




JES2 Debugging

Adam Nadel
anadel@us.ibm.com
IBM - Poughkeepsie, NY

Thursday, August 7, 2014
Session Number 16167



#SHAREorg



Copyright (c) 2014 by SHARE Inc. Except where otherwise noted, this work is licensed under <http://creativecommons.org/licenses/by-nc-sa/3.0/>



Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

- IBM®
- MVST™
- Redbooks®
- RETAIN®
- z/OS®
- zSeries®

The following are trademarks or registered trademarks of other companies.

- Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.
- All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM Business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

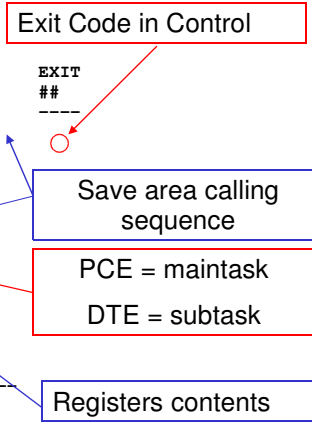


Get To Know Your Error With \$HASP088s



```

$HASP088 JES2 ABEND ANALYSIS
$HASP088 -----
$HASP088 FMID = HJE7780 LOAD MODULE = HASJES20
$HASP088 SUBSYS = JES2 z/OS1.13
$HASP088 DATE = 2013.343 TIME = 10.23.45
$HASP088 DESC = DISASTROUS ERROR AT LABEL KBLOBERR
$HASP088 MODULE      MODULE      OFFSET      SERVICE     ROUTINE
$HASP088 NAME        BASE        + OF CALL   LEVEL       CALLED
$HASP088 -----
$HASP088 HASPRAS     0003B480 + 0005E4   OA41318     *ERROR $DIS
$HASP088 HASPCKPT    1D8C44D0 + 0058F8   OA36155     $DISTERR
$HASP088 HASPCKPT    1D8C44D0 + 004BAE   OA36155     KBLEMPY
$HASP088 HASPCKPT    1D8C44D0 + 0005CE   OA36155     KBLOB
$HASP088 PCE = CKPT (1DA4E3B8)
$HASP088 R0 = E0000264 1DA4AFD4 00000000 1DBF18C8
$HASP088 R2 = 00000000 00000000 00000000 00000000
$HASP088 R4 = 00000000 1D8C9DC8 00000000 00000000
$HASP088 R6 = 00000F00 003A1700 00000000 1DA4AE70
$HASP088 R8 = 00000F00 00000000 00000000 00000000
$HASP088 R10 = 00000000 00000000 00000000 00007000
$HASP088 R12 = 00000000 0003B660 00000000 1DA4E3B8
$HASP088 R14 = 00000000 8003B9C0 00000000 0003C390
$HASP088 -----
    
```



Most JES2 abends will be accompanied by diagnostic \$HASP088 messages. These will be preceded by the JES2 message indicating whether the dump is catastrophic (\$HASP095) or disastrous (\$HASP096) and followed by message indicating what level of recovery has occurred otherwise termination options.

Depending on the type of error produced, the \$HASP088 messages may also contain other useful information such as the jobname (and jobid) being processed at time of error, the home/primary/secondary ASID at time of error (not guaranteed that JES2 is the primary ASID) etc.

The NETRV address space (JES2Snnn) has an equivalent version via \$HASP5088

Disastrous vs Catastrophic



\$HASP096 DISASTROUS ERROR AT SYMBOL **TIMERROR** IN CSECT HASPJQS

- Spool control block related
 - \$IOT, \$JCT, \$HDB etc
- Real I/O error reading from spool
 - IOS error details accompanying \$HASP064
 - Logical error associated with a spool control block
 - Control block does not match expectations
- Minimal Impact

label in JES2 code

Which is worse – Disastrous or Catastrophic errors? Since most forms of disastrous errors are logical errors in which part of control block does not match expectations, they are typically far less severe than catastrophic in terms of impact. In logical error cases, there is typically no loss of JES2 function and the impact of error confined to the JOB in-hand. Real I/O errors are far less common, but under those circumstances the impact would not necessarily be confined to a single JOB etc.

While less severe in terms of impact, disastrous errors can often be more difficult to debug – because it often entails reviewing the entire lifespan of the JOB(s) affected to understand what may have caused the spooled block to not match expectations. Did something prevent an IO from completing successfully such as an error/abend within the job itself? Or was there an disruption to JES2 overall (not a clean shutdown etc)?

What Is a CBIMPL4?

```
$HASP096 DISASTROUS ERROR AT SYMBOL CBIMPL4 IN CSECT HASPNUC,
MQTR=040000A1B90C, UNIT=A056, VOLSER=JES11
```

```
$HASP088 JES2 ABEND ANALYSIS
$HASP088 -----
$HASP088 FMID = HJE7790 LOAD MODULE = HASJES20
$HASP088 SUBSYS = JES2 z/OS 2.1
$HASP088 DATE = 2014.005 TIME = 09.20.19
$HASP088 DESC = DISASTROUS ERROR AT LABEL CBIMPL4
$HASP088 MODULE      MODULE      OFFSET      SERVICE  ROUTINE      EXIT
$HASP088 NAME        BASE        + OF CALL   LEVEL    CALLED       ##
$HASP088 -----
$HASP088 HASPRAS     00022E30   + 0005E4   OA37847  *ERROR $DIS
$HASP088 HASPNUC     00007000   + 0095A4   OA37654  $DISTERR
$HASP088 HASPTRAK    1A630EE0   + 000DC6   OA37847  $CBIOM
$HASP088 HASPTRAK    1A630EE0   + 0002DA   OA37847  PURSAF
$HASP088 HASPTRAK    1A630EE0   + 003432   OA37847  $PURGER
$HASP088 HASPTRAK    1A630EE0   + 002C54   OA37847  VIOTPRG
$HASP088 PCE = PURGE (1A84F0A0) JOB12345 ADAM1
...
```

CBIMPL4 is the most common of JES2 errors. It is a logical error (disastrous) in which JES2 is attempting to access a control block for job A (in above case JOB12345 – ADAM1) and instead reads in a block for job B. The buffer identifying job B can be found within the respective PCE save area chain (which we will cover shortly) or in the respective SYMREC indicating SPOOL TRACKGROUP RECOVERY.

If this is a single/isolated instance, then there is no cause for high alarm and we are looking for some kind of disruption within the lifespan of the job(s) in question that could have prevented an IO from completing. If this is a one of MANY errors of the same/similar nature, then would be greater concern as it could be a reflection of adverse impact to spool and/or checkpoint – such as accidentally starting JES2 with wrong spool or checkpoint volume(s) etc.

What Is a CBIMPL4?

- Recovery is confined to job identified
 - Job is purged
 - Track recovered by spool trackgroup reclamation (SNIFFER)
 - SYMREC produced
 - Can be controlled/expedited via
SPOOLDEF,GCRATE=NORMAL/FAST

```

COMPONENT ID: 5752SC1BH
COMPONENT RELEASE LEVEL: Z113
SERVICE RELEASE LEVEL: OA38671
DESCRIPTION OF FUNCTION: SPOOL TRACKGROUP RECOVERY
PROBLEM ID: SYMTAB SUBSYSTEM ID: JES2Z113
...

```

SYMREC type

MTTR/MQTR

```

FREE FORMAT COMPONENT INFORMATION:
KEY = 010D LENGTH = 000003 (0003)
+000 0304A4
KEY = 010E LENGTH = 000008 (0008)
+000 00000000 00000000
KEY = 010F LENGTH = 000008 (0008)
+000 80004F2A B6B46F7F
KEY = 0110 LENGTH = 000256 (0100)
+000 C8C4C240 ... ..

```

Control block contents
(HDB, IOT, etc)



Session 16167

The impact is that the affected job will be purged and the trackgroup in question (that contained residual data for a different job) will be temporarily be marked as not owned. Thereafter, the JES2 trackgroup reclamation PCE (aka SNIFFER) will run and clean up the track group, restoring it to the track group map for future reuse. SNIFFER defaults to NORMAL setting – it will cycle through all tracks within ~7 days. It can be increased to \$TSPoolDEF,GCRATE=FAST which causes SNIFFER to interrogate all tracks immediately (and after that it automatically returns to NORMAL rate).

