



JES2 Checkpoint Reconfiguration

Permission is granted to SHARE Inc. to publish this presentation in the SHARE proceedings. IBM retains its right to distribute copies of this presentation to whomever it chooses.

SHARE, Summer 2004

Project: JES2

Session: 2662

Tom Wasik
JES2 Development
wasik@us.ibm.com

This presentation describes the JES2 checkpoint reconfiguration process and its usage. It also describes the relationship between system managed rebuild and the JES2 checkpoint reconfiguration process. A basic understanding of JES2 and the JES2 checkpoint is assumed.

This presentation was created using examples from the z/OS release 4 level of JES2. Some differences may exist in older levels of JES2.

Overview



This presentation will cover the following:

- ▶ Why should I care about CKPT reconfiguration?
- ▶ Types of checkpoint reconfiguration
- ▶ Operator initiated reconfiguration
- ▶ Error reconfiguration
- ▶ Reconfiguration at initialization
- ▶ System managed rebuild vs. reconfiguration
- ▶ Hints and tips

In this presentation I will go over the various ways a checkpoint reconfiguration can be entered. For each form of reconfiguration I will go through an examples of how to respond to the various messages to complete the task.

There will also be a discussion of System managed Rebuild and how that relates to checkpoint reconfiguration. Finally, as time permits, I will review some situations that came up over the years and how checkpoint reconfiguration was used to handle the situation.

The examples in this presentation were done at the z/OS 1.4 level of JES2. There may be some minor differences in the message text when compared to other releases.

Why do I care about CKPT?



- Ignoring it does not make it go away
- JES2 CKPT contains major JES2 data areas
 - ▶ Representations of jobs, Output, SPOOLS
- Loss of CKPT implies loss of all data on SPOOL
- Size of CKPT limits number of things JES2 supports (jobs, OUTPUT, SPOOL, etc.)
- Poor CKPT performance can impact entire SYSPLEX

Before we start looking at the reconfiguration process, let's pause to discuss why we care about the CKPT data set.

The JES2 CKPT is not one of those data sets you can just create and forget. It stores the information JES2 needs to track jobs, SYSOUT, and the SPOOL volumes themselves. Loss of the checkpoint will cause the loss of all data that JES2 is maintaining. To prevent the loss of the JES2 checkpoint, JES2 maintains 2 checkpoint data sets in parallel and has the checkpoint reconfiguration process to recover a lost data set.

However, the best code is useless if no-one knows how to use it. You do not want to learn about the checkpoint reconfiguration process when a JES2 checkpoint is lost. It is recommended that you become familiar with the process by reviewing the documentation and trying it on a test system.

The checkpoint data set size also limits the number of jobs, SYSOUT data sets, and SPOOL space a system can support. It may be necessary to use the reconfiguration process to move the checkpoint to a larger data set (perhaps on a larger volume).

One caution, checkpoint data set performance can have a significant impact on overall system and SYSPLEX performance. Careful consideration must be given to where a checkpoint data set is placed. The reconfiguration process may be needed to move a checkpoint to a better performing device.

What is checkpoint reconfiguration?



- Checkpoint reconfiguration is the only supported way to move a JES2 checkpoint data set to a new location.
- Can also suspend and resume use of a checkpoint data set.
- Ensures no loss of data if there is a checkpoint error
 - ▶ Data is written from memory thus ensuring currency
 - ▶ If you cancel out of error reconfiguration, JES2 terminates, losing some data, possibly requiring a cold start

Checkpoint reconfiguration provides a mechanism to alter the current checkpoint configuration. Use it to move a checkpoint data set to a new location, suspend use of an existing checkpoint data set, or resume using a data set that was suspended. Reconfiguration can move a checkpoint to a different device type including from DASD to coupling facility and back. Checkpoint reconfiguration gets control when there is an error on a checkpoint data set (either a physical error accessing the device or a logical error with the data being accessed). Since it uses the instorage copy of the checkpoint data to create a new checkpoint data set, it can recover from a total failure of the checkpoint data sets (as long as at least one member is still active). If instead of completing an error reconfiguration, you cancel out of it, the members that experienced the error will ABEND, causing any checkpoint data that member may have in memory to be lost. This could result in jobs failing or systems that must be IPLed. In the worst case, a cold start may be needed to recover from the error.

Forms of Checkpoint Reconfiguration



- Operator initiated
 - ▶ The result of a `$TCKPTDEF,RECONFIG=YES`
 - ▶ Used to move, suspend, or resume a CKPT
- Error reconfiguration
 - ▶ The result of
 - I/O error reading CKPT
 - Logic error reading or processing CKPT
 - ▶ Used to move or suspend CKPT with error
 - ▶ Can be performed "automatically"

An operator initiated reconfiguration is started with a `$TCKPTDEF,RECONFIG=YES` command. This command is entered on any member of the MAS and places all members into the reconfiguration. This form of the reconfiguration is used to move, suspend or resume a checkpoint data set. Either data set can be processed at this time. Canceling out of this reconfiguration will return JES2 to normal operation.

