentation before attempting to clear the issue. Without documentation, support will be unable to verify the cause of the problem.

FTP the following documentation to the IBM support center for further diagnostics:

- DUMPs
- OPERLOG (or SYSLOG from all plex systems)
- LOGREC (from all systems)
- JOBLOGS (any which may be pertinent)

To clear the issue, start by investigating the task/lock/latch that the diagnostic command indicate is holding up the system and then attempt to clear the specific resource by cancelling the affected job /

VSAM RLS problem determination

transaction / region / or system. Info APAR II14597 provides a detailed step by step set of instructions for this and other common scenarios.

Best practices

SMSVSAM will not start up

Symptoms

If the SMSVSAM address space is active, but no IGW414I message is displayed.

How to investigate

Issue a D SMS,SMSVSAM,ALL command. The output will show the current status of the SMSVSAM instances around the plex. The STEP should indicate SmsVsamInitComplete if there were no problems. For any other issue, contact the support center and provide the OPERLOG showing the output of D SMS,SMSVSAM command

```
IGW420I DISPLAY SMS,SMSVSAM,ALL
DISPLAY SMS,SMSVSAM - SERVER STATUS
  SYSNAME: SYSTEM1 UNAVAILABLE ASID: 00C5 STEP: SHC_Ph2_Init
  ....
```

Recovery actions

Best practices

Share Control Datasets not specified

Symptoms

Receiving IGW611A and IGW609A

```
*08.41.23 SYSTEM1          *IGW611A SHARE CONTROL DATA SET NEVER ASSIGNED
*08.41.23 SYSTEM1          *IGW609A NO SPARE SHARE CONTROL DATA SETS
* EXIST. IMMEDIATE ACTION REQUIRED
```

How to investigate

Issue:

```
D SMS,SMSVSAM,ALL
```

If it displays a STEP of SHC_PH2_Init, this means that SMSVSAM is waiting for the user to enter the name of a Share Control Dataset (SHCDS) for SMSVSAM.

Recovery actions

In a normal setup we require 2 active SHCDS datasets, and 1 spare. If the SHCDS have already been defined, then simply issue the following commands to add them to SMSVSAM:

```
V SMS,SHCDS(shcdfs.name),NEW
V SMS,SHCDS(shcdfs.sparename),NEWSPARE
```

If they are not yet defined, refer to the *DFSMSdfp Storage Administration* manual, section “Defining sharing control data sets”.

Best practices

Ensure SHCDS are properly defined prior to starting SMSVSAM.

Chapter 25. XES and XCF operational problem determination

Cross system coupling facility (XCF) and cross-system extended services (XES) are z/OS components that provide services for a sysplex and Parallel Sysplex environment respectively. These components:

- Manage resources unique to a sysplex environment
- Monitor the health of systems in the sysplex
- Monitor the applications that exploit sysplex services and resources

XES and XCF find and report hung or unresponsive systems in the sysplex, XCF group members, or XES structure connectors, taking action if necessary to isolate the impact on the sysplex.

Basic XES and XCF problem determination functions

XES and XCF provides the following functions to help with problem determination, data collection, and recovery:

- System commands:
 - Use the DISPLAY XCF and DISPLAY CF commands to query the status of:
 - Systems in the sysplex
 - Signaling paths
 - Groups
 - Couple data sets
 - Policies
 - Coupling facilities and structures
 - Use the SETXCF command to control XCF. For example, you can use SETXCF to start or stop signaling paths, start or stop a new policy, or put a CF into maintenance mode.

See the DISPLAY and SETXCF commands in *z/OS MVS System Commands*.

- Use SVC dump SDATA options COUPLE and XESDATA to gather XES/XCF information in a dump. See SVC dump information in *z/OS MVS Diagnosis: Tools and Service Aids*.
- Use CTRACE options SYSXCF and SYSXES to gather detailed CTRACE information for XCF/XES requests and events. See component trace information in *z/OS MVS Diagnosis: Tools and Service Aids*
- XCF and XES issue messages with IXC and IXL prefixes, respectively. See *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.
- Make sure that all IBM Health Checker for z/OS XCF checks are activated and running to warn you of impending XES and XCF problems. For details, see the XCF check descriptions in *IBM Health Checker for z/OS User's Guide*.

XCF and XES component-specific problems

- [“Sysplex sympathy sickness” on page 322](#)
- [“XCF group member stall and signaling sympathy sickness” on page 322](#)
- [“Hangs in structure connectors using XES services ” on page 324](#)
- [“Unable to bring new or upgraded coupling facility into use ” on page 326](#)
- [“High CPU usage in the XCF address space \(XCFAS\)” on page 328](#)

Sysplex sympathy sickness

If you find that more than one system in the sysplex is experiencing problems, the first thing to check for is sysplex sympathy sickness. Sysplex sympathy sickness refers to a state where one unhealthy, non-responsive system impacts the health of other systems in the sysplex. For example, sysplex sympathy sickness might occur as a result of a system hanging while holding global resources or not completing sysplex partitioning.

XCF detects and reports that a system is hanging when it has not updated the sysplex couple data set at regular intervals. The system issues messages IXC427A or IXC426D or IXC101I indicating 'status update missing' for a hanging system. The system may also issue write to operator (WTOR) messages IXC102A , IXC402D, or IXC409D prompting the operator to reset the sick system and reply 'DOWN' if the system cannot be successfully fenced. If the 'DOWN' reply is delayed, other systems usually experience sympathy sickness, with the symptoms listed in ["Symptoms"](#) on page 322.

Symptoms

Symptoms of sysplex sympathy sickness include:

- Message ISG361A indicating GRS list lock contention
- Message \$HASP263 indicating JES checkpoint contention
- Message ISG633I indicating that GRS is running impaired
- Messages IOS071I and IOS431I indicating START PENDING status for devices because of reserves held by the sick system
- Global ENQ resource contention
- Multiple XCF group members or structure connectors detected hung with accompanying system messages such as IXC431I, IXC631I, IXC640E, IXL040E, IXL041E, and IXL045E.

How to investigate

Check for system message IXC101I followed by outstanding WTOR message IXC102A, IXC402D, or IXC409D prompting an operator to reset the sick or hung system. Reset the system and reply 'DOWN' to the WTOR so that the system can be partitioned - sysplex partitioning cleanup of resources by various functions and products does not occur until after the WTOR reply 'DOWN'.

Best practices

- Activate an SFM policy specifying ISOLATETIME(0) and CONNFAIL(YES) to automate sysplex partitioning.
- Enable the system status detection (SSD) function SYSTATDETECT and configure the sysplex with the base control program internal interface (BCPii). The SSD partitioning protocol exploits BCPii interfaces using z/Series hardware services to determine the status of failed systems in a sysplex. If SSD determines that a system has failed, it is automatically removed from the sysplex. SYSTATDETECT and BCPii almost eliminate the need for operator intervention and ensures that failing systems are removed promptly from the sysplex. See the topic on using the SSD partitioning protocol and BCPii in [z/OS MVS Setting Up a Sysplex](#).

XCF group member stall and signaling sympathy sickness

Signalling sympathy sickness occurs when a system tries to send signals to a target system that has run out of I/O buffers because the buffers are full of undelivered messages for an XCF group member that is not precessing signals in a timely manner (XCF group member stall). XCF monitors signal delivery at the XCF group member level to notify and prevent signaling sympathy sickness.

Symptoms

The symptoms of XCF group member stalls and signalling sympathy sickness include the following system messages:

- XCF issues the following message when it detects that signals are stalled for a group member or members for about 3 minutes:

```
IXC431I GROUP IMSPGRP1 MEMBER CSQ4 JOB CSQ4MSTR ASID 0062
        STALLED AT 07/06/2011 06:12:48.569782 ID: 0.29
        LAST MSGX: 07/06/2011 10:24:14.714602 6 STALLED 0 PENDINGQ
        LAST GRPX: 07/06/2011 09:32:39.773389 0 STALLED 0 PENDINGQ
```

When the stall condition is relieved, XCF issues message IXC432I.

- If the stall condition continues for about 5 minutes, XCF issues the following message:

```
*IXC430E SYSTEM P104 HAS STALLED XCF GROUP MEMBER
```

- If the XCF member or members stall condition is allowed to persist, it may escalate into an XCF inbound signaling buffer full condition, as reported in messages IXC467I and IXC440E:

```
IXC467I RESTARTING PATHOUT STRUCTURE IXCSTR1 LIST 8
        USED TO COMMUNICATE WITH SYSTEM SYS01
        RSN: I/O APPARENTLY STALLED
        DIAG073: 08200208 003E4609 003E4608 00000015

IXC440E STALLED XCF GROUP MEMBERS ON SYSTEM SYS01 IMPACTING SYSPLEX
```

- If an XCF member, such as GRS for example, is defined as a critical member and XCF detects that it is stalled or impaired, XCF issues system messages IXC631I, IXC635E, and IXC633I.
- If a system tries to join a sysplex, but cannot establish signalling connectivity with all the systems because signalling sympathy sickness is occurring, XCF issues message IXC445D.