Disastrous vs Catastrophic



\$HASP095 JES2 CATASTROPHIC ABEND. CODE = S0C4 (RC = 00000004)

- **CODE=ERROR**

- JES2 detected error condition
 - \$Knn – CKPT read/write errors – module HASPCKPT
 - \$Qnn – problem with job (JQE) – module HASPJQS
 - \$Jxx – problem with output (JOE) – module HASPJOS
- Error regs found in \$ERROR save area
- JES2 internal Ctraces useful in diagnosis

JES2 error condition \$nnn
-or- MVS ABEND

- **CODE=ABEND**

- MVS detected error (0C4, 878, B00, etc)
- JES2 maintask ESTAE gets control for recovery
- RTM2WA generated
- System trace table



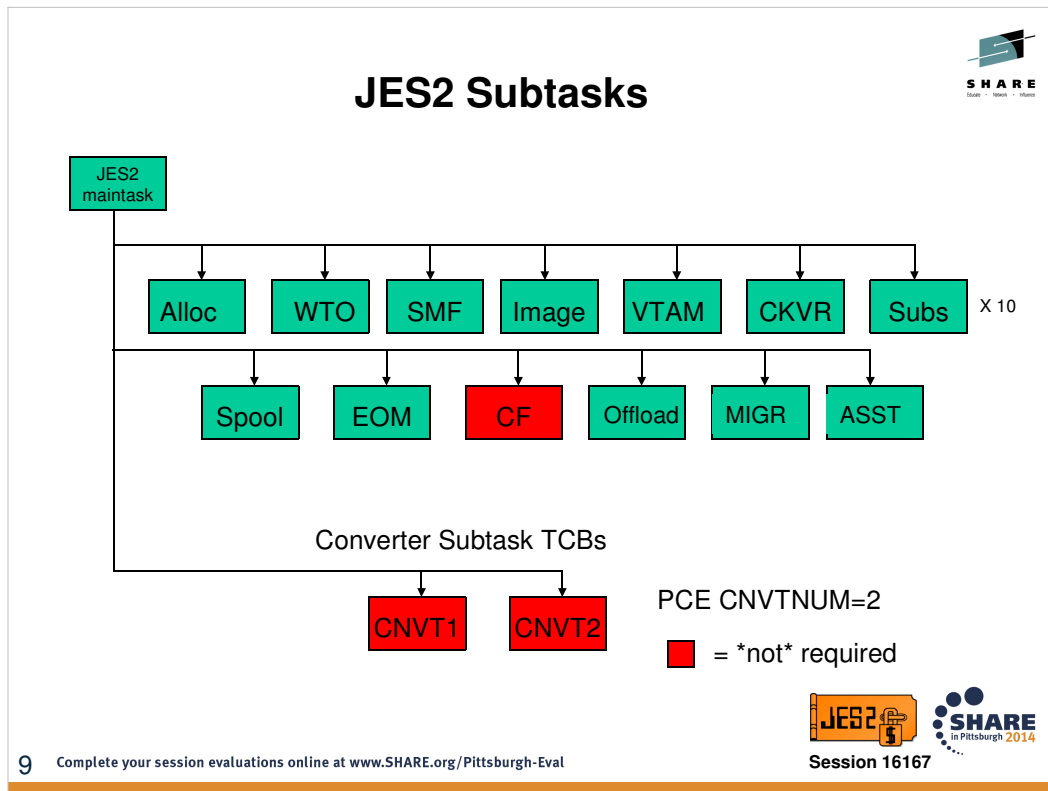
Catastrophic errors are unexpected, logically detected errors. They encompass both JES2 detected errors as well as general MVS abends encountered under JES2. For MVS abends, it is appropriate to approach their diagnosis as you would any other MVS type abend – using RTM2WA, systrace, SUMM FORMAT, etc.

For JES2 detected errors, there are diagnostics available within JES2 such as \$ERROR save area calling sequence and internal Ctraces.

PCE Recovery (or not)

- \$HASP098 Enter Termination Option – worst
 - Required PCE failed and could not be recovered
- \$HASP073 Recovery Successful – best
 - Normal processing resumes
 - May be confined to job in hand
- \$HASP068 Partial Recovery Successful – good enough?
 - PCE has terminated and will not run again
 - Processing continues without that PCE
 - How many PCEs remain of that type
 - Is function impacted
 - How can I recover PCE

Depending on the severity of the error, there are varying degrees of JES2 recovery. Partial recovery is intended to keep JES2 operating and stable and allow time to schedule hoststart/IPL at your nearest convenience to recover lost PCE. The type of impact may vary based on the specific type of PCE affected. A device PCE (PRT1) means the device will not function (may be critical). Other PCEs such as Sysout API (SAPI) interface may have far smaller impact depending on the number of PCEs defined. When JES2 terminates a PCE, it produces a message indicating the PCE has terminated and also how many of that type remain. An ended PCE will prevent a clean shutdown of JES2 and can be identified via \$DPCE(*),ENDED.



JES2 maintask does not like to MVS wait. JES2 creates separate subtask TCB's to invoke services that may result in an MVS wait. There are 14 different JES2 subtask types of which one is for conversion.

The number of converter subtasks corresponds to PCEDEF CVNTNUM parameter. The default is 2. The MVS converter is linked to in order to converter the JCL images.

The MVS converter also performs the PROC expansions.

Brief Summary of Subtask functions:

ALLOC- used to perform dynamic allocations

WTO – issues MVS WTO to put out JES2 messages

SMF – writes SMF records to SMF dataset

IMAGE – allocates and opens SYS1.IMAGELIB (only done during JES2 startup)

VTAM – used to open or close VTAM ACB

CKVR – checkpoint versions and WLM sampling

SUBS – general purpose subtasks most often used for performing SAF calls (there are 10 of these TCB's)

SPOOL – handles spool volume allocations etc

EOM – z4 and up. Processes \$SJB placed on the EOM queue for end of memory SSI processing

CF – used when CKPT is on coupling facility to interface with the CF to read/write CKPT data

OFFLOAD – used to perform I/O etc to offload datasets

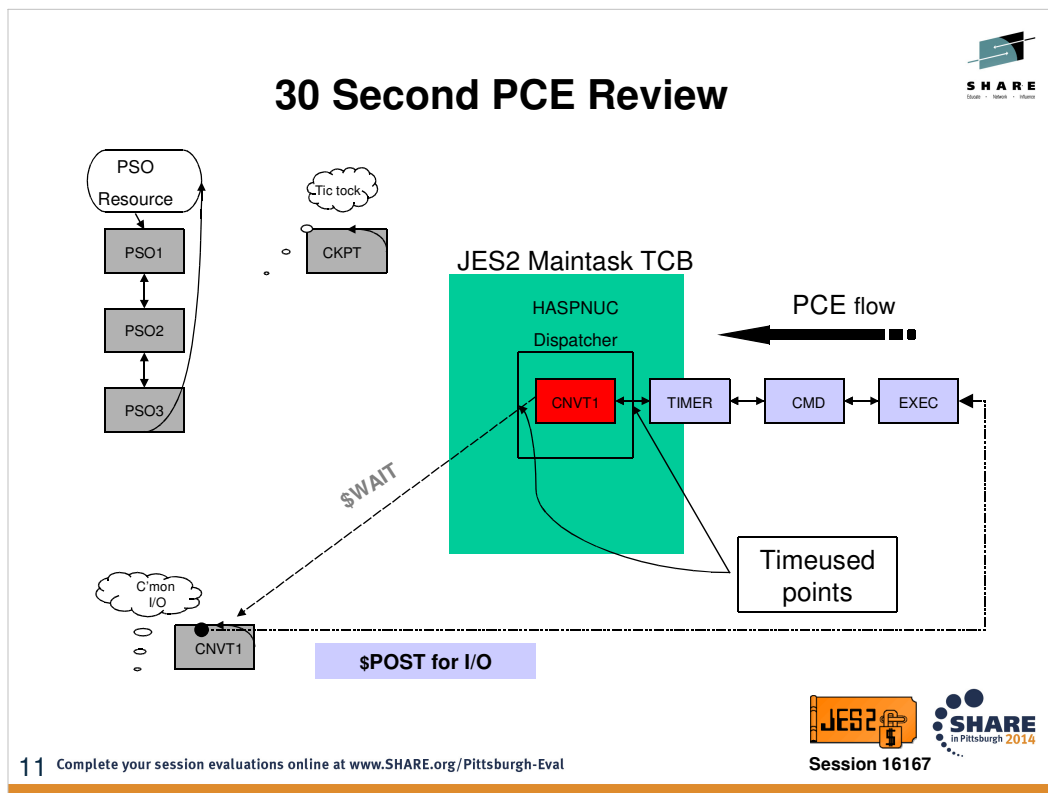
MIGR and ASST– involved in spool migration processing

JES2 Subtask Recovery (or not)



- \$HASP078 Subtask failed
 - Indicates the failing JES2 subtask
 - Always MVS abend code
- \$HASP095 error \$Z03 issued if a required subtask cannot be recovered
- Potential function loss when subtask terminates

Impact may vary depending on the type of subtask that was impacted. Barring the *not* required subtasks, the overall health JES2 is typically going to be in trouble if it loses a subtask. For instance, losing a CNVT subtask may not be critical if you have 10 defined. However, the loss of the CKPT version subtask may prevent the updating of checkpoint versions (copies) – which could affect respective exploiters like SDSF. The loss of the VTAM subtask would impact SNA communications etc.



The Processor Control Element (PCE) represents an instance of a “process” running under the control of JES2 main task – each PCE is a dispatchable unit of work controlled by the JES2 dispatcher. “Process” is synonymous for JES2 service – such as EXEC, CMD, SAPI, PSO, CNVT, etc. There are one or more PCEs for each process, some dictated by PCEDEF statement definitions.

Above illustrates basic flow of PCE’s being dispatched by JES2 maintask: When a PCE has work to do, it is moved into the ready queue (awaiting their turn to be dispatched). When the PCE’s runs through the dispatcher its entry and exit into and out of it is framed with TIMEUSED macros. This allows JES2 to capture CPU time information that shows up in internal traces and PERFDATA.