An error reconfiguration is entered when JES2 detects an error on the checkpoint data set (either with the device or with the data that was read). This reconfiguration is used to process an error with a specific data set. The data set can be moved to a new location or use of the data set can be suspended. Optionally, an installation can set up to automatically forward the error data set in the event of an error. Canceling out of an error reconfiguration will cause all members that encountered the error to be ABENDED.

Forms of Checkpoint Reconfiguration *(continued)*



- Initialization reconfiguration
 - ▶ Processing based on reason for entering
 - ▶ RECONFIG start parameter
 - Checkpoint specifications are forced to match what is specified
 - Little verification is done
 - ▶ Initialization cannot locate correct CKPT
 - Initialization does normal processing with data sets specified
 - Full verification done
 - Can switch to RECONFIG format

The final form of checkpoint reconfiguration is the initialization format. This is very different from the other reconfiguration processes. During initialization there is no instorage copy of the checkpoint available to create a new checkpoint data set. You must specify a recently used checkpoint to read the data from. Initialization reconfiguration is also the only form that is a single member processes. In general, it can only be used on an all-member warm start or if there are no JES2 address spaces active in a MAS. There are 2 forms of the initialization reconfiguration. The first is the result of a RECONFIG being specified on the JES2 start PARM=. This form allows installations to change the checkpoint specifications JES2 uses. When you complete this dialog, JES2 will use the current checkpoint specifications (the ones displayed in the messages) in preference to what is currently defined in the checkpoint control record. Getting this wrong has resulted in cold starts.

The other form is entered as the result of JES2 initialization being unable to locate the correct data sets to use. In this situation, JES2 will not start unless it can allocate (and read) the same checkpoint data sets it was using when JES2 was last active. You can convert from this format to a RECONFIG format if there is a problem accessing the data sets that were previously used (and this is an all member warm start). When you complete this dialog, JES2 will use the current checkpoint specifications just as though you specified them in the initialization deck. Normal processing and verification is done (including any forwarding processing).

Operator Initiated Reconfiguration



`$TCKPTDEF, RECONFIG=YES`

*\$HASP233 REASON FOR JES2 CHECKPOINT RECONFIGURATION IS OPERATOR REQUEST

*\$HASP285 JES2 CHECKPOINT RECONFIGURATION STARTED - DRIVEN BY MEMBER IBM1

*\$HASP271 CHECKPOINT RECONFIGURATION OPTIONS

VALID RESPONSES ARE:

'1' - FORWARD CKPT1 TO NEWCKPT1

'2' - FORWARD CKPT2 TO NEWCKPT2

'5' - SUSPEND THE USE OF CKPT1

'6' - SUSPEND THE USE OF CKPT2

'CANCEL' - EXIT FROM RECONFIGURATION

CKPTDEF (NO OPERANDS) - DISPLAY MODIFIABLE SPECIFICATIONS

CKPTDEF (WITH OPERANDS) - ALTER MODIFIABLE SPECIFICATIONS

*05 \$HASP272 ENTER RESPONSE (ISSUE D R,MSG=\$HASP271 FOR RELATED MSG)

The `$TCKPTDEF RECONFIG=YES` option is issued on one member of the MAS. All messages are directed to the console that issued the command. As a result, the messages may be issued by a JES2 MAS member that is not normally displayed on the console. Since not all messages are highlighted (and highlighted messages scroll in many installations), you should choose your console carefully. A console that does not get much traffic or an extended console that is used only for commands would be ideal. There are a total of 6 options that can appear on the HASP271 message (sometimes referred to as the *gateway to the west* message). These options are:

'1' - FORWARD CKPT1 TO NEWCKPT1

'2' - FORWARD CKPT2 TO NEWCKPT2

'5' - SUSPEND THE USE OF CKPT1

'6' - SUSPEND THE USE OF CKPT2

'7' - UPDATE AND START USING CKPT1

'8' - UPDATE AND START USING CKPT2

Only those options that are currently valid (based on the environment) are displayed. Options 3 and 4 are not used (retired in SP 5.1.0).

A common error at this point is looking for an option that is not there. If it is missing, that is because it is not valid at this time. For example, you cannot suspend CKPT1 if CKPT2 has been suspended. That would leave you with no checkpoints.