How to investigate

1. Check for outstanding WTOR messages IXC102A, IXC402D, and IXC409D indicating that XCF is trying to partition a system. A stall condition can be caused by sysplex-wide sympathy sickness resulting when resources are held by systems waiting to be partitioned. Reset the system and reply 'DOWN' to these WTORs as soon as possible.
2. Start Runtime Diagnostics to analyze the system for critical messages, high CPU usage and loops, local lock contention, and GRSENQ and latch contention using the following commands:

```
S HZR, SUB=MSTR
F HZR, ANALYZE
```

See [Chapter 5, “Runtime Diagnostics,” on page 35.](#)

3. Issue a DISPLAY GRS, C command and look for the name of the stalled job to see whether it is waiting for a resource or because of latch contention.
4. Issue the DISPLAY XCF, GROUP, *grpname* command. In the output, stalled group members are flagged with an asterisk (*) if the stall condition persists for 30 seconds or more, and with an exclamation point (!) if the stall condition is impacting the sysplex:

```
IXC331I 14.52.49 DISPLAY XCF 752
* INDICATES STALL, ! INDICATES STALL IMPACTING SYSPLEX
GROUPS(SIZE):  SYSGRS(2)      SYSIEFTS(2)      SIGW00(2)
                SIGW01(2)    SYSIKJBC(2)      SYSIOS01(1)
                SYSIOS02(1)  SYSJES(2)        !*SYSMCS(7)
                SYSMCS2(9)   SYSTTRC(2)      *SYSWLM(2)
                SYSXCF(2)    XTTXGRP(2)      ZTTXNGRP(2)
```

5. Issue the DISPLAY XCF, PI, DEV=ALL command to display signaling paths that are stalled and causing sympathy sickness with a status of STALL-SS:

```
IXC356I 15.41.16 DISPLAY XCF 805
LOCAL DEVICE  REMOTE  PATHIN  REMOTE  LAST  MXFER
PATHIN        SYSTEM  STATUS  PATHOUT  RETRY  MAXMSG  RECDV  TIME
0CEE          D13ID72  STALL-SS  08E1    10    3000  17092  27671
```

Recovery actions

- Issue the DUMP command to request a dump of the stalled member's address space. If the IXC431I stall condition is currently active, set the following SLIP to trigger a dump on a recurrence:

```
SLIP SET,COMP=00C,RE=020F0006,AL=(H,P,S),ML=3,ID=STAL,  
SDATA=(RGN,COUPLE,XESDATA,ALLNUC,ALLPSA,CSA,GRSQ,SQA,TRT),END
```

See [z/OS MVS System Commands](#) for information about the SLIP command.

- Cancel or recycle the job or jobs causing the XCF signaling stall condition.

Best practices

Implement a sysplex failure management (SFM) Policy with the MEMSTALLTIME(*nnn*) option, where *nnn* is a value between 600 and 900. This allows the system to take automatic actions to resolve the XCF signalling stall condition when it begins to impact other systems in the sysplex. When the MEMSTALLTIME interval expires, the system requests a dump of the sysplex member consuming the most buffers and that member is then terminated to resolve the problem. See [z/OS MVS Setting Up a Sysplex](#).

Hangs in structure connectors using XES services

In a Parallel Sysplex environment, coupling facility structure connectors must participate in and respond to structure related events, such as rebuilds and other recovery processes within about 2 minutes. XES monitors connectors for a timely response.

Symptoms

The symptoms of a hang in a connector using XES services are the following system messages:

- If a structure connector does not respond to structure or rebuild events, XES issues system messages IXL040E or IXL041E with information about the delay. For example:

```
IXL040E CONNECTOR NAME: connector-name, JOBNAME: jobname, ASID: asid HAS  
text. process FOR STRUCTURE structure-name CANNOT CONTINUE.  
MONITORING FOR RESPONSE STARTED: mondate montime. DIAG: x
```

```
IXL041E CONNECTOR NAME: connector-name, JOBNAME: jobname, ASID: asid HAS  
NOT RESPONDED TO THE event FOR SUBJECT CONNECTION:  
subject-connector-name. process FOR STRUCTURE structure-name  
CANNOT CONTINUE. MONITORING FOR RESPONSE STARTED: mondate  
montime. DIAG: x
```

- If the condition signalled by messages IXL040E or IXL041E still persists after 2 minutes, XES issues abend code X'026' with reason code X'08118001' accompanied by a dump.
- XES issues message IXL045E indicating that an XES connector to a lock structure may be encountering delays due to limited XES SRB scheduling.
- When the connector finally provides the required response and XES monitoring is no longer necessary for the event, XES issues system message IXL042I, IXL043I, IXL047I, or IXL048I:

```
IXL042I CONNECTOR NAME: connector-name, JOBNAME: jobname, ASID: asid HAS  
action. THE REQUIRED RESPONSE event FOR STRUCTURE structure-name  
IS NO LONGER EXPECTED.
```

```
IXL043I CONNECTOR NAME: connector-name JOBNAME: jobname ASID: asid HAS  
action. THE REQUIRED RESPONSE FOR THE event FOR SUBJECT  
CONNECTION subject-connector-name, STRUCTURE structure-name IS  
NO LONGER EXPECTED.
```

```
IXL047I THE RESPONSE REQUIRED FROM CONNECTOR NAME: conname TO STRUCTURE  
strname, JOBNAME: jobname, ASID: asid responsetype IS NO LONGER  
EXPECTED. REASON: reason
```

```
IXL048I THE RESPONSE REQUIRED FROM CONNECTOR NAME: conname TO STRUCTURE  
strname, JOBNAME: jobname, ASID: asid FOR THE event FOR SUBJECT  
CONNECTION subjectconname IS NO LONGER EXPECTED. REASON: reason
```

- If the system takes action to resolve the hang, the system issues message IXL049E and IXL050I:

```
IXL049E HANG RESOLUTION ACTION FOR CONNECTOR NAME: conname TO STRUCTURE
strname, JOBNAME: jobname, ASID: asid: actiontext
```

```
IXL050I CONNECTOR NAME: conname TO STRUCTURE strname, JOBNAME: jobname,
ASID: asid HAS NOT PROVIDED A REQUIRED RESPONSE AFTER
noreponsetime SECONDS. TERMINATING termtarget TO RELIEVE THE
HANG.
```

How to investigate

1. Start Runtime Diagnostics to analyze the system for critical messages, high CPU usage and loops, local lock contention, and GRSEQ and latch contention using the following commands:

```
S HZR, SUB=MSTR
F HZR, ANALYZE
```

See [Chapter 5, “Runtime Diagnostics,”](#) on page 35.

2. Check for system message IXC518I indicating loss of connectivity to the coupling facility. Loss of connectivity to the coupling facility drives rebuilds for multiple structures at once, so contention for the coupling facility resource management (CFRM) couple data set might cause delays. Before taking action, wait to see if the condition resolves, as indicated by system messages IXL042I, IXL043I, IXL047I, or IXL048I.
3. If system messages IXL040E or IXL041E report just one connector or class of connectors as hung, do the following:
 - Issue system command `DISPLAY XCF, STR, STRNAME=strname` and look in the output for other hung connectors that might be affecting this connector. An asterisk (“*”) next to a connector indicates that the connector owes a response to another system, which might mean that it is hung.
 - Do general dispatchability checks for connectors that are hung:
 - Is the system running at or close to 100% capacity, causing the connector address spaces to contend for resources?
 - Is the connector running at too low a workload management service class importance? See the information about service classes in [z/OS MVS Planning: Workload Management](#).
 - Check for resource or latch contention in the affected address spaces by issuing the `DISPLAY GRS, LATCH, C` command.

See [z/OS MVS System Commands](#)

Recovery actions

If a signaling structure connector is hung and not responding, do the following:

1. Issue the following system command on all systems in sysplex to stop using the affected structures:

```
SETXCF STOP, PATHOUT | PATHIN, STRNM=structurename, UNCOND
```

2. Issue the following system command to bring the structure back into use:

```
SETXCF START, PI/P0, STRNM=structurename
```

3. If the hang persists, recycle or cancel the structure connectors address space.
4. If the system does not initiate a dump, collect dumps of connector address spaces, specifying `SDATA=(XESDATA, RGN, COUPLE, . . .)` as well as other options and address spaces required for the hung connector.

Best practices

- Implement CFSTRHANGTIME in your SFM policy to automate XES recovery for hung connectors. The recovery actions kicked off by CFSTRHANGTIME might include stopping a rebuild, terminating the

XES and XCF problem determination

connector's address task or address space, or partitioning the affected system. The system may initiate more than one of these actions if the first action taken does not relieve the hang. See the topic on SFM parameters in *z/OS MVS Setting Up a Sysplex*.

- Specify message based processing in your CFRM couple data set using the ITEM NAME(MSGBASED) parameter. With MSGBASED enabled, the system uses XCF signalling to coordinate recovery and rebuild staging data sets to reduce serialization contention on CFRM policy records. This improves recovery time and minimizes delays for structure connectors. See *z/OS MVS Setting Up a Sysplex*.

Unable to bring new or upgraded coupling facility into use

If you cannot bring a new or upgraded coupling facility into use, structures may fail in allocation. This can happen if you do not make the appropriate logical and physical configuration changes before bringing a new or upgraded coupling facility into use. Use the best practices for updating a coupling facility documented in *z/OS MVS Setting Up a Sysplex* to minimize the likelihood of problems.