JES2 Component Panels



```
IPCS JES2 Format Trace Debug
----- JES2 Component Data Analysis
Option ==> 2;6;S JES2;
Enter JES2 name ==> JES2

Select desired option for JES2 dump:
 1 JES2 base display
 2 JES2 job control blocks
 3 JES2 job output control blocks
 4 JES2 devices
 5 JES2 processors
 6 JES2 subtasks
 7 JES2 control blocks
 8 JES2 NJE/RJE control blocks
 9 JES2 MAS member data
10 JES2 checkpoint control blocks
11 JES2 BERT control blocks
12 JES2 monitor data

These panels are for
JES2 FMID: HJE7780
Service level: 0
```

Issued from IPCS primary menu

JES2 subsystem name

12 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Session 16167

The following slides assume that the JES2 IPCS Support modules (SHASPARM, SHASMIG, SHASPNL0) have been loaded into the requisite concatenations on your system. For specific information on JES2 IPCS Support modules, please refer to z/OS JES2 Diagnosis manual (chapter: Using IPCS for Diagnosis).

From this panel, you have various formatting options based on what you are attempting to debug; however, option “1 JES2 base display” is often the best place to begin diagnosis as it surfaces an abundance of pertinent information. Most of the options place you in another panel with prompting fields for additional information. The panels do have help screens to assist in navigation and data entry.

The subsystem name defaults to JES2, but is an overtype field for alternative JES2 subsystem names (JESA etc)

Also on this panel (but not illustrated above), is option “101 – Select JES2 control blocks for non-JES2 address space”. These panels may be useful for JES2-related abends that occur within a user address space – allowing formatting of JES2 control blocks that reside in common storage

JES2 Base Display



*** JES2 Base Display ***

Subsystem "JES2" is in address space ASID(X'002D')
Dump for JES2 release="z/OS 2.1", Product level=43, Service level=0
(pointed to by SSCTSUSE); CVTPRODI=HBB7790
Maximum extended region size for "JES2" is 1,395M (per LDAELIM)

*** **WARNING**: ASCBDSP1=80

System set non-dispatchable and this ASCB is not exempt (per
ASCBSSND bit)

*** **WARNING**: DEBUG BERT=NO specified (per \$DBGBERT bit off in
\$DEBGOPS in \$HCT)

*** **WARNING**: \$EVENT(s) exist (PCBEVNTF=0 in \$PERFCB)

*** **NOTICE**: \$QSUSE is NOT in effect (per \$QSONDA bit in \$STATUS in
\$HCT)

*** **NOTICE**: SPOOLDEF FENCE=ACTIVE=YES in effect (per CCTSMVFN
bit in CCTSTUS in \$HCCT)



13 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

Session 16167

This is the top portion of the Base Display. It includes the JES2 product information along with WARNING, NOTICE, and ERROR alert messages. You will always find alert messages in a dump, so their presence alone is not indicative of any particular problem. However, WARNING and NOTICE messages draw attention to key pieces of information that will assist the debugger in understanding the state of JES2 at time of dump. Some examples are:

- JES2 is abending/abended
- JES2 is quiesced via \$P or \$PXEQ
- \$ZAPJOB has been issued
- \$EVENTS exist (produced by JES2MON)
- JES2 ASCB is not dispatchable
- etc

ERROR alerts often indicate that certain areas are not able to be formatted. These may be rather innocuous and simply reflect that some storage area was not dumped, or can shed insight into control block overlay scenarios etc.

JES2 Base Display



```

$PCE: 1AEB46E0
+0000 PCEEYE... PCE
+0000 PSVID.... PCE          PSVPREV.. 00000000    PSVNEXT.. 2B1C8A28
+00EC RSV..... 00000000
      ***** INTERNAL READER *****
+0000 RDWTEMP.. C2404040 40404040
+0460          40404040 40404040 40404040 40404040
    
```

Current PCE

```

$PSV: 2B1C8A28
+0000 PSVID.... SAVE          PSVPREV.. 1AEB46E0    PSVNEXT.. 1AEB46E0
+000C PSVR14... 800EE48A    PSVR15... 000F99EE    PSVR0.... 1AEBAB50
+005A RSV..... 00000000 0000
+0060 PSVSTCK.. CD05E208    AF3D690C
04/18/2014 09:33:32.008406
Routine name: RERROR
      000F9A06: HASPRDR (X'000ED1C8') + X'0000C83E'
Address routine called from (assuming normal linkage):
      000EE48A: HASPRDR (X'000ED1C8') + X'000012C2'
1 $PSV(s) processed
    
```

\$Save Area /
Calling Sequence

```

$DCT: 1A1BF570
+0000 DCTID.... DCT          DCTPCE... 1AEB46E0    DCTSTAT.. 90
+0028 DCTDEVN.. INTRDR     DCTUCB... 00000000    DCTTOKA.. 1A1CF5F0
+010E RIDFLAG3. 00          RIDRSV3.. 00
    
```

\$DCT from
PCEDCT field

```

** $JQE Address=1BCC49B0, Offset=0000E998, Index=000256
** $JQX Address=1CFD1C40, Offset=00008C28
** Address of first $BERT for this $JQA is 20ECEB98
** BERT lock is not held
** NOTE: $JQA incomplete, all fields past label JQABERT are zero
$JQA: (Composite of $JQE and $JQX)
      JQE.....
+0000 JQEPRIO.. FF JQETYPE.. 20 JQEJOBNO. 1ED2
    
```

\$JQE from
PCEJQE field



This section is towards the bottom of the Base Display panels. I have omitted the middle section which also displays the \$HCT and \$HCCT control blocks. All of this information can also be formatted via other JES2 panels (such as PCE panels, job display panels, subtask panels, etc).

The PCE Save Area (\$PSV) can be thought of as the JES2 version of a linkage stack – one entry produced per PCE to represent the state of processing as it issued a \$SAVE (but not yet issued the \$RETURN). Once the \$RETURN is issued, the PSV is dechained and available for reuse. It provides a lot of insight into the path leading up to the error (including register contents) and will match up to the calling sequence identified in the \$HASP088 messages.

It will also format and display other control blocks that are active/in-hand at time of error such as device blocks (device control table \$DCT), job blocks (job queue element \$JQE, output blocks (job output element \$JOE), etc.

Useful Commands & Module Background



```
Command ==> IP CBF 00091A0 STR($MODLOC)
***** TOP OF DATA *****
000091A0: HASPNUC (X'00007000') + X'000021A0' OA36155/UA68055
```

- HASCnnnn → JES2 module in **C**ommon storage
 - Maintenance hitting module typically requires WARMstart (IPL)
- HASPnnnn → JES2 module in **P**rivate storage
 - Maintenance hitting module typically requires HOTstart

```
Command ==> IP CBF 072E3050 STR($PCE)
***** TOP OF DATA *****
$PCE: 072E3050
+0000 PCEEYE... PCE
+0000 PSVID.... PCE      PSVPREV.. 00000000  PSVNEXT.. 072E3050
+0018 PSVR1.... 069CC230 PSVR2.... 069CC138  PSVR3.... 00003000
```

Browsing raw storage, JES2 module eyecatcher information is at the beginning of each module; however, the maintenance level information is at the end of the module. For this reason, is often very helpful to use the \$MODLOC formatter to verify if/where an address is in JES2 code. The formatters will also work within the user address space for common modules.

JES2 common modules HASCnnnn are primarily responsible for:

- SSI calls
- Extended status
- Sysout allocation / open / close / PUT / GET / POINT
- SAPI / PSO

The mainline recovery for common modules is HASCLINK

The second example shows formatting an address as a \$PCE. JES2 has formatters for many control blocks (\$JQE, \$DTE, \$JOE, etc), so it may be worthwhile to attempt a CBF against that respective block to assist in formatting (rather than dealing with raw storage).

JES2 Ctraces



```
----- CTRACE DISPLAY PARAMETERS -----  
COMMAND ==> 2;7;1;d  
System   ==>                               (System name or blank)  
Component ==> SYSJES2                       (Component name (required))  
Subnames ==> JOE
```

Issued from IPCS primary menu

- Component is SYSnnn
 - nnn = JES2 subsystem name (JES2, JESA, etc)
- Subnames
 - DISP
 - JQE
 - JOE
 - SAPI ****new**** (delivered via APAR OA43882)



The JES2 Ctraces are component traces that are always running internally. They are in-storage only and cannot be put out to external writer etc. The installation does not control the size of the trace, and they are rolling traces. These traces can be displayed via the IPCS component trace facility as displayed above. Alternatively, you can use the “TRACE” drop down menu from the JES2 primary panel (shown on slide 11).