Symptoms

If problems occur in bringing a new or upgraded coupling facility into use, the following symptoms might be reported:

- IXLCONN service requests to connect to a structure fail with return code X'C', reason code X'02010C08', indicating a lack of connectivity to the coupling facility.
- The new or upgraded coupling facility is not physically available, as indicated by message IXL158I indicating that the PATH is NOT OPERATIONAL or the lack of message IXL157I confirming that the path to the coupling facility is operational OR :

```
IXL158I  PATH chpid IS NOW NOT-OPERATIONAL TO CUID: cuid COUPLING
        FACILITY type.mfg.plant.sequence PARTITION: partition side
        CPCID: cpcid
```

OR

```
IXL157I  PATH chpid IS NOW OPERATIONAL TO CUID: cuid COUPLING FACILITY
        type.mfg.plant.sequence PARTITION: partition side CPCID: cpcid
```

- The new or upgraded coupling facility is not logically available, as indicated by the lack of message IXL517I confirming that the sysplex can now use the coupling facility.

```
IXC517I  SYSTEM sysname ABLE TO USE COUPLING FACILITY
        type.mfg.plant.sequence PARTITION: partition side CPCID: cpcid
        NAMED cfname
```

- If there is a loss of connectivity to a CF during normal operations for some reason, XES issues messages IXL158I and IXC518I:

```
IXL158I  PATH chpid IS NOW NOT-OPERATIONAL TO CUID: cuid COUPLING
        FACILITY type.mfg.plant.sequence PARTITION: partition side
        CPCID: cpcid
```

```
IXC518I  SYSTEM sysname NOT USING COUPLING FACILITY
        type.mfg.plant.sequence PARTITION: partition side CPCID: cpcid
        NAMED cfname REASON: text
```

- The first system IPLed into a disaster recovery site or into a sysplex that was taken down and is in GRS STAR mode encounters wait state X'A3' with reason code X'09C'. See the information about wait state X'A3' in *z/OS MVS System Codes*.

How to investigate

1. Verify that the paths to the coupling facility are physically online and operations, by issuing the following commands:
 - D CF,CFNAME=*cfname*

- Because you must define the CF LPAR link connectivity with Hardware Configuration Definition (HCD), you can use the CHPID number from the HCD definition in the following command:

```
d m=chp(f0)
RESPONSE=MCETA1
IEE174I 11.20.45 DISPLAY M 985
CHPID F0: TYPE=23, DESC=INTERNAL COUPLING PEER, ONLINE
COUPLING FACILITY 002817.IBM.51.000000F0DF6
PARTITION: 36 CPCID: 00
NAMED CF1 CONTROL UNIT ID: FFFB

SENDER PATH      PHYSICAL      LOGICAL      CHANNEL TYPE
F0               ONLINE       ONLINE       ICP

COUPLING FACILITY SUBCHANNEL STATUS
TOTAL: 14 IN USE: 14 NOT USING: 0 NOT USABLE: 0
OPERATIONAL DEVICES / SUBCHANNELS:
FFC8 / 0E61      FFC9 / 0E62      FFCA / 0E63      FFCB / 0E64
FFCC / 0E65      FFCD / 0E66      FFCE / 0E67      FFCF / 0E68
FFD0 / 0E69      FFD1 / 0E6A      FFD2 / 0E6B      FFD3 / 0E6C
FFD4 / 0E6D      FFD5 / 0E6E
```

If the paths to the coupling facility are physically offline, verify path status from a hardware perspective. On the Hardware Management Console (HMC):

- Find the image owning or sharing the channel path, right click on the image and select 'Chpids' from the pop-up menu.
- Select the target chpid.
- From the task list, open 'Chpid Operations'.
- From the list displayed, select 'Channel Problem Determination'
- Select 'Analyze Channel Information' to display details. From here, verify the status, node, and flag information presented.

If you find problems with the path status, work with the customer engineer (CE) to resolve them.

2. Compare node descriptor information from the actual coupling facility hardware with the node descriptor information specified in the CFRM active policy by comparing the output from the two following commands:

- For the hardware node descriptor information, issue:

```
D CF,CFNAME=cfname
```

- For the logical or CFRM policy node descriptor settings, issue:

```
D XCF,CF,CFNAME=cfname
```

3. Issue the following command to see if there is an old instance of the coupling facility with the same name that is pending deallocation:

```
D XCF,CF,CFNAME=cfname
```

This typically results when structures are not removed from the coupling facility before the coupling facility is taken out of use physically and logically with a CFRM policy change.

You can verify that structures are pending deallocation from an old coupling facility and have no active connectors with the following command:

```
D XCF,STR,STRNAME=strname
```

See [z/OS MVS System Commands](#).

Recovery actions

1. If there are no systems active in the sysplex, try re-IPLing a system, specifying one of the following:
 - GRS=TRYJOIN (to avoid wait state X'0A3')

XES and XCF problem determination

- PLEXCFG=XCFLOCAL
2. Use the IXCMIAPU administrative utility to verify and correct the CFRM active policy coupling facility node and descriptor information. Then re-IPL the system.
 3. If you find paths that are offline, issue the CONFIG command from the z/OS console or issue the CONFIGURE command from the HMC attached to the coupling facility LPAR to configure the paths off and back online.
 4. If you find coupling structures pending deallocation, verify that the structures are no longer in use by any system in the sysplex and force deletion with the following command:

```
SETXCF FORCE,PNDSTR,CFNAME=cfname
```

This command removes any structures that are pending-deallocation from the CFRM active policy because of the loss of connectivity to the coupling facility.

Best practices

See *Best practices for updating a coupling facility in z/OS MVS Setting Up a Sysplex*.

High CPU usage in the XCF address space (XCFAS)

The XCF address space (XCFAS) address space is a service provider for applications and subsystems that communicate and serialize with other members in a sysplex. As a result, a rise in the number of sysplex service requests also increases the CPU usage of the XCFAS address space.

The following are some factors that may contribute to an increase in XCFAS CPU usage:

- **Increase in the number of requests to coupling facilities:** If there is an increase in the number of requests for structures in the coupling facilities, there will be a corresponding increase in the number or running time of subchannel completion service request blocks (SRBs) driven on the backend in the XCFAS address space.
- **Increase in XCF signalling structure requests:** When XCF signalling paths between systems are defined in structures, completion of the signalling request runs in the XCFAS address space. An increase in the number of signals between XCF group members raises the CPU usage in the XCFAS. The challenge is to determine which XCF groups and/or members are causing the increase. A review of the RMF XCF Activity reports from before and after the high XCF signalling structure request situation can help identify the culprit.
- **Structures defined with system-managed duplexing:** When a structure is duplexed, the system creates 2 copies of the structure on 2 different coupling facilities. Updates to a duplexed structure drives an operation to both coupling facilities, resulting in higher XCFAS CPU usage than for a simplex structure.
- **Monitoring for structures reaching full threshold conditions:** When a structure is defined in the CFRM policy with a non-zero FULLTHRESHOLD parameter and ALLOWAUTOALTER(YES), the XCFAS checks on the structure usage every 60 seconds. If XCF detects that the structure has reached the FULLTHRESHOLD limit specified, the XCFAS kicks off the alter processing and increases the monitoring frequency to every 30 seconds and continues to monitor the structure every 30 seconds until the structure usage drops below the FULLTHRESHOLD value.
- **Contention for IXLSHR or IXLSHELL Facility spin locks with performance monitor products:** Performance monitor products such as DB2/SAP performance collector, SAPCL or RMF might issue IXLMG HWSTATISTICS=YES requests to gather CF and structure statistical information from the CF. This requires serialization using the spin locks and can cause contention for other serialized functions running in XCFAS address space.

See the information on RMF reports and the XCF activity report in *z/OS MVS Setting Up a Sysplex*.

How to investigate

- Begin by comparing RMF CF Activity reports with comparable workloads, one from a time of normal XCFAS CPU usage, and one from a high CPU usage time.

- For each coupling facility, look in the Coupling Facility Usage Summary section of the report and look for a structure that shows a significant increase in the rate of requests. The RMF spreadsheet reporter may be helpful in this task, see [RMF Spreadsheet Reporter for Windows \(www.ibm.com/systems/z/os/zos/features/rmf/tools/rmftools.html#spr_win\)](http://www.ibm.com/systems/z/os/zos/features/rmf/tools/rmftools.html#spr_win).
- If the increased request rate is observed for the signaling structures, then the next step is to review an RMF XCF Activity report using SMF 74.2 records collected for all systems in the sysplex. Review the Usage by Member report to determine which XCF group (members) are driving an higher rate of signaling requests.
- Sometimes high XCFAS CPU usage is due to repetitive coupling facility path recovery, which can occur when a coupling facility is defined with a shared CP. Issue the following command to determine whether the coupling facility is using shared or dedicated CPs:

```
D XCF,CF,CFNAME=cfname
```

See Redbook tip [Use of Shared Engines for Coupling Facilities - Implementation and Impact \(www.redbooks.ibm.com/abstracts/TIPS0237.html\)](http://www.redbooks.ibm.com/abstracts/TIPS0237.html).

- Try turning off any performance collection products that may collect hardware statistical information for the coupling facilities and/or structures to determine if performance collection is inadvertently leading to high XCFAS CPU usage by causing spinlock contention.

Recovery actions

1. Ensure that the coupling facility LPAR is defined with sufficient CP resource. If the coupling facility is in a production environment where response times are critical, IBM suggests that you use dedicated CPs.
2. If the preceding steps do not reveal a problem, request a console initiated SVC dump specifying:

```
JOBNAME=XCFAS,SDATA=(RGN,XESDATA,COUPLE,RGN,SQA,CSA,TRT,GRSQ)
```

Report the problem to the IBM Support Center.

Part 6. Diagnosis reference material

Before calling IBM, it is important to gather the correct information. The following topics can help you find specific diagnosis information for z/OS base elements and features and help you have the correct information available to discuss with the IBM support specialist.

Chapter 26. Diagnosis information for z/OS base elements and features

In addition to reading the topic about [Part 5, “Diagnosing component-specific problems,”](#) on page 257, find specific diagnosis information for z/OS base elements, features, and products that run on z/OS in [IBM Knowledge Center \(www.ibm.com/support/knowledgecenter/SSLTBW/welcome\)](#).

Chapter 27. Reporting problems to IBM

Before you begin: Be familiar with the information in this document; know how to collect the data that your software specialist needs to solve your problem.

This chapter covers the following topics:

- [“Software support service checklist” on page 335](#) including severity levels and examples.
- [“Automatic problem reporting” on page 337](#)
- [“Invoking IPCS as a background job” on page 337](#) which includes [“Step for invoking IPCS as a background job” on page 337](#)

Software support service checklist

In order to understand and resolve your software support service request in the most expedient way, it is important that you gather information about the problem and have it on hand when discussing the situation with the software specialist. The following information is required:

- Definition of the problem

Note:

It is very important that you are as specific as possible in explaining a problem or question to our software specialists. Our specialists want to be sure that they provide you with exactly the right solution so, the better they understand your specific problem scenario, the better they are able to resolve it. To assist you with problem identification, see the [Chapter 28, “Problem diagnostic worksheet,” on page 339](#) (in Appendix A).

- Background information

Note:

If possible, obtain all data about a problem soon after the problem occurs. Otherwise, updates to the system can cause discrepancies in the data. Ask yourself the following questions:

- What levels of software were you running when the problem occurred? Please include all relevant products, for example: operating system as well as related products.
- Has the problem happened before, or is this an isolated problem?
- What steps led to the failure?
- Can the problem be recreated? If so, what steps are required?
- Have any changes been made to the system? (workload, hardware, network or software)
- Were any messages or other diagnostic information produced? If yes, what were they?
- It is helpful to have the message number(s) of any messages received when you place the call for support or to document in the ETR.
- Define your technical problem statements in specific terms and provide the version and release level of the product(s) in question.

- Relevant diagnosis information

Note:

It is often necessary that our software support specialists analyze specific diagnostic information, such as storage dumps and traces, in order to resolve your problem. Gathering this information is often the most critical step in resolving your problem. Product specific diagnostic documentation can be very helpful in identifying what information is typically required to resolve problems. You should keep all problem data until the problem is resolved or until the data is successfully transmitted to IBM. The following are examples of problem data that the support center might ask you to provide:

Software support service checklist

- Any changes made to the system recently, preceding when the problem began occurring (for example, PTFs or new products installed or new hardware).
- Problem type (for example: abend, hang, loop)
- Search arguments
- Dump data, see [“Invoking IPCS as a background job” on page 337](#)
- Failing input request: macro, command, or statement
- SDWAVRA keys, lengths, and contents
- Offset of the failing instruction into the module or CSECT
- Accompanying messages: identifiers and texts
- Logrec report, if used
- All printed output and output data sets related to the problem
- Data on any related problems
- Module name and level
- Name and level of the operating system(s) with a list of program temporary fixes (PTF) applied at the time of the problem and all installation modifications, exits, and products with other than Class A service
- Other problem data developed while using the diagnosis book for the component, subsystem, or program
- Severity level

You need to assign a severity level to the problem when you report it, so you need to understand the business impact of the problem you are reporting. A description of the severity levels is in the following table.

Table 24. Severity levels and examples

Severity	Definition
1	Critical Impact/System Down: Business critical software component is inoperable or critical interface has failed. This indicates you are unable to use the program resulting in a critical impact on operations. This condition requires an immediate solution.
2	Significant impact: A software component is severely restricted in its use, causing significant business impact. This indicates the program is usable but is severely limited.
3	Moderate impact; A noncritical software component is malfunctioning, causing moderate business impact. This indicates the program is usable with less significant features.
4	Minimal impact; A noncritical software component is malfunctioning, causing minimal impact, or a nontechnical request is made.

- Mention the following items that apply to your situation:
 - If you are under business deadline pressure
 - When you are available (for example, when you will be able to work with IBM Software Support)
 - Where you can be reached
 - A knowledgeable alternate contact with whom IBM can speak
 - Other open problems (PMRs/Incidents) with IBM regarding this service request
 - If you are participating in an early support program (ESP)
 - If you have researched this situation prior to calling IBM and have detailed information or documentation to provide for the problem.

Automatic problem reporting

Parts of the system automatically report the need for service to IBM; for example, the central processor complex (CPC) reports problems directly to IBM. If the system contains a Hardware Management Console (HMC), you should be aware that problems in the Sysplex Timer and in direct access storage devices (DASD) might be automatically reported, even though the problems are recorded by MVS:

- MVS captures information about the problems and creates the following logrec records:
 - ETR record: For problems in the Sysplex Timer
 - DASD-SIM record: For problems in DASD
- For a unique Sysplex Timer or DASD error, HMC creates a problem record (PMR) in RETAIN to notify IBM that service is needed.

Invoking IPCS as a background job

Before calling IBM, format the dump using the IPCS subcommand recommended in the appropriate procedure. Some IPCS subcommands take time to run for a large dump and transferring the dump to IBM is often time consuming. You can start an IPCS session before calling IBM. Then, during the call, the output can be browsed as needed. See [“Step for invoking IPCS as a background job” on page 337](#).

Sometimes you might discover it is easier to create a second dump with a subset of the original dump. Send the second dump to IBM for initial diagnosis. For example, sometimes only the address spaces of the suspected problem job are necessary to diagnose system hangs. Of course, this is not always the case, so if you do send a subset of the dump to IBM, do not delete the original dump.

Step for invoking IPCS as a background job

Use the IPCS COPYDUMP subcommand to reduce the size of a very large dump, such as a stand-alone dump. An initial review of ASIDs 1-10 and others that are known to be involved in the problem, greatly reduces the size and transfer time of the dataset to be sent to the support center for initial diagnosis.

Use COPYDUMP to extract the problem address spaces that you want to analyze. COPYDUMP always includes address spaces 1 through 4 in the new dump data set, as well as any data spaces associated with the address spaces. Use the LISTDUMP subcommand to see the address spaces available in a dump data set.

The IPCS COPYDUMP subcommand can copy a single unformatted dump from one data set to another. Use the following example JCL to guide your creation of a batch job that invokes IPCS as a background job. This job opens the dump and extracts the desired ASIDs, using the IPCS COPYDUMP command, and saves the result in another data set.

```
//RLWG JOB '796634,?,S=I', 'RL WRIGHT', MSGLEVEL=(2,1),
// CLASS=2, NOTIFY=RLW, MSGCLASS=H
//IKJEFT01 EXEC PGM=IKJEFT01, REGION=4096K, DYNAMNBR=50
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
IPCS DEFER NOPARM
COPYDUMP INDSNAME('IPCS1.SYSOS18.SV03235') +
  OUTDSNAME('IPCS1.SYSOS18.SV03235.SUBSET') +
  ASIDLIST(1:10,15) SPACE(15000 10000) +
  NOCONFIRM
END
```

Figure 62. JCL for invoking IPCS as a background job

The options on the example IPCS command do the following:

DEFER

Defers the use of a dump directory. In this example COPYDUMP had no need for the dump directory.

NOPARM

Indicates not to use an IPCSPRnn parmlib member. This eliminates the allocation of the problem and data set directories named by IPCSPR00 on the production system.

For a complete list of IPCS options, see [z/OS MVS IPCS Commands](#).

Chapter 28. Problem diagnostic worksheet

Use this worksheet when calling IBM Technical Support to help you resolve your problem.

Table 25. What is the impact of your problem?

Impact

-
1. Critical business impact
 2. Significant business impact
 3. Some business impact
 4. Minimal business impact
-

Is there a system outage? If yes, how many systems are affected:

Is the problem repetitive? Can you recreate the problem?

Number of occurrences:

Details:

Table 26. How is your system configured?

System environment and level

CP model and serial number:

z/OS level:

How many systems are involved with the problem?

- LPAR
 - VM
 - NATIVE
 - Sysplex
-

What other hardware devices are involved?

Table 27. What are the external symptoms

External symptoms

- Coded system wait state
 - System hung or partitioned from sysplex
 - Loop or high system overhead
 - Loop or high CP usage by job
 - Job/subsystem/application/function failure
 - Job/subsystem/application/function hang
 - Output incorrect or missing
 - Performance or slowdown
 - Error message issued
-

Details:

Table 28. What symptom information did you collect?

Symptoms extracted from diagnostic information

Dump title:

ABEND code(s):

Wait state code:

Message ID(s):

Module name(s) and rmid:

Component ID(s):

Other:

Table 29. Which type of documentation did you obtain?

Documentation obtained

Dump produced:

- SYSABEND, SYSUDUMP, CEEDump (formatted dump)
 - SLIP, Console, SVC, SYSDUMP, TDUMP (unformatted dump)
 - SADUMP
-

Joblog, SYSLOG, OPERLOG or other:

EREP report of SYS1.LOGREC:

GTF data set:

CTRACE data set:

Other:

Table 30. What recovery actions did you attempt?

Recovery actions

- Program terminated
 - Job canceled
 - Job restarted
 - Job forced
 - Device taken offline
 - Restart key on HMC selected
 - System partitioned or re-ipld
 - Sysplex restarted
-

Other:

Appendix A. Accessibility

Accessible publications for this product are offered through [IBM Knowledge Center \(www.ibm.com/support/knowledgecenter/SSLTBW/welcome\)](http://www.ibm.com/support/knowledgecenter/SSLTBW/welcome).

If you experience difficulty with the accessibility of any z/OS information, send a detailed message to the [Contact z/OS web page \(www.ibm.com/systems/z/os/zos/webqs.html\)](http://www.ibm.com/systems/z/os/zos/webqs.html) or use the following mailing address.

IBM Corporation
Attention: MHVRCFS Reader Comments
Department H6MA, Building 707
2455 South Road
Poughkeepsie, NY 12601-5400
United States

Accessibility features

Accessibility features help users who have physical disabilities such as restricted mobility or limited vision use software products successfully. The accessibility features in z/OS can help users do the following tasks:

- Run assistive technology such as screen readers and screen magnifier software.
- Operate specific or equivalent features by using the keyboard.
- Customize display attributes such as color, contrast, and font size.