There are four types of traces/subames: DISP, JOE, JQE, SAPI

JES2 DISP Ctrace



SYSA DISP 00000421 21:59:16.610981 Dispatch PCE

PCE Address->1AE8B638 Exit->00 JOB#/offset->00000000 00000000
Module/seq#->HASPPSO 01960000 Wait time->00000000 0027E5AD
\$POST type-->0000

Dispatch point

Time length PCE \$WAITed till Dispatch

PCE description:PROCESS SYSOUT PROCESSOR
\$WAIT Events: POST
\$WAIT Resource: PSO
\$WAIT Options:
\$POST Reason: Resource post

\$POST information

SYSA DISP 00000420 21:59:16.611114 PCE \$WAIT

PCE Address->1AE8B638 Exit->00 JOB#/offset->00001ED1 0000EC54
Module/seq#->HASPNUC 17000000 Run time->00000000 00000085
CPU time---->00000000 00000085

\$WAIT point

JOB# that PCE is working on

PCE description:PROCESS SYSOUT PROCESSOR
\$WAIT Events: IO
\$WAIT Options:

\$WAIT information

17 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Session 16167

The JES2 dispatcher rolling ctrace shows information on each PCE as it respectively enters/exits the JES2 dispatcher. It also will show when JES2 encounters an MVS WAIT.

Things to consider while reviewing DISP ctrace:

- Are there abnormally large time gaps between entries or large MVS waits?
- Are one (or more) specific PCE unexpectedly monopolizing the dispatching?
- Any PCEs appear to be looping?
- Any unusual \$WAIT conditions?
- Is an exit in control (related to any of the above)?

JES2 JQE Ctrace



```
-----  
SYSAS JQE 00000203 21:59:16.610444 $QMOD ← Macro traced  
  
PCE Address->1AE88148 Exit->00 JOB#/offset->00001ED1 0000EC54  
Original Queue->02 New Queue->01 Busy->00 Lock->01  
Artificial JQE ← $JQETYPE changed  
  
PCE description:OUTPUT PROCESSOR  
-----  
SYSAS JQE 0000020C 21:59:16.610458 $DOGJQE  
  
PCE Address->1AE88148 Exit->00 JOB#/offset->00001ED1 0000EC54  
Original Queue->01 New Queue->01 Busy->00 Lock->01  
  
PCE description:OUTPUT PROCESSOR  
-----  
SYSAS JQE 00000207 21:59:16.610478 $FREJLOK  
PCE Address->1AE88148 Exit->00 JOB#/offset->00001ED1 0000EC54  
Original Queue->01 New Queue->01 Busy->00 Lock->00  
Artificial JQE ← $JQEBUSY indicator  
  
PCE description:OUTPUT PROCESSOR  
-----
```

\$JQEBUSY indicator



Session 16167

18 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

The JQE rolling ctrace The above shows information about job state a job state changes – particularly the (un)busying of the job block, (un)locking of the job, and transitioning of job from queue-to-queue. The above case illustrates a job moving from the output queue (being serviced by a OUTPUT PCE), to the hardcopy queue. As part of this process we can observe the joblock is obtained and then freed.

Things to consider while reviewing JQE ctrace:

- Are there any large gaps in processing?
- Are you looking for a specific job?
- Are you looking to see that a particular queue/phase is being serviced (backlog?)?
- Is an exit in control (related to any of the above)?

JES2 JOE Ctrace



```
SYSA JOE 00000319 21:59:16.611063 $#BUSY ← Macro traced
PCE Address->1AE8B638 Exit->00 Job number->00001ED1 JOE offset->00003FC8
Original Class->D3 New Class->D3 Busy->01 Type->80
PCE description:PROCESS SYSOUT PROCESSOR ← $JOEBUSY indicator
-----
SYSA JOE 00000312 21:59:16.826295 $#REM
PCE Address->1AE86638 Exit->00 Job number->00000000 JOE offset->00003FC8
Original Class->D3 New Class->D3 Busy->00 Type->C0
PCE description:PROCESS SYSOUT PROCESSOR ← Offset into JOT
-----
SYSA JOE 0000031A 21:59:18.625218 $#GET
PCE Address->2B1D9180 Exit->00 Job number->00001ED1 JOE offset->00004510
Original Class->D8 New Class->D8 Busy->01 Type->80
PCE description:NJE SYSOUT TRANSMITTER
```




Session 16167

19 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

The JOE rolling ctrace The above shows information about job output state a job state changes – particularly the (un)busying of the output block, (un)locking of the output, and transitioning of output from queue-to-queue. The above case illustrates two pieces of output within the same job being processed. This is evident by the two different offsets into the JOT along with each JOE being within different classes. The first piece of output is processed by PSO and purged (noted by JOETYPE=C0). The second piece of output is then selected by a NJE sysout transmitter Lnn.STn.

Things to consider while reviewing JOE ctrace:

- Are there any large gaps in processing?
- Are you looking for a specific output?
- Are you looking to see that a particular queue/phase is being serviced (backlog)?
- Is an exit in control (related to any of the above)?





JES2 SAPI Ctrace

```

SY1      SAPI      05000033  20:54:22.058909  Bulk Modify
SAPI name--> ADAM1.      PCE Address->0BD82200
Job number-->00000039  JOE offset->00000410  SAPID->00001000
SSS2SELx 1->E4 2->00 3->00 4->00 5->00 6->00
CPU Time---->00000000  0000009D  Run Time----->00000000  000000A1
$QSUSE Time->00000000  0000009F  Elapsed Time->00000000  0003E516
$#GET Time-->00000000  0000001C  $RQUE Time---->00000000  00000013
Application name
I/O Count--->00000000
A JOE was returned
SAPID assigned a JOE
-----
SY1      SAPI      05000031  16:15:23.924338  Put/Get call
SAPI name--> ARCHIVE2    PCE Address->0BD72B28
Job number-->00000140  JOE offset->00000478  SAPID->00001000
SSS2SELx 1->08 2->00 3->00 4->00 5->00 6->00
CPU Time---->00000000  0000009E  Run Time----->00000000  000000E6
$QSUSE Time->00000000  000000EB  Elapsed Time->00000000  00000207
$#GET Time-->00000000  00000016  $RQUE Time---->00000000  00000049
Time obtaining the JOE
I/O Count--->00000001
A JOE was returned
SAPID assigned a JOE
-----

```

20 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

Session 16167

The SAPI rolling ctrace captures the last 2000 SAPI requests at time of dump. It identifies the requestor, the type of request and breakdown of the overhead of the request. Much of the same information is captured in the JES2 id traces (ID=28,29 for SAPI, ID=20 for \$#GET) – depending on the duration of the SAPI problem and timeliness of dump, the internal Ctrace may be sufficient. JES2 id traces are more appropriate for capturing data across a wider timeframe.

Things to consider while reviewing SAPI ctrace:

- Are there any unexpected SAPI applications involved in the processing of job output?
- Any SAPI application appear to be looping/processing same output?
- Any requests taking long in duration (wall clock or CPU time)?
- What is the specific request type and criteria (related to any of the above)?

Merging Ctraces



```
----- MERGE SPECIFICATION -----  
Command ==> 2;7;5;c ← Issued from IPCS primary menu  
  
Enter/verify trace specifications for this MERGE operation.  
In the left column, type C/G/R: ( C = CTRACE G = GTFTRACE R = reset )  
  
C/G/R--Trace Invocation Parameters -----  
  
1. CTRACE COMP(SYSJES2) SUB((SAPI)) FULL DSNAME('ADAM.SYSA.DUMP1')  
2. CTRACE COMP(SYSJES2) SUB((JQE)) FULL DSNAME('ADAM.SYSA.DUMP1')  
3. CTRACE COMP(SYSJES2) SUB((DISP)) FULL DSNAME('ADAM.SYSB.DUMP2')  
4. CTRACE COMP(SYSJES2) SUB((JOE)) FULL DSNAME('ADAM.SYSB.DUMP2')  
  
ENTER = continue MERGE definition.  
END/PF3 = return to the MERGE GLOBAL PARAMETERS panel.  
S = start MERGE.
```



Sometimes it may be beneficial to merge ctraces in order to gain a better understanding of the processing flow. This can be achieved via the MERGE trace facility. Note that you can specify multiple datasets which can be handy if processing for a job spanned different JES2 MAS members.

PERFDATA



- Most useful in diagnosing JES2 performance problems
- Undocumented command(s) that capture various JES2 performance statistics
- Proper PERFDATA Collection
 - Reset statistics via \$TPERFDATA(*),RESET
 - Wait interval that covers problem timeframe (10-15 minutes)
 - Display all statistics via \$DPERFDATA(*)
- Gather several samples
 - Good vs Bad timeframe ?

22 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



When diagnosing JES2 performance concerns, it is of paramount importance that PERFDATA is collected correctly. PERFDATA statistics are always running/accumulating, so resetting the statistics is always the first step to collecting an accurate sample. Without the reset, you may be investigating a problem in which JES2 CPU utilization drastically increased over a 15 minute timeframe using statistics covering an interval of 20+ days! In those cases, the data is considered oversaturated in that there is no way to discern what actually happened in those specific 15 minutes. Typically Level 2 recommends gathering samples in 10-15 minute increments (and you can always gather multiple/back-to-back samples).

It is also worth consideration to occasionally capture sample(s) when processing is good/normal. These can become handy to compare and contrast if JES2 performance drastically changes for the worse. A debugger can view the samples side by side to observe differences in PCE processing, job throughput, checkpoint cycling, etc.