Consult assistive technologies

Assistive technology products such as screen readers function with the user interfaces found in z/OS. Consult the product information for the specific assistive technology product that is used to access z/OS interfaces.

Keyboard navigation of the user interface

You can access z/OS user interfaces with TSO/E or ISPF. The following information describes how to use TSO/E and ISPF, including the use of keyboard shortcuts and function keys (PF keys). Each guide includes the default settings for the PF keys.

- [*z/OS TSO/E Primer*](#)
- [*z/OS TSO/E User's Guide*](#)
- [*z/OS ISPF User's Guide Vol I*](#)

Dotted decimal syntax diagrams

Syntax diagrams are provided in dotted decimal format for users who access IBM Knowledge Center with a screen reader. In dotted decimal format, each syntax element is written on a separate line. If two or more syntax elements are always present together (or always absent together), they can appear on the same line because they are considered a single compound syntax element.

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that the screen reader is set to read out punctuation. All the syntax elements that have the same dotted decimal number (for example, all the syntax elements that have the number 3.1) are mutually exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, your syntax can include either USERID or SYSTEMID, but not both.

The dotted decimal numbering level denotes the level of nesting. For example, if a syntax element with dotted decimal number 3 is followed by a series of syntax elements with dotted decimal number 3.1, all the syntax elements numbered 3.1 are subordinate to the syntax element numbered 3.

Certain words and symbols are used next to the dotted decimal numbers to add information about the syntax elements. Occasionally, these words and symbols might occur at the beginning of the element itself. For ease of identification, if the word or symbol is a part of the syntax element, it is preceded by the backslash (\) character. The * symbol is placed next to a dotted decimal number to indicate that the syntax element repeats. For example, syntax element *FILE with dotted decimal number 3 is given the format 3 * FILE. Format 3* FILE indicates that syntax element FILE repeats. Format 3* * FILE indicates that syntax element * FILE repeats.

Characters such as commas, which are used to separate a string of syntax elements, are shown in the syntax just before the items they separate. These characters can appear on the same line as each item, or on a separate line with the same dotted decimal number as the relevant items. The line can also show another symbol to provide information about the syntax elements. For example, the lines 5.1*, 5.1 LASTRUN, and 5.1 DELETE mean that if you use more than one of the LASTRUN and DELETE syntax elements, the elements must be separated by a comma. If no separator is given, assume that you use a blank to separate each syntax element.

If a syntax element is preceded by the % symbol, it indicates a reference that is defined elsewhere. The string that follows the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 means that you must refer to separate syntax fragment OP1.

The following symbols are used next to the dotted decimal numbers.

? indicates an optional syntax element

The question mark (?) symbol indicates an optional syntax element. A dotted decimal number followed by the question mark symbol (?) indicates that all the syntax elements with a corresponding dotted decimal number, and any subordinate syntax elements, are optional. If there is only one syntax element with a dotted decimal number, the ? symbol is displayed on the same line as the syntax element, (for example 5? NOTIFY). If there is more than one syntax element with a dotted decimal number, the ? symbol is displayed on a line by itself, followed by the syntax elements that are optional. For example, if you hear the lines 5 ?, 5 NOTIFY, and 5 UPDATE, you know that the syntax elements NOTIFY and UPDATE are optional. That is, you can choose one or none of them. The ? symbol is equivalent to a bypass line in a railroad diagram.

! indicates a default syntax element

The exclamation mark (!) symbol indicates a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicate that the syntax element is the default option for all syntax elements that share the same dotted decimal number. Only one of the syntax elements that share the dotted decimal number can specify the ! symbol. For example, if you hear the lines 2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the default option for the FILE keyword. In the example, if you include the FILE keyword, but do not specify an option, the default option KEEP is applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, the default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1! (KEEP), and 2.1.1 (DELETE), the default option KEEP applies only to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.

*** indicates an optional syntax element that is repeatable**

The asterisk or glyph (*) symbol indicates a syntax element that can be repeated zero or more times. A dotted decimal number followed by the * symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be repeated. For example, if you hear the line 5.1* data area, you know that you can include one data area, more than one data area, or no data area. If you

hear the lines 3* , 3 HOST, 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

Notes:

1. If a dotted decimal number has an asterisk (*) next to it and there is only one item with that dotted decimal number, you can repeat that same item more than once.
2. If a dotted decimal number has an asterisk next to it and several items have that dotted decimal number, you can use more than one item from the list, but you cannot use the items more than once each. In the previous example, you can write HOST STATE, but you cannot write HOST HOST.
3. The * symbol is equivalent to a loopback line in a railroad syntax diagram.

+ indicates a syntax element that must be included

The plus (+) symbol indicates a syntax element that must be included at least once. A dotted decimal number followed by the + symbol indicates that the syntax element must be included one or more times. That is, it must be included at least once and can be repeated. For example, if you hear the line 6.1+ data area, you must include at least one data area. If you hear the lines 2+, 2 HOST, and 2 STATE, you know that you must include HOST, STATE, or both. Similar to the * symbol, the + symbol can repeat a particular item if it is the only item with that dotted decimal number. The + symbol, like the * symbol, is equivalent to a loopback line in a railroad syntax diagram.

Notices

This information was developed for products and services that are offered in the USA or elsewhere.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
United States of America*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

This information could include missing, incorrect, or broken hyperlinks. Hyperlinks are maintained in only the HTML plug-in output for the Knowledge Centers. Use of hyperlinks in other output formats of this information is at your own risk.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
Site Counsel
2455 South Road*

Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or

reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's name, email address, phone number, or other personally identifiable information for purposes of enhanced user usability and single sign-on configuration. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at ibm.com/privacy and IBM's Online Privacy Statement at ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at ibm.com/software/info/product-privacy.

Policy for unsupported hardware

Various z/OS elements, such as DFSMS, JES2, JES3, and MVS, contain code that supports specific hardware servers or devices. In some cases, this device-related element support remains in the product even after the hardware devices pass their announced End of Service date. z/OS may continue to service element code; however, it will not provide service related to unsupported hardware devices. Software problems related to these devices will not be accepted for service, and current service activity will cease if a problem is determined to be associated with out-of-support devices. In such cases, fixes will not be issued.

Minimum supported hardware

The minimum supported hardware for z/OS releases identified in z/OS announcements can subsequently change when service for particular servers or devices is withdrawn. Likewise, the levels of other software products supported on a particular release of z/OS are subject to the service support lifecycle of those products. Therefore, z/OS and its product publications (for example, panels, samples, messages, and product documentation) can include references to hardware and software that is no longer supported.

- For information about software support lifecycle, see: [IBM Lifecycle Support for z/OS \(www.ibm.com/software/support/systemsz/lifecycle\)](http://www.ibm.com/software/support/systemsz/lifecycle)
- For information about currently-supported IBM hardware, contact your IBM representative.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at [Copyright and Trademark information \(www.ibm.com/legal/copytrade.shtml\)](http://www.ibm.com/legal/copytrade.shtml).

Index

Special Characters

\$HASP263 [274](#)

A

abend

- finding the module [187](#)
- flowchart [182](#)
- hardware-detected [179](#)
- message text [180](#)
- obtaining the reason code [183](#)
- overview [179](#)
- searching databases [190](#)
- software-detected [179](#)
- symptoms [179](#)

abend analysis

- process [182](#)

abend code

- obtaining [183](#)
- steps for obtaining [183](#)
- VERBEXIT MTRACE [186](#)
- where to find [184](#)

abend example

- determining the dump type [195](#)
- dump title [186](#)
- finding the abend code [184](#), [195](#)
- finding the job name [193](#)
- finding the program name [193](#)
- free-format search argument [190](#)
- identifying the module prefix [189](#)
- obtaining the dump title [187](#)
- RSMDATA SUMMARY report [197](#)
- searching for multiple problems [194](#)
- SELECT output [193](#)
- STATUS CPU [188](#)
- STATUS FAILDATA [184](#)
- STATUS SYSTEM [195](#)
- STATUS WORKSHEET [187](#)
- SUMMARY TCBERROR [194](#)
- SYSTRACE output [191](#), [192](#)
- variable recording area [197](#)
- VERBEXIT MTRACE output [195](#)
- VERBX LOGDATA [185](#)
- VERBX MTRACE [186](#)
- VERBX SYMPTOM [186](#)

abend task

- abend code, obtaining
 - steps [183](#)
- component, identifying
 - steps [187](#)
- data, collecting
 - steps [194](#)
- data, gathering
 - steps [191](#)
- databases, searching
 - steps [190](#)

abend task (*continued*)

- logrec, collecting
 - steps [193](#)
- messages, collecting
 - steps [193](#), [194](#)
- module, identifying
 - steps [187](#)
- trace data, gathering
 - steps [191](#)

ABENDOF4

- RLS [318](#)

abnormal wait [222](#)

accessibility

- contact IBM [343](#)
- features [343](#)

address

- where to find [187](#)

address space

- status [218](#)

AIRH187E

- tracked jobs exception report [118](#), [163](#)

AMATERSE [26](#)

Assist On-Site [26](#)

assistive technologies [343](#)

ATR248E

- RRS hang symptom [287](#)

ATR249E

- RRS hang symptom [287](#)

authorization

- UPDATE [88](#)

automatic problem reporting [337](#)

automatic spin loop [234](#)

automation

- COMMNDxx for HZR [36](#)
- Runtime Diagnostics [36](#)

B

best practices

- stand-alone dump [23](#)

BLOCKER

- Runtime Diagnostics [41](#)

bpxmtext

- example [16](#)

C

catalog

- commands
 - problem determination [259](#)

catalog component

- problem determination [259](#)

CBFORMAT CVT

- example [14](#)

check

- common storage [95](#)
- enqueue request rate [108](#)

- check (*continued*)
 - JES spool usage [118](#)
 - Logrec dump rate [125](#)
 - message arrival rate [133](#)
 - private storage exhaustion [147](#)
 - SMF arrival rate [163](#)
- checkpoint contention
 - JES2 [274](#)
- checkpoint reconfiguration [271](#)
- checks
 - RRS [285](#)
- coded wait state
 - guidance for diagnosis [181](#)
- collecting data
 - performance problem [248](#)
 - using JES2 commands [250](#)
- collecting documentation
 - RRS [285](#)
- command
 - no system response [225](#)
- commands
 - differences [83](#)
 - for PFA [83](#)
 - IBM Health Checker for z/OS [83](#)
 - pfa,display [85](#)
 - pfa,modify [88](#)
 - Predictive Failure Analysis (PFA) [85](#), [88](#)
 - Runtime Diagnostics dump [53](#)
 - update [88](#)
- common storage
 - output report [95](#)
- communication
 - console
 - locked out [225](#)
- component
 - diagnosis [257](#)
- component diagnosis
 - RRS [257](#)
- component owner
 - identifying [5](#)
- compressing
 - dump data sets [26](#)
- console
 - disabled loop
 - locked out [225](#)
 - locked out [225](#)
- contact
 - z/OS [343](#)
- contention
 - JES2 [271](#)
 - Runtime Diagnostics [41](#), [199](#)
- control interval (CI)
 - splits [25](#)
- COPYDUMP
 - syntax [25](#)
 - using [25](#)
- coupling facility
 - volatile [274](#)
- CPU analysis
 - Runtime Diagnostics [41](#)
- CPU usage
 - analysis
 - Runtime Diagnostics [35](#)
- creating PFA directories

- creating PFA directories (*continued*)
 - example [81](#)
- critical messages
 - analysis
 - Runtime Diagnostics [35](#)
- cross-system coupling facility (XCF) [321](#)
- cross-system extended services (XES) [321](#)
- CSA + SQA
 - output [95](#)
- CSECT [229](#)
- CTIRRSxx
 - creating [286](#)
- CTRACE
 - RRS [286](#)

D

- DAE [10](#)
- data set allocation [239](#)
- data set output
 - diagnosing [239](#)
- database
 - for problem reporting [12](#)
- databases
 - abend symptoms [190](#)
 - problem reporting [8](#)
- debug options
 - HZR [54](#)
 - SVC dump [54](#)
 - target system home [54](#)
 - target system not equal to home [54](#)
- description
 - GRS [263](#)
 - RRS [285](#)
 - XCF [321](#)
 - XES [321](#)
- detection Loop [41](#)
- determining a loop [41](#)
- Determining hardcopy settings [31](#)
- diagnose
 - component problem [257](#)
- diagnosing
 - job or subsystem hang
 - steps [215](#)
- diagnosis
 - best practices [23](#)
 - diagnosis [199](#)
 - hang [199](#)
 - information for elements [333](#)
 - loop [199](#), [225](#)
 - output problem [237](#)
 - overview [xix](#)
 - performance problem [247](#)
 - Runtime Diagnostics [199](#)
 - starting [xix](#)
 - steps [4](#)
 - synopsis [xix](#)
 - system hang [199](#)
 - wait [199](#)
- diagnosis data
 - gathering [5](#)
- diagnostic worksheet [208](#)
- Diagnostics report [41](#)
- DISABLED FOR PER [228](#)

- disabled loop
 - diagnosing [229](#)
- disabled wait [200](#)
- disabled wait state [207](#)
- DISPLAY
 - GRS,CONTENTION [259](#)
 - RRS [285](#)
 - STATUS [85](#)
 - system logger [307](#)
 - XES and XCF [321](#)
- DISPLAY command
 - examples [85](#)
- documentation
 - where to find [21](#)
- dummy wait [200](#)
- dump
 - for loop [228](#)
 - RRS [285](#)
 - RRS and system logger [285](#)
 - Runtime Diagnostics [53](#)
- dump directory
 - defining [24](#)
- dump title
 - obtain the title [186](#)
- dumps
 - compressing [26](#)
- dynamic exits [240](#)

E

- enabled loop
 - diagnosing [230](#)
- enabled wait [200](#)
- enabling
 - Runtime Diagnostics [36](#)
- encrypting dump data [26](#)
- ENQ
 - analysis
 - Runtime Diagnostics [35](#)
 - contention
 - Runtime Diagnostics [41](#)
 - determining contention [41](#)
- ENQ BLOCKER Runtime Diagnostics [41](#)
- ENQ WAITER Runtime Diagnostics [41](#)
- enqueue request rate
 - high [108](#)
- ENQUEUE request rate
 - prediction report [108](#)
- event HIGHCPU [31](#)
- events
 - Runtime Diagnostics [40](#)
- example
 - ANALYZE RESOURCE [218](#)
 - ASCB control block [218](#)
 - bpxmtext [16](#)
 - CBFORMAT srb [220](#)
 - creating PFA directories [81](#)
 - finding the abend address [189](#)
 - finding the abend code [186](#)
 - free-format search argument [190](#)
 - hang
 - abend during NIP [204](#)
 - CBFORMAT CVT [200](#)
 - IEE844W [209](#)

- example (*continued*)
 - hang (*continued*)
 - STATUS CPU [210](#)
 - identifying control [222](#)
 - identifying the module [189](#)
 - IEA794I [5](#)
 - IEA911E [5](#)
 - IEA995I [180](#)
 - IPCS STATUS WORKSHEET [14](#)
 - ISMF ABEND PANEL [180](#)
 - ISPF abend panel [181](#)
 - loop
 - VERBEXIT LOGDATA output [235](#)
 - VERBEXIT LOGDATA Output [233](#)
 - VRA [233](#)
 - LSE [222](#)
 - message indicating an abend [180](#)
 - mount point [82](#)
 - mount points [81](#)
 - MVS Diagnostic Worksheet [234](#)
 - program check interruptions [188](#)
 - program name [193](#)
 - RB [221](#)
 - RSMDATA SUMMARY report [197](#)
 - Runtime Diagnostics
 - debugging [31](#)
 - SELECT output [193](#)
 - spin loop messages [232](#)
 - STATUS CPU REGISTERS [188](#)
 - subsystem hang [215](#)
 - SUMMARY FORMAT [221](#)
 - SUMMARY FORMAT JOBNAME [218](#)
 - SUMMARY TCBERROR [194](#)
 - SVC dump request [180](#)
 - SVC dumps [5](#)
 - symbolic link (symlink) [81](#)
 - SYSTEM status [195](#)
 - SYSTEM TRACE TABLE [230](#)
 - TSO/E abend message [181](#)
 - VERBEXIT MTRACE output [195](#)
 - VERBEXIT SYMPTOM [186](#)
 - VERBX SRMDATA report [223](#)
 - VRA [197](#)
 - zFS attributes [81](#)
- example high CPU [56](#)
- examples
 - DISPLAY [85](#)
 - UPDATE [88](#)
- exception handling
 - Predictive Failure Analysis (PFA) [75](#)
- excessive spin
 - diagnosing [232](#)
- excessive spin loop
 - steps for diagnosing [233](#)
 - time out [226](#)
- EXSPATxx [226](#)
- external symptoms [5](#)
- extracting
 - search argument [8](#)

F

- failure message for Runtime Diagnostics
 - log stream [40](#)

- fast tape [24](#)
- feedback [xxi](#)
- Flash
 - URL [21](#)
- flowchart
 - abend analysis [182](#)
 - job or subsystem hang [216](#)
 - loop [227](#)
- FMID (function management identifier)
 - for operating system [14](#)
- format
 - for symptoms [9](#)
- free-format search
 - building [8](#)

G

- GRS
 - problem determination [263](#)

H

- hang
 - dispatchable TCB or SRB [222](#)
 - during IPL
 - steps [203](#)
 - job not running [222](#)
 - logrec, gathering
 - steps [212](#)
 - messages, gathering
 - steps [212](#)
 - overview [199](#)
 - problem description, collecting
 - steps [202](#)
 - problem reporting, searching
 - steps [210](#)
 - symptom [199](#)
 - system trace activity [223](#)
 - wait states during IPL [203](#)
- hang example
 - ABEND issued [205](#)
 - ANALYZE EXCEPTION [206](#)
 - CBFORMAT [207](#)
 - COPYCAPD [212](#)
 - COUPLE SYSPLEX EXCEPTION [210](#)
 - CP disabled coded wait [202](#)
 - CP no work wait [202](#)
 - IPL statistics [203](#)
 - IPLDATA STATUS [203](#)
 - MVS diagnostic worksheet [208](#)
 - NIP statistics [203](#)
 - SELECT ALL [203](#)
 - STATUS CPU [201](#), [205](#)
 - wait [204](#), [205](#)
 - wait states during IPL [203](#)
- hardware-detected
 - abend [179](#)
- health exception event
 - Runtime Diagnostics [47](#)
- HIGHCPU event [31](#)
- HIGHCPU report [41](#)
- Hot Topics [21](#)
- how PFA works [73](#)

- HZR
 - Runtime Diagnostics [40](#)
- HZR integration with PFA [75](#)
- HZR0201I [38](#)
- HZROUT [36](#)