PERFDATA

- \$T PERFDATA(*),RESET – resets performance data
- \$D PERFDATA(INITSTAT) – JES2 initialization stats
- \$D PERFDATA(QSUSE) – PCE \$QSUSE summary
- \$D PERFDATA(PCESTAT) – 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(WSSTAT) – work selection (\$#GET & \$#POST) stats
- \$D PERFDATA(*) –all of the above

The \$TPERFDATA(*),RESET is absolutely essential to ensuring a healthy sample is gathered. Thereafter, the type of display requested may be dictated by the specific problem being investigated; although Level 2 most commonly asks for all data \$DPERFDATA(*). The WSSTAT option is a newly added section delivered In APAR OA43882. Disclaimer – the output of these displays can be rather abundant!

\$D PERFDATA(CKPTSTAT)



```
$HASP660 CKPT PERFORMANCE STATISTICS SYS1-INTERVAL=11:10:12.320961,  
$HASP660 AVGHOLD=0.318337,AVGDORM=45.305289,TOT$CKPT=3284,  
$HASP660 WRITE-4K=0,WRITE-CB=788,OPT$CKPT=2496,OPT4K=0,  
$HASP660 IO=R1,COUNT=875,AVGTIME=0.010943,  
$HASP660 IO=R2,COUNT=0,AVGTIME=0.000000,TOTAL4K=0,TOTALCB=118,  
$HASP660 IO=PW,COUNT=876,AVGTIME=0.004066,TOTAL4K=39,TOTALCB=0,  
$HASP660 IO=IW,COUNT=878,AVGTIME=0.003776,TOTAL4K=0,TOTALCB=670,  
$HASP660 IO=FW,COUNT=876,AVGTIME=0.003888,TOTAL4K=0,TOTALCB=118
```

AVGHOLD & AVGDORM = average HOLD and DORMANCY values

TOT\$CKPT = total number of \$CKPTs

AVGWAIT = average I/O times to CKPT



24 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

Session 16167

The CKPT statistics section will illustrate whether your CKPT cycling is what you would expect based on HOLD and DORMANCY settings as well as provide information about the relative health of CKPT I/Os. On each member in the MAS you can compare AVGHOLD and AVGDORM versus the HOLD and DORMANCY coded values to verify if JES2 is cycling the CKPT as expected. Large differences in these values may suggest additional tuning is needed based on workload distribution etc (eg should some members be favored more because it does most of the job submit? ... or archiving?)

Comparing TOT\$CKPT on each member of the MAS is a quick way to assess which members have the most checkpoint activity.

The AVGWAIT values associated with primary/intermediate/final write can be used to assess relative health of the I/Os. These times can vary (particularly depending on checkpoint placement on DASD vs coupling facility); however, the numbers should be consistent for I/O on the device.

\$D PERFDATA(CPUSTAT)



\$HASP660 CPU PERFORMANCE STATISTICS SYS1 - INTERVAL=14:49.926816,
\$HASP660 CPU=3.067221,

INTERVAL = length of time data has been accumulating

CPU = CPU time used by all of JES2 over that interval

\$HASP660 PCENAME=SPI, CPU%=6.16 , CPU=0.120145, TIME=0.151743,

\$HASP660 QSUSE_TIME=0.098023, IOCOUNT=599, CKPT_COUNT=28280,

PCENAME = name of the group of PCE's captured (see \$DPCE for details)

CPU% = percentage of total JES2 maintask time this subset of PCE's used

CPU = total CPU time used by this subset of PCE's

TIME = Wall clock time this subset of PCE's was disp

25 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Session 16167

This will show a breakdown of CPU utilization by PCE type. These numbers may change greatly/frequently based on configuration and workload. For instance, a member that is primarily used for archiving may demonstrate SAPI utilization far higher than a member used for communications/NJE (in which that member may show higher utilization for NET.SR etc). It is typically very reasonable to see CKPT PCE towards the top of this list.

Additional fields:

QSUSE_TIME = Wall clock time subset of PCE's ran when they acquired queue (\$QSUSE)

IOCOUNT = Total number of I/O's issued by this subset of PCE's

CKPT_COUNT = Number of \$CKPT's issued by this subset of PCE's

\$D PERFDATA(PCESTAT)



```
$HASP660 PCENAME=NET.SR,TIME=43.011997,CPU=40.892083,CPU%=7.63,  
$HASP660 QSUSE_TIME=0.277515,IOCOUNT=278356,CKPT_COUNT=7631,  
$HASP660 WAIT=IO,MOD=HASPNUC,SEQ=17000000  
$HASP660 COUNT=4066,AVGWAIT=0.001505,  
$HASP660 POST=IO,COUNT=4066,AVGWAIT=0.001505,  
$HASP660 WAIT=BUF,INHIBIT=NO,MOD=HASPNSR,SEQ=70272000  
$HASP660 COUNT=4982,AVGWAIT=0.001099,  
$HASP660 POST=RESOURCE,COUNT=4970,AVGWAIT=0.000878,  
$HASP660 POST=IO,COUNT=12,AVGWAIT=0.092677,  
$HASP660 WAIT=CKPT,MOD=HASPNUC,SEQ=28410000  
$HASP660 COUNT=221,AVGWAIT=0.342762,CMOD=HASPJQS,CSEQ=03330000,  
$HASP660 POST=RESOURCE,COUNT=221,AVGWAIT=0.342762,
```



Session 16167

26 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

This section starts with the CPU statistics (CPUSTAT), and then follows with a breakdown of activity of that PCE type. Specifically, it will breakdown the PCE \$WAITS and \$POSTs by type and count. The AVGWAIT for WAIT=CKPT help give insight into CKPT access time. All wait times include queue time for the PCE type – the time it is on the ready queue awaiting dispatch.

Fields:

WAIT = wait type(s) passed on the \$WAIT macro

MOD/SEQ = module and sequence number where \$WAIT was issued

COUNT = the number of \$WAITS (or \$POSTs) issued from this location

AVGWAIT = Average time the PCE spent at this location waiting

POST = Post type that woke the PCE up from this \$WAIT

PERFDATA - CPU Increase Example



Problem: JES2 CPU spike! \$DPERDATA for a 30 minute interval provided:

```
$HASP660 PCENAME=STAC,TIME=9:05.678529,CPU=5:58.996610,CPU%=95.67,  
$HASP660 QSUSE_TIME=9:01.582558,IOCOUNT=0,CKPT_COUNT=0,  
$HASP660 WAIT=CKPT,MOD=HASPNUC,SEQ=28410000  
$HASP660 COUNT=1653,AVGWAIT=0.756200,CMOD=HASPSTAC,CSEQ=13100000,  
$HASP660 POST=RESOURCE,COUNT=1653,AVGWAIT=0.756200,  
$HASP660 WAIT=STAC,INHIBIT=NO,MOD=HASPSTAC,SEQ=09900000  
$HASP660 COUNT=1646917,AVGWAIT=0.008434,
```

CPU% & PCENAME identify STAC (Status/Cancel) as likely culprit.

WAIT=STAC is the STAC PCE wait for work.

COUNT with STACNUM=2 on PCEDEF indicates 823,458 SSI requests were made with a relative rate of **AVGWAIT**.



The above demonstrates a scenario in which STAC PCE is monopolizing the JES2 activity – using 95% of the total CPU consumed by JES2. It is interesting to observe TIME vs CPU divergence. In the above case CPU is ~2/3 of the wall clock TIME – indicating JES2 is ready to run but is not getting CPU cycles. The CPU cycles it is getting are clearly funneling into STAC. It is also helpful to review the AVGWAIT associated with CKPT to understand checkpoint access is healthy. From here we would begin focusing on whether there was a loop within STAC processing or whether someone was continuously driving STAC requests – possibly requiring separate traces and/or dumps.

PERFDATA - Throughput Analysis



```
PCENAME=JQRP,TIME=0.903077,CPU=0.700405,CPU%=3.13,  
QSUSE_TIME=0.278596,IOCOUNT=2196,CKPT_COUNT=35317,  
WAIT=WORK,INHIBIT=NO,MOD=HASPJQS,SEQ=70910000  
COUNT=2061,AVGWAIT=0.329332,  
POST=IO,COUNT=936,AVGWAIT=0.000606,  
POST=$$POST,COUNT=1125,AVGWAIT=0.602831,  
WAIT=CKPT,INHIBIT=NO,MOD=HASPJQS,SEQ=70923300  
COUNT=3542,AVGWAIT=0.095666,  
POST=RESOURCE,COUNT=1528,AVGWAIT=0.188931,  
POST=IO,COUNT=947,AVGWAIT=0.000848,
```