I

- IBM Health Checker for z/OS
 - for PFA commands [83](#)
 - RRS [285](#)
- IBM Service
 - instructions Runtime Diagnostics [53](#)
- IBMLink [12](#)
- ICH408I
 - insufficient access authority
 - Runtime Diagnostics [40](#)
- identifying
 - system problems [5](#)
- IEA611I [6](#)
- IEADMCxx
 - RRS [285](#)
- IEE844W [209](#)
- incomplete output [239](#)
- Index trap
 - VSAM [311](#)
- information
 - where to find more [21](#)
- input
 - analyzing [239](#)
- install
 - Predictive Failure Analysis (PFA) [78](#)
- installation exits
 - analyzing [240](#)
- installing
 - PFA [77](#)
 - shared file system [80](#)
- introduction
 - problem determination [3](#)
- IPCS
 - COPYDUMP [25](#)
 - performance [25](#)
- IPCS STATUS WORKSHEET
 - example [14](#)
- IPCS subcommand
 - by problem [195](#)
- IPL
 - hang [203](#)
 - information [13](#)
- ISGGWAIT [222](#)
- ISMF
 - ABEND PANEL [180](#)
- ISPF
 - abend panel [181](#)
- IXC246E [42](#)
- IXC431I [42](#)
- IXC467I [42](#)
- IXC585E [42](#)
- IXG231I
 - log stream connectivity [40](#)
- IXL013I [42](#)
- IXL030I [42](#)

J

- Java
 - level [77](#), [82](#)
- Java 7.1
 - example [82](#)
- Java path
 - example [82](#)
 - updating [82](#)
- JCL
 - for dumps [27](#)
- JES spool usage
 - prediction report [118](#)
- JES2
 - checkpoint contention [271](#), [274](#)
 - checkpoint reconfiguration [274](#)
 - collecting data [250](#)
 - contention [271](#)
 - high CPU [271](#)
 - performance [271](#)
 - problem determination [271](#)
- job
 - dispatchability [218](#)
- job hang
 - overview [215](#)
 - symptom [215](#)
- job name
 - finding [218](#)
 - Runtime Diagnostics [36](#)
- job or subsystem hang
 - gathering data
 - step [217](#)
- job querying
 - analysis
 - Runtime Diagnostics [35](#)
- job status
 - determine [217](#)
 - waiting for resources [217](#)

K

- keyboard
 - navigation [343](#)
 - PF keys [343](#)
 - shortcut keys [343](#)
- keyword list [10](#)
- known problem [10](#)

L

- latch contention
 - Runtime Diagnostics [41](#)
- local lock suspension
 - Runtime Diagnostics [41](#)
- Local lock suspension [41](#)
- lock
 - Runtime Diagnostics [41](#)
- lock suspension, Local [41](#)
- log stream
 - failure for Runtime Diagnostics [40](#)
 - failure message [40](#)
- logrec
 - analyzing for loop [235](#)

- Logrec
 - abend example [181](#)
- LOGREC
 - prediction report [125](#)
- loop
 - analyzing interactively [231](#)
 - analyzing logrec
 - steps [235](#)
 - analyzing printed output [231](#)
 - determining an enabled loop [228](#)
 - determining type [228](#)
 - diagnosis [225](#)
 - disabled
 - definition [225](#)
 - diagnosing [229](#)
 - symptom [225](#)
 - enabled
 - definition [225](#)
 - diagnosing [230](#)
 - symptom [225](#)
 - flowchart [227](#)
 - gathering data [227](#)
 - Runtime Diagnostics [41](#)
 - spin
 - definition [225](#)
 - symptom [226](#)
 - steps for diagnosis [226](#)
 - TCB
 - Runtime Diagnostics [41](#)
- Loop detection [41](#)
- loop example
 - determining a loop [233](#)
 - DUMP command [228](#)
 - finding an excessive spin loop [234](#)
 - finding CSECT [229](#)
 - finding the dump title [234](#)
 - finding the module name [231](#)
 - reading system trace [230](#)
 - STATUS CPU [228](#), [229](#), [231](#)
 - SYSTRACE ALL [230](#)
 - viewing the FRR stack [235](#)
 - WHERE [230](#)
- LOOP report [41](#)

M

- managing PFA
 - command [83](#)
- MAS [274](#)
- message analysis
 - Runtime Diagnostics [41](#)
- messages
 - HZR [40](#)
 - Runtime Diagnostics [40](#)
 - Runtime Diagnostics analysis
 - BPX prefix [57](#)
 - IEA prefix [59](#)
 - IEE prefix [60](#)
 - IOS prefix [61](#)
 - IRA prefix [61](#)
 - IXC prefix [62](#), [65](#)
 - IXL prefix [68](#)
- migration
 - PFA [76](#)

missing
 data set output [239](#)
MODIFY
 CATALOG [259](#)
module
 identify [187](#)
 identifying [5](#)
multiprocessor [14](#)
MVS diagnostic worksheet [208](#)
MVS Diagnostic Worksheet [187](#)
MVS symptom [10](#)

N

navigation
 keyboard [343](#)
no work wait [200](#)

O

OnTop [26](#)
operating system
 level
 determining [13](#)
 name
 determining [13](#)
OPERLOG
 assigning hardcopy [31](#)
 overview [31](#)
 Runtime Diagnostics [37](#), [40](#)
 setting up
 steps [31](#)
output
 analyzing [239](#)
 using the message prefix [241](#)
output problem
 data set allocation [239](#)
 determine path [243](#)
 diagnosis [237](#)
 installation exit [240](#)
 teleprocessing [243](#)
output problem example
 identifying the component [241](#)
 message prefix for identification [241](#)
 obtaining the program name [241](#)
 viewing an abend [241](#)
output problems
 logrec [243](#)
 messages [243](#)
 reporting [244](#)
output task
 additional data, gathering
 steps [243](#)
 data set allocation, analyzing
 steps [239](#)
 data, collecting
 overview [238](#)
 step [239](#)
 inputs and outputs, analyzing
 steps [239](#)
 installation exits, analyzing
 steps [240](#)
 problem reporting, searching

output task (*continued*)
 problem reporting, searching (*continued*)
 steps [242](#)
 problems, analyzing
 overview [237](#)
 problems, reporting
 information [244](#)
 program or component, identifying
 steps [241](#)
overview
 enabled wait state [199](#)
 hang [199](#)
 job hang [215](#)
 job or subsystem hang [215](#)
 loop [225](#)
 performance problem [247](#)
 problem resolution [3](#)

P

PDSE
 specific problems [278](#)
PDSE component
 problem determination [277](#)
performance
 cause of decline [250](#)
 JES2 [271](#)
performance problem
 diagnosis [247](#)
 overview [247](#)
performance problem example
 displaying the contention report [251](#)
 obtaining the program name [252](#)
 requesting an SVC dump [254](#)
 viewing a SLIP summary [249](#)
 viewing held jobs [250](#)
 viewing printer status [251](#)
 viewing task status [250](#)
performance task
 additional data, collecting
 steps [248](#)
 additional data, gathering
 steps [253](#)
 additional JES2 data, gathering
 steps [250](#)
 dump, analyzing
 steps [254](#)
 problem databases, reporting
 steps [252](#)
 problems, reporting
 steps [254](#)
PFA
 how it works [73](#)
 Java level [77](#), [82](#)
 managing [83](#)
 migration [76](#)
PFA DISPLAY command [85](#)
PFA MODIFY command [88](#)
PFA_COMMON_STORAGE_USAGE
 check [95](#)
 creep [95](#)
 output report [95](#)
 spike definition [95](#)
PFA_ENQUEUE_REQUEST_RATE

PFA_ENQUEUE_REQUEST_RATE (*continued*)
 check [108](#)
 parameters [108](#)

PFA_JES_SPOOL_USAGE
 check [118](#)
 parameters [122](#)

PFA_LOGREC_ARRIVAL_RATE
 check [125](#)

PFA_MESSAGE_ARRIVAL_RATE
 check [133](#)
 parameters [133](#)

PFA_PRIVATE_STORAGE_EXHAUSTION
 check [147](#)

PFA_SMF_ARRIVAL_RATE
 check [163](#)
 parameters [172](#)

PR/SM (Processor Resource/Systems Manager)
 determine use [14](#)

prediction file
 PFA_COMMON_STORAGE_USAGE
 .csv format [105](#)
 .csv values [105](#)
 values [131](#)

prediction report
 ENQUEUE request rate [108](#)
 JES spool usage [118](#)
 LOGREC [125](#)
 private storage exhaustion [157](#)
 SMF arrival rate [163](#)

predictive failure
 common storage usage [95](#)
 ENQUEUE request rate check [108](#)
 JES spool usage [118](#)
 LOGREC arrival rate check [125](#)
 message arrival rate check [133](#)
 private storage exhaustion check [147](#)
 SMF arrival rate check [163](#)