IOCOUNT / 2 = number of jobs created during this interval

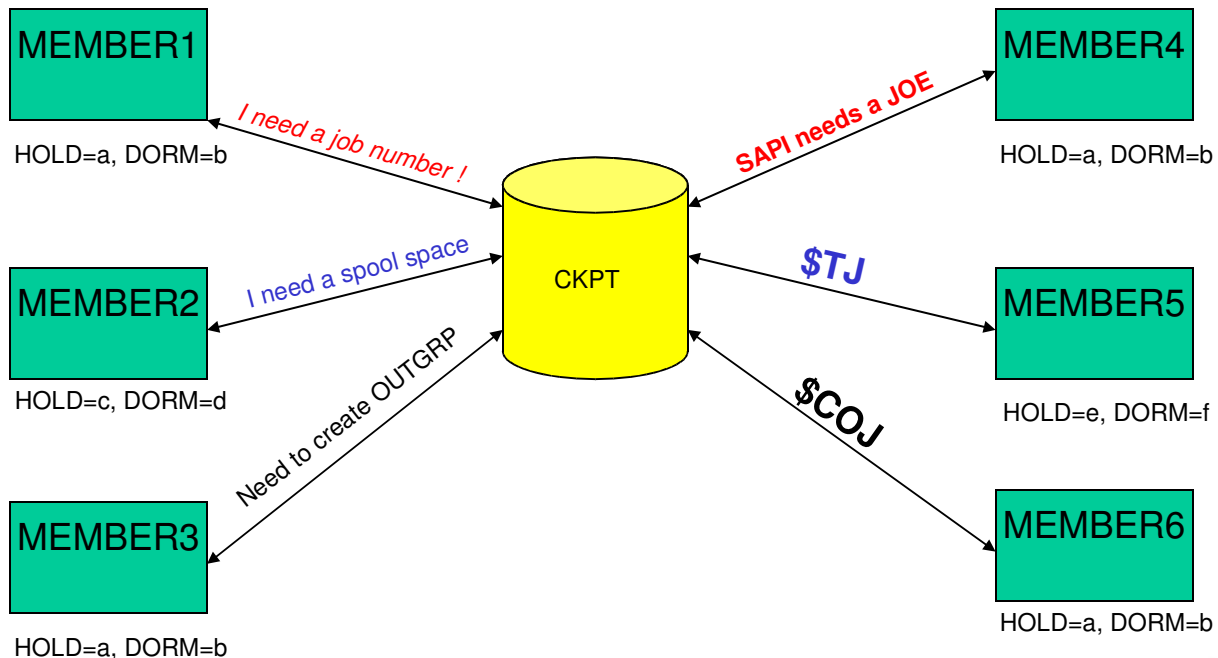
AVGWAIT = average time PCE waiting for CKPT

JQRP PCE shows the total I/O count for the interval. Since JES2 performs 2 I/Os for each job created, the IOCOUNT divided by 2 yields the count of how many jobs were created during the PERFDATA interval. In the above example $2196 / 2 = 1098$. If the interval were 6 minutes, then that would suggest approx ~3 jobs were created per second.

The AVGWAIT time gives insight in any problems surrounding CKPT access. Generally, we view this health based on order-of-magnitude where anything larger than 0.10 seconds *could* indicate contention for checkpoint.

30 Second Checkpoint Review

Checkpoint access in a MAS looks like...



Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

How equitably the checkpoint is shared amongst MAS members is controlled by the following MASDEF parameter (which have a scope of member):

HOLD= The minimum length of time a member will hold the checkpoint before it will try to release it

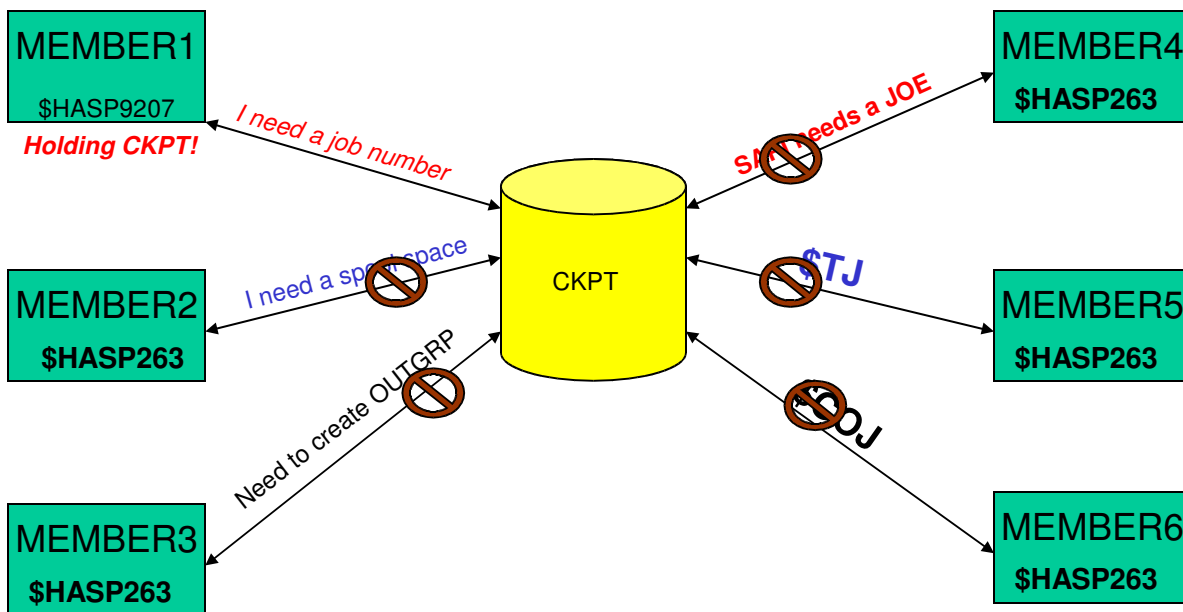
DORMANCY= The length of time a member will wait before attempting to reacquire the CKPT

Notification of a checkpoint lockout condition is based on the MASDEF parameter:

LOCKOUT=The length of time a member needing the CKPT will wait before issuing \$HASP263

Another 30 Seconds about Checkpoint...

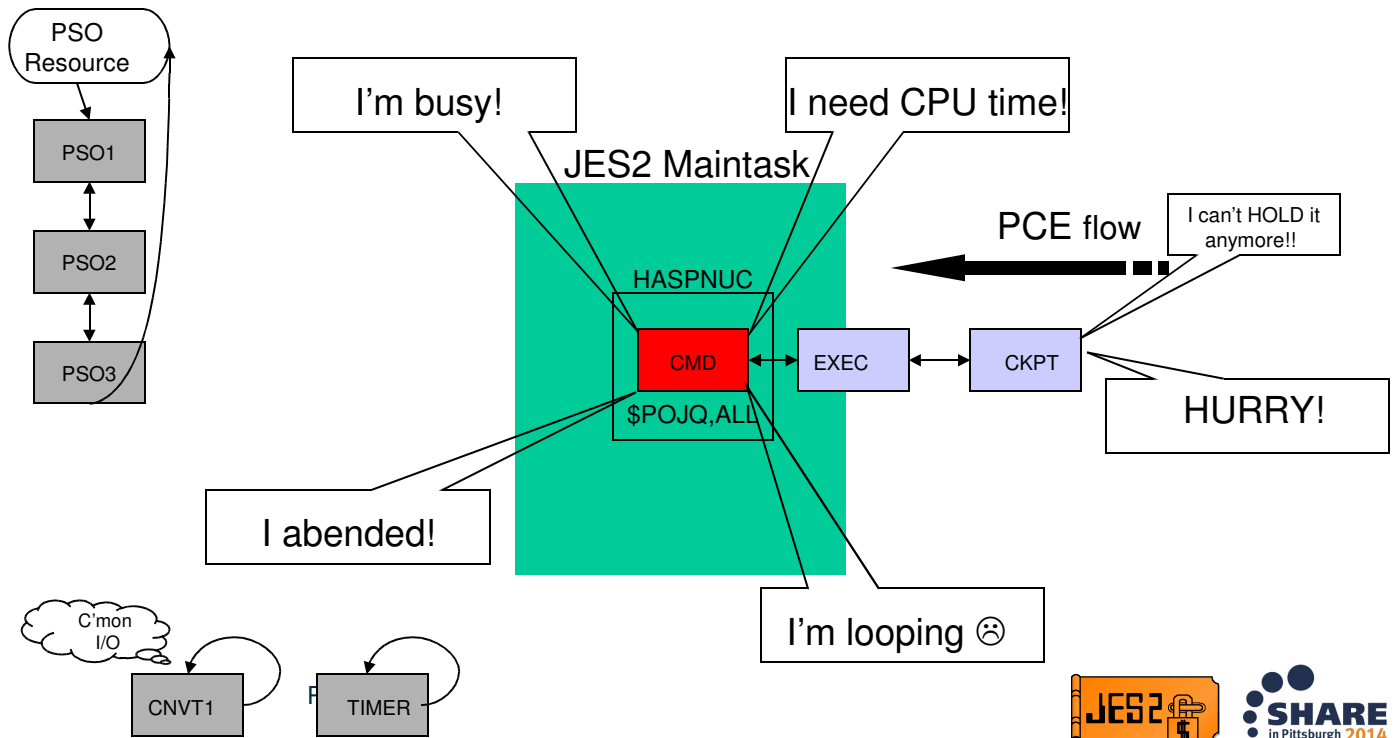
Checkpoint lockout in a MAS looks like...



Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

Just 30 More Seconds about Checkpoint (don't lockout on me...)

Checkpoint lockout on the holding member could look like...



Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

The first predicament that the CKPT PCE can find itself in, is being unable to get dispatched under the maintask. There are several reasons this could occur.

1. The PCE currently processing under maintask is busy doing valid work. It may be the nature of the work that is resulting in the excessive processing time. Any command that requires the scanning and/or filtering of a large number of jobs or outgrps can take some time to complete. The \$POJQ with a filter command is an example if it needs to process tens of thousands of jobs. It is cpu intensive and could result in \$HASP263's on other members depending on the coded LOCKOUT value.
2. JES2 is currently CPU restricted. That is the maintask TCB is not getting any or enough cycles to get through the chain of PCE's in a timely fashion to allow the CKPT PCE to run.
3. A PCE abended and has issued \$HASP098 for a termination option. If this abend occurred while the CKPT was held and the WTOR is not replied to in a timely fashion then the CKPT will not be released
4. A PCE is in a loop in which no \$WAIT is issued so it will never give up control of the maintask TCB

The first two conditions can be transient in nature the \$HASP263's will be issued but then stop as either the PCE completes its work or JES2 gets the needed CPU cycles. For second two, the \$HASP263's will be issued until the causing condition is resolved.

Diagnosing Checkpoint Lockout



- Diagnosis must occur on the system that is HOLDing the checkpoint
 - The system that is HOLDing the checkpoint...
 - Will **not** issue \$HASP263
 - Will **not** issue IOS071I 016E,**,*MASTER*, START PENDING
 - Will issue \$HASP9207 JES CHECKPOINT LOCK HELD DURATION xyz
 - Possibly other JES2 monitor \$HASP92xx messages too
 - The system that is a victim not HOLDing the checkpoint...
 - Will issue \$HASP263
 - LOCK HELD BY MEMBER abc (if CKPT on CF)

32 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Once the \$HASP263's have started the first step is determining which MAS member is holding the checkpoint.

The HASP263's and IOS071I's really indicate who is NOT holding it. Absence of these messages on a member would suggest that it is the one holding the CKPT. The \$HASP9207 message issued by the JES2 monitor identifies the system holding the CKPT.

\$HASP263 WAITING FOR ACCESS TO JES2 CHECKPOINT. LOCK HELD BY **SYSTEM** is a special case. The **SYSTEM** referred to in the message means that XES has indicated to JES2 that no member holds the lock but it is currently in the hands of XES. If this message persists and no JES2 member is showing signs of getting any access then a system with XCF/XES errors occurring is likely the problem. A Vary out of the plex of the system should force XES to release the lock.

Diagnosing Checkpoint Lockout



Persistent \$HASP263s

- Check health of overall system
 - MVS commands responding?
- Check JES2 CPU usage
- Check for outstanding JES2 WTOR's
- Check for indications of JES2 functioning
 - \$HASP250 (jobs purging)?
 - JES2 commands working?
- Diagnostics
 - Console dump JES2 on system holding the CKPT –or- slip on \$HASP9207
 - PERFDATA samples

33 Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



The first step is to determine whether or not the problem is at the system or JES2 level. If MVS commands are not responding then JES2 is likely not releasing the CKPT due to problems outside of JES2 so terminating or taking other actions on the JES2 asid are likely not to resolve the lockout condition.

If MVS appears healthy then the focus can shift to the JES2 asid: Has JES2 abended? Are there an outstanding WTOR's for JES2? Is JES2 using a lot or any CPU? Resource monitors and SDSF can assist with this or a D A, JES2 followed by another D A, JES2 will show how much CPU was used between commands. If no, CPU is used then JES2 is simply not getting a chance to run so other higher priority tasks may need to be examined. Is JES2 responding to commands? Yes, then this would indicate that JES2 PCE's are running which is predicament number two.

If JES2 is not responsive to commands and there are no other messages being issued such as \$HASP100, \$HASP250 (\$HASP395's do not count) and there is high CPU then JES2 is likely looping. Scanning the syslog looking for the last commands or messages issued by JES2 may give an indication of whether it is related to a CPU intensive command being issued.

Diagnosing Checkpoint Lockout



Transient \$HASP263s

- Messages appear on one or more systems but do not persist
- Normally caused by JES2 being temporarily busy with work/commands or short on CPU
- Diagnostics
 - Console dump JES2 on system holding the CKPT –or- slip on \$HASP9207
 - PERFDATA samples

These transient \$HASP263's come in two flavors: the first is “every once in while”. The second is of a more “roaming” nature. The first type is usually the result of a temporary condition that either resulted in JES2 being busy or not being dispatched. The second is a little more troublesome. The messages appear consistently however the system identified as holding the lock changes and may cycle through all of the members of the MAS. This is much more difficult to isolate to any specific type of problem or system and could be require a small amount of tuning of HOLD/DORMANCY

Resource Shortages



- \$HASP050 message issued
 - Not all resources are critical
- BERTs, JQEs, JOEs, JNUMs, TGs are MAS wide resources – critical!
 - BERTs - **DO NOT RUN OUT OF BERTS!!!!**
 - BERTs – use \$DCKPTSPACE,BERTUSE to identify usage
 - TGs – use \$DJOBQ,SPOOL=(% > nn) to identify usage
 - TGs - should be viewed at a job level not output level
 - unless output is SPIN
- BSCB, BUFX, CKVR, CMB, CMD, ICES, LBUF,NHB,SMFB, TTAB, VTAMB are member specific – not as critical

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



JES2 will produce \$HASP050 message indicating resource shortage – message will repeat until condition is relieved.

BERTs are one of the most critical JES2 resources. They represent non-contiguous pieces of storage on checkpoint that back/comprise other JES2 blocks such as \$CAT, \$DJB, \$JQA, etc. There are some processes in JES2 that cannot wait for BERTs to become available, thus it is imperative to avoid complete exhaustion. For the BERT resource, in addition to \$HASP050 you may also encounter:

\$HASP051 EXTREME BERT SHORTAGE detected ...

\$HASP052 JES2 BERT resource shortage is critical -- IMMEDIATE action required...

Ideally you want to have enough BERTs defined such that you would exhaust any resource it is backing (such as JOEs, JQEs, \$CATs) first rather than exhausting BERTs themselves. When in a BERT shortage condition, you want to identify and address any offending job/output –and/or- increase BERTNUM definition.

For TG shortages, you again want to identify and address any offending job/output –and-or- add additional pool space. It is important to approach TG usage at a job level because trackgroups are not restored to the TG map until the entire job is purged (exception being SPIN output, in which TGs restored when output is processed). Consider a job ADAM1 that has two pieces of non-SPIN output: JOE1 using 1 TG and the other JOE2 using 9K TGs. If you purge JOE2, it will not restore the 9K TGs because JOE1 still exists with 1 TG!

Spare spool volumes can be formatted in advance and then volume simply started \$\$ if needed –or- can be formatted dynamically on the \$\$\$PL command.

Resource Shortages



- \$JDHISTORY command will show historical usage
 - Since JES2 warmstart/IPL
 - Hourly time slices of usage (interval at the top of each hour)
 - Limit/current/low/high/average usage
 - Same resources as \$HASP050 message
- SDSF equivalent
 - RM panel
 - JH command on the MAS panel
- MSGID slip or console dump JES2 *if* needed

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



When approaching resource shortages one of the key pieces of information is whether the resource utilization is a sudden unexpected spike vs slow creep vs simply running at a fairly constant number too close to the warning limit etc. The above displays provide the answers, broken down in 1hr intervals. It also helps illustrate any relationships between resource trends - eg. Is job output growing at a 3x rate while jobnums are not growing at all? It may prompt the debugger to scrutinize a particular time interval in syslog/operlog– does it correspond with a peak or change in workload?

Questions?



Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

JES2 Service Information



<u>yy/mm/dd</u>	<u>APAR</u>	<u>COMMENTS</u>
14/02/27	OA44612	SECURITY APAR
14/05/27	OA45292	SECURITY APAR
14/02/13	OA44474	\$HASP313 & DUPLICATE ICH408I IF INCORRECT CREDENTIALS ON JOBCARD
14/02/17	OA44511	\$HASP896 RC83 TGSIZE MISMATCH FOR STUNT VOLUME DURING WARMSTART
14/02/24	OA44564	VARIOUS ERRORS RUNNING IN 64-BIT AMODE WHEN DATA LEFT HIGH HALF OF REG15
14/03/12	OA44724	AFTER OA43541, ABEND0C4 IN IFGDEBCK DURING SYSIN PUT REQUEST
14/03/27	OA44852	ABENDS0C4 ALLOCATING INTERNAL READER IN SECONDARY SUBSYSTEM
14/03/31	OA44908	INCORRECT REASON CODE RETURNED TO SAPI APPLICATION
14/04/08	OA44967	DYNAMIC PROCLIB INCORRECTLY OVERRIDING STATIC PROCLIB AT 2.1
14/04/23	OA45057	HASCSISC REASSEMBLY TO GET APAR OA44094 POST MACRO CORRECTIONS
14/04/28	OA45088	INCORRECT OR MISSING SIGNATURE RECORD IN JES2 SYMREC
14/04/29	OA45118	EXIT53 FAILS TO UPDATE EXECUTION CLASS IN \$JCT AT 2.1
14/05/05	OA45123	ABEND0C4 IN HASCPool RECOVERY
14/05/06	OA45138	ENHANCEMENTS TO JES2 EXIT32
14/05/19	OA45232	\$CBIO FAILS WITH RC=04 DURING \$VERIFY
14/05/22	OA45259	EXIT44 FAILS TO OVERRIDE JOBCLASS SCHENV
14/06/02	OA45337	JOBS AWAITING CONVERSION AFTER WLM SERVICE DEFINITION CHANGE