Predictive Failure Analysis (PFA)
 components [74](#)
 directory [78](#)
 example
 typical check [74](#)
 exception handling [74](#), [75](#)
 how to update [88](#)
 install [78](#)
 installing [77](#)
 overview [73](#)
 pfa,display [85](#)
 pfa,modify [88](#)
 prerequisites [77](#)
 Runtime Diagnostics reports [75](#)
 storage [78](#)
 UID [78](#)

private storage exhaustion
 prediction report [157](#)

problem
 categories [6](#)

problem data
 checklist [337](#)
 collecting [3](#)
 report to IBM [337](#)

problem determination
 ASIDs consuming excessive storage below 2GB [304](#)
 catalog commands [259](#)
 problem determination (*continued*)
 damaged or broken catalogs [260](#)
 GRS [263](#)
 hang in the Catalog address space [259](#)
 hang in VSAM record management code [313](#)
 HANG/WAIT in RLS/TVS [319](#)
 IGW038A message [279](#)
 introduction to [3](#)
 Invalid Track Format error on a volume [302](#)
 issues opening, closing or extending VSAM data sets
[314](#)
 JES2 commands [271](#)
 loop in VSAM record management code [313](#)
 message ANTX5104E [303](#)
 microcode issue impacting concurrent copy [305](#)
 mismatch of storage control sessions [304](#)
 No Record Found error on a volume [302](#)
 PDSE ABEND 0F4 failures [278](#)
 PDSE data set corruption [281](#)
 RRS [285](#)
 RRS commands [285](#)
 Share Control Datasets not specified [320](#)
 slow performance for catalog address space [261](#)
 SMSPDSE address space failure [283](#)
 SMSPDSE1 address space failure [283](#)
 SMSVSAM will not start up [320](#)
 StorageControlTimeout [303](#)
 system logger [307](#)
 system logger commands [307](#)
 unable to determine PPRC paths [301](#)
 unexpected return codes from VSAM record
 management [314](#)
 XCF [321](#)
 XES [321](#)
 XES and XCF commands [321](#)
 problem determination, catalog component [259](#)
 problem determination, JES2 [271](#)
 problem determination, PDSE component [277](#)
 problem determination, System Data Mover (SDM)
 component [301](#)
 problem determination, VSAM component [311](#)
 problem determination, VSAM record-level sharing (RLS)
 component [317](#)
 problem diagnostic worksheet [339](#)
 problem indicators
 overview [6](#)
 table [6](#)
 problem reporting
 additional data for output [243](#)
 automatic [337](#)
 databases [8](#)
 output problems [242](#)
 to IBM [337](#)
 problem reporting database
 search [12](#)
 problem resolution
 overview [3](#)
 problem source identification
 definition [3](#)
 introduction to [3](#)
 problem type
 indicators [6](#)
 processing
 UPDATE [88](#)

- processor
 - example [14](#)
- processors
 - How many are active? [14](#)
- profile
 - Runtime Diagnostics [36](#)
- program check [204](#)
- program check interruptions
 - example [188](#)
- program name [193](#)
- PSI
 - definition [3](#)

Q

- qualified success
 - Runtime Diagnostics [40](#)

R

- real storage manager
 - usage [197](#)
- reason codes
 - determining meaning [16](#)
 - z/OS UNIX [16](#)
- reconfiguration
 - checkpoint [274](#)
- remote access [26](#)
- report HIGHCPU [41](#)
- report LOOP [41](#)
- reporting example
 - finding the system level [14](#)
 - number of active processors [14](#)
- reports
 - Runtime Diagnostics [40](#)
- requirements Runtime Diagnostics [36](#)
- resource contention
 - performance [251](#)
- Resource Recovery Services (RRS) [285](#)
- restrictions
 - UPDATE [88](#)
- RETAIN symptom [9](#)
- RMF
 - analyze output [251](#)
- RRS
 - collecting documentation [285](#)
 - commands
 - problem determination [285](#)
 - CTRACE options [286](#)
 - description [285](#)
 - dump [285](#)
 - problem determination [285](#)
- RSMDATA SUMMARY report
 - example [197](#)
- RSVSF [26](#)
- Runtime Diagnostics
 - contention [41](#)
 - CPU analysis [41](#)
 - CPU usage [35](#)
 - critical WTO message analysis
 - example [41](#)
 - ENQ contention report [41](#)
 - event

- Runtime Diagnostics (*continued*)
 - event (*continued*)
 - example [40](#)
 - explanation [40](#)
 - F, HZR [53](#)
 - guidelines [36](#)
 - high CPU and lock message [41](#)
 - HZR
 - user ID [36](#)
 - HZRINIT [36](#)
 - JES2
 - health exception event [47](#)
 - JOBNAME parameter [36](#)
 - local lock suspension [41](#)
 - log stream [40](#)
 - loop [41](#)
 - message
 - log stream failure [40](#)
 - message analysis [41](#)
 - OPERLOG [40](#)
 - output [40](#)
 - overview [35](#)
 - partial analysis
 - system [38](#)
 - Predictive Failure Analysis (PFA) [75](#)
 - profile [36](#)
 - requirements [36](#)
 - security [36](#)
 - server
 - health exception event [47](#)
 - STARTED class [36](#)
 - status message [40](#)
 - SYS1.PROCLIB data set [36](#)
 - target system
 - example [38](#)
 - tasks [35](#)
 - TCB loop [35](#)
 - understanding messages [38](#)
 - understanding report messages [40](#)
 - using [36](#)
 - z/OS UNIX
 - latch contention [41](#)
- Runtime Diagnostics LOOP and HIGHCPU report [41](#)
- Runtime Diagnostics proc
 - SYS1.PROCLIB [36](#)

S

- SDM
 - specific problems [301](#)
- search
 - problem reporting database [12](#)
- search argument
 - developing [8](#)
 - stand-alone dump [8](#)
 - SVC dump [8](#)
 - symptoms [10](#)
- searching
 - loop problems [235](#)
 - problem reporting databases [8](#)
- sending to IBM
 - reader comments [xxi](#)
- serialization summary [46](#)
- shared file system

- shared file system (*continued*)
 - installing [80](#)
- shortcut keys [343](#)
- SLIP
 - setting for abend [194](#)
- SMF
 - analyze output [251](#)
- SMF arrival rate
 - prediction report [163](#)
- soft failures
 - description [73](#)
- software-detected
 - abend [179](#)
- spin loop
 - diagnosing [232](#)
- SRB status
 - examining [220](#)
- stand-alone dump
 - automation [27](#)
 - creating [24](#)
 - dump directory [24](#)
 - example [28](#)
 - fast tape [24](#)
 - guidelines [24](#)
 - merging [25](#)
 - performance [23](#)
 - planning [23](#)
 - sample jobs [27](#)
 - using IPCS [25](#)
- standardized symptom keyword list [10](#)
- STARTED class Runtime Diagnostics [36](#)
- STATUS SYSTEM [195](#)
- status update [209](#)
- STATUS WORKSHEET
 - example [187](#)
- storage
 - exhaustion [147](#)
 - Predictive Failure Analysis (PFA) [78](#)
- subsystem hang
 - symptom [215](#)
- subsystem Runtime Diagnostics [36](#)
- summary of changes
 - V2R2 [xxiii](#)
- supervised learning
 - configure [91](#)
- suspension, Local lock [41](#)
- SVC
 - dump message [180](#)
- SVC dump
 - job or subsystem hang [217](#)
- symbolic link (symlink)
 - example [81](#)
- symlink [81](#)
- symptom
 - classifications [3](#)
 - disabled loop [225](#)
 - enabled loop [225](#)
 - free format [9](#)
 - hang [199](#)
 - in search argument [10](#)
 - MVS format [9](#)
 - RETAIN format [9](#)
 - spin loop [226](#)
 - wait [199](#)

- symptoms
 - abend [179](#)
 - search argument [8](#)
 - selecting [8](#)
- syntax
 - HZR [53](#)
- sysplex
 - system partitioned [209](#)
- System Data Mover (SDM) component
 - problem determination [301](#)
- system hang
 - diagnosis [199](#)
- system level
 - determining [14](#)
- system logger
 - commands
 - problem determination [307](#)
- System logger
 - problem determination [307](#)
- system partitioned [209](#)
- system problem
 - identification [5](#)
 - what to collect [5](#)
- system trace
 - analyzing for abend [192](#)
- SYSTEM TRACE TABLE [230](#)
- SYSTRACE
 - wait entries [223](#)

T

- task
 - abend
 - overview [179](#)
- tasks
 - hang, diagnosing
 - steps [199](#), [201](#)
 - loop, diagnosing
 - steps [41](#)
 - resource contention, checking
 - steps [251](#)
- TB status
 - examining [220](#)
- TCB loop
 - Runtime Diagnostics [35](#)
- tech notes
 - URL [21](#)
- techdocs
 - URL [21](#)
- Technical help database [12](#)
- technotes
 - URL [21](#)
- Technotes [12](#)
- testing
 - dump operations [26](#)
- tip
 - query trace [5](#)
 - routing information [4](#)
- tools
 - IBM Health Checker for z/OS [3](#)
- trace
 - query [5](#)
- trace data
 - gathering for abends [191](#)

trademarks [350](#)
TRSMAIN [26](#)
TSO/E abend message
example [181](#)

U

Understanding PFA commands [83](#)
uniprocessor [14](#)
UPDATE
examples [88](#)
EXCLUDED_JOBS
example [88](#)
usage
UPDATE [88](#)
user ID
Runtime Diagnostics [36](#)
user interface
ISPF [343](#)
TSO/E [343](#)

V

VERBEXIT MTRACE output [195](#)
VERBEXIT SYMPTOM
example [186](#)
subcommands [195](#)
volatile
coupling facility [274](#)
VSAM
Index trap [311](#)
problems [311](#)
VSAM component
problem determination [311](#)
VSAM record-level sharing (RLS)
problems [318](#)
VSAM record-level sharing (RLS) component
problem determination [317](#)

W

wait
coded [200](#)
disabled
non-restartable [200](#)
restartable [200](#)
enabled
causes [200](#)
definition [200](#)
dummy [200](#)
no work wait [200](#)
missing interrupt [200](#)
non-restartable [200](#)
restartable [200](#)
symptom [199](#)
wait state
coded
diagnosing [207](#)
WAITER
Runtime Diagnostics [41](#)
worksheet
reporting problems [339](#)

X

XCF
description [321](#)
problem determination [321](#)
XES
description [321](#)
problem determination [321](#)
XES and XCF
commands
problem determination [321](#)

Z

z/OS UNIX
reason codes [16](#)



SC23-6844-40