LEGEND:

HiPer APARs (Hi Impact, or Pervasive)

Security/Integrity

Special Attention



Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

Session 16167

OA44612 – SECURITY APAR



Rating → CVSS values:
CVSS/B=7.2
CVSS/T=6.3
CVSS/V=(AV:L/AC:L/Au:N/C:C/I:C/A:C/E:ND/RL:OF/RC:C)

Bypass → None

Fix → JES2 z/OS 1.13 and 2.1

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Refer to CVSS Scoring Site:

<http://nvd.nist.gov/cvss.cfm>

Refer to z/OS Security portal:

<http://www-03.ibm.com/systems/z/advantages/security/integrity.html>

OA45292 – SECURITY APAR



Rating → CVSS values:
CVSS/B=7.2
CVSS/T=6.3
CVSS/V=(AV:L/AC:L/Au:N/C:C/I:C/A:C/E:ND/RL:OF/RC:C)

Bypass → None

Fix → JES2 z/OS 1.12

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Session 16167

Refer to CVSS Scoring Site:

<http://nvd.nist.gov/cvss.cfm>

Refer to z/OS Security portal:

<http://www-03.ibm.com/systems/z/advantages/security/integrity.html>

OA44474 - \$HASP313 & DUPLICATE ICH408I IF INCORRECT CREDENTIALS ON JOBCARD



Problem → Incorrect credentials (eg. password) specified on jobcard

JES2 checks security environment two times resulting in ICH408I and IRR013I issued twice

Failure is counted twice increasing chances that user get revoked

Bypass → Correct the credentials to avoid recurrence of scenario and/or revocation

Fix → JES2 z/OS 2.1



Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

Session 16167

OA44511 - \$HASP896 RC83 TGSIZE MISMATCH FOR STUNTED VOLUME DURING WARMSTART



- Problem** → Volume is (or was) in stunted status
 - JES2 warmstart is performed
 - Track group size is miscalculated for the volume
 - \$HASP896 message produced & percentage spool used by jobs incorrect since total SPOOL space value incorrect
- Bypass** → Avoid putting volumes into stunted status
- Fix** → JES2 z/OS 1.12, 1.13, and 2.1

OA44564 - VARIOUS ERRORS RUNNING IN 64-BIT AMODE WHEN DATA LEFT HIGH HALF OF REG15



Problem → Processing for JES2 binary tree service -and- Extended Status and Job Modify SSIs run in 64-bit mode

Macros/routines/services that are invoked within run in 31-bit mode

JES2 improperly manages switching of amode such that extraneous data is left in part of R15

Results in Abend0C4s etc

Bypass → None


Fix → JES2 z/OS 2.1

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Session 16167

OA44724 - AFTER OA43541, ABEND0C4 IN IFGDEBCK DURING SYSIN PUT REQUEST




OA43541 installed


JES2 calls IFGDEBCK to verify DEB

DEB is on page boundary, DEBCK looks at prefix area at -4 crossing onto previous page

Abend0C4 can result



None



JES2 z/OS 1.12, 1.13, and 2.1



OA44852 - ABENDS0C4 ALLOCATING INTERNAL READER IN SECONDARY SUBSYSTEM



Problem → Job submitted to JES2 subsystem

Requests allocation of INTRDR on a different subsystem and requests list of symbols

JES2 allocation code assumes job submission occurred on this 2ndry subsystem and does not have addressability to requisite blocks

Abend0C4 can result

Bypass → Avoid submitting job attempting INTRDR allocation on a different subsystem than which it was submitted

Fix → JES2 z/OS 2.1

OA44908 - INCORRECT REASON CODE RETURNED TO SAPI APPLICATION



Problem → SAPI application does PUT/GET for SYSOUT,
collects token, and returns dataset as KEEP
Associated job is marked HOLD by operator
SAPI application attempts to delete the SYSOUT
SSS2RENM (no matching output) returned instead of
SSS2RENS (matching output not selectable)

Bypass → Avoid holding jobs for which tokens were previously
collected

Fix → JES2 z/OS 2.1

OA44967 - DYNAMIC PROCLIB INCORRECTLY OVERRIDING STATIC PROCLIB AT 2.1



Problem → PROCLIB initialization statement PROC01 with DDs on VOL=ABC. JES2 init deck contains PROC01 statement with DD on VOL=XYZ. \$DPROCLIB shows init deck PROC01 as dynamic proclib DDs on VOL=ABC.

JES2 running with static PROC00. A dynamic proc TEMP is added via \$ADDPROCLIB and then switched to PROC00 via \$TPROCLIB. \$DELPROC(PROC00) does not fall back to static instead finds no matching entries

Bypass → Jobs submission will fail for jobs accessing PROC in question
Hotstart or reallocate the dynamic proclibs explicitly specifying all appropriate keywords

Fix → JES2 z/OS 2.1

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Session 16167

OA45057 - HASCSISC REASSEMBLY TO GET APAR OA44094 POST MACRO CORRECTIONS



Problem → POST macro at z/OS 2.1 contained defects as identified by BCP APAR OA44094

HASCSISC was assembled with this macro when built for z/OS 2.1

This APAR forces reassembly of HASCSISC to pick up corrected POST macro

Bypass → Avoid reassembling any level of HASCSISC with defective POST macro

Fix → JES2 z/OS 2.1

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Session 16167

Note: If you reassembled a level of HASCSISC previous to z/OS 2.1 (HJE7790) using the defective z/OS 2.1 POST macro, then you should reassemble the module using the version of the POST macro that has been provided by APAR OA44094.

OA45088 - INCORRECT OR MISSING SIGNATURE RECORD IN JES2 SYMREC



Problem → Following a JES2 \$DISTERR, JES2 produces a symptom record in LOGREC containing diagnostic information such as MTTR/MQTR contents etc. Wrong MTTR passed to \$SIGIO routine resulting in missing or erroneous SYMREC data

Additionally, wrong MTTR/MQTR is passed to routine(s) responsible for reclaiming track groups previously marked in bad trackgroup map; thus, track cannot be reclaimed

Bypass → None

Fix → JES2 z/OS 1.12, 1.13, and 2.1

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



OA45118 - EXIT53 FAILS TO UPDATE EXECUTION CLASS IN \$JCT AT 2.1



Problem → No CLASS= coded in JCL and exit53 assigns single character jobclass via JCTCLAS. This worked at pre-2.1, but no longer honored at 2.1 due to introduction of 8 character jobclass functionality.

Also applies to exit3

Bypass → Update exit to alter both JCTJCLAS as well as JCXJCLA8 (in \$JCTX)

Fix → JES2 z/OS 2.1



Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

Session 16167

Modifying JCT field will now be honored via exit3 and exit53 respectively. However, new XPL fields (X003JCLS and X053JCLS) were also added to assist in job class modification rather than exit directly accessing JCT or JCTX.

OA45123 - ABEND0C4 IN HASCPPOOL RECOVERY



Problem → Attempt to return a cell pool fails in CPFREE routine. JES2 is in 64-bit mode and cell pool recovery fails as it expected 31-bit mode.

Alternative cell pool created and old one orphaned. Then abend0C4 or abend0C1 within job or STC making use of the specific cell.

Bypass → None

Fix → JES2 z/OS 1.12, 1.13, and 2.1

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



OA45138 - ENHANCEMENTS TO JES2 EXIT32



Problem → Exit 32 enhanced to support job eviction – specifically to allow exit to return job to execution queue and the job start from last known step (per job eviction)

Bypass → None

Fix → JES2 z/OS 1.13, and 2.1

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



OA45232 - \$CBIO FAILS WITH RC=04 DURING \$VERIFY



Problem → Installation coded own \$CBIO macro invocation with VERIFY=HDB. JES2 fails with RC=04 and message \$HASP364

Bypass → None

Fix → JES2 z/OS 1.12, 1.13, and 2.1

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Session 16167

OA45259 - EXIT44 FAILS TO OVERRIDE JOBCLASS SCHENV



Problem → Exit44 being used to override JOBCLASS SCHENV via field JQASCHE. This worked at pre-2.1, but no longer honored at 2.1 due to changes to conversion processing

Bypass → None

Fix → JES2 z/OS 2.1

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Modifying JQA field will now be honored via exit44. However, new XPL fields (X044JCLS and X044SCHE) were also added to assist in job class and scheduling environment modification respectively

OA45337 - JOBS AWAITING CONVERSION AFTER WLM SERVICE DEFINITION CHANGE



Problem → Last active MAS member terminated and then WLM Service Definition change made

Upon next hotstart JES2 will detect mismatch against that which resides on checkpoint

Jobs rejected in conversion processing

Bypass → Update WLM Service Definition, activate via V WLM,POLICY=xyz, then reset policy to desired one -or- perform JES2 all-member warmstart

Fix → JES2 z/OS 1.12, 1.13, and 2.1