Select forward option



```

r 5,1
IEE600I REPLY TO 05 IS;1
*$HASP273 JES2 CKPT1 DATA SET WILL BE ASSIGNED TO NEWCKPT1
SYS2.JESCKPT1 ON CKPTPK

VALID RESPONSES ARE:

'CONT' - PROCEED WITH ASSIGNMENT
'CANCEL' - EXIT FROM RECONFIGURATION
CKPTDEF (NO OPERANDS) - DISPLAY MODIFIABLE
                        SPECIFICATIONS
CKPTDEF (WITH OPERANDS) - ALTER MODIFIABLE
                        SPECIFICATIONS

*06 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP273 FOR RELATED MSG)

```

In this example, I decided that CKPT1 was to be forwarded. Forwarding is the process of creating a new copy of the checkpoint, writing all the current checkpoint data to the new checkpoint, updating all internal pointers to the new checkpoint, and suspending the use of the old data set. This process is coordinated with all members of the MAS. JES2 will ensure that all MAS members can access the new checkpoint data set before completing the process. Any problem will result in an error message and this message being redisplayed.

In this case, there was a NEWCKPT1 data set specified. Before JES2 allocates the new data set, it will confirm that this is the data set you want to forward to. At this point you can change the target for the forwarding. A CKPTDEF command with NEWCKPT1 (the file name from the HASP273 message) can be used as a response to change the target of the forwarding. Some common errors at this point are:

- Trying to specify CKPTDEF CKPT1=.... instead of NEWCKPT1
- Using a \$TCKPTDEF command to try to alter NEWCKPT1
- Simply forgetting the CKPTDEF and specifying NEWCKPT1 or CKPT1

Continue with Forwarding



r 6,cont

```
IEE600I REPLY TO 06 IS;CONT
$HASP414 MEMBER IBM1 -- OBTAIN FAILED FOR NEWCKPT1 SYS2.JESCKPT1 ON
CKPTPK WITH CC 8
*$HASP278 UNABLE TO LOCATE OR UNABLE TO USE NEWCKPT1
SYS2.JESCKPT1 ON CKPTPK
DOES NOT EXIST OR IS NOT USABLE - REFER TO PREVIOUS
MESSAGE(S)
```

VALID RESPONSES ARE:

```
'CANCEL' - EXIT FROM RECONFIGURATION
'CREATE' - CREATE DATA SET
CKPTDEF (NO OPERANDS) - DISPLAY MODIFIABLE
                        SPECIFICATIONS
CKPTDEF (WITH OPERANDS) - ALTER MODIFIABLE
                        SPECIFICATIONS
*07 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP278 FOR RELATED MSG)
```

Even though the data set was specified as NEWCKPT1, the data set does not exist on the volume. Before JES2 creates the data set, it will again confirm that it is using the correct data set specification. If a response of CREATE is given JES2 will allocate a new data set with the correct amount of space for the checkpoint (note: this works for DASD specifications only). This is useful if you need to place a checkpoint in an unexpected location. However, it is recommended that sufficiently large NEWCKPT1 and NEWCKPT2 data sets exist so that the space on the volume is not accidentally used by some other data set.

At this point, you can also change the target data set specification. If you do, JES2 will confirm the new specifications before allocating a new data set. In general, if you respond to any reconfiguration message with a CKPTDEF operand, you will get another WTO/WTOR to confirm the new values you have specified.

One thing to note, the \$HASP414 message is not a highlighted message and will scroll off the screen if there is a large amount of message traffic on this console. This should be considered before issuing the \$TCKPTDEF command.

We're done!



r 07,create

```
IEE600I REPLY TO 07 IS;CREATE
$HASP280 JES2 CKPT1 DATA SET (SYS2.JESCKPT1 ON CKPTPK) IS NOW IN USE
IEE400I THESE MESSAGES CANCELLED - 08.
*$HASP256 FUTURE AUTOMATIC FORWARDING OF CKPT1 IS SUSPENDED UNTIL
NEWCKPT1 IS RESPECIFIED.
ISSUE $T CKPTDEF,NEWCKPT1=(...) TO RESPECIFY
$HASP255 JES2 CHECKPOINT RECONFIGURATION COMPLETE
```

\$ckptdef

```
$HASP829 CKPTDEF
$HASP829 CKPTDEF CKPT1=(DSNAME=SYS2.JESCKPT1,VOLSER=CKPTPK,
$HASP829 INUSE=YES,VOLATILE=NO),
$HASP829 CKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=CKPTPK,
$HASP829 INUSE=YES,VOLATILE=NO),NEWCKPT1=(DSNAME=,
$HASP829 VOLSER=),NEWCKPT2=(DSNAME=SYS2.JESCKPT2,
$HASP829 VOLSER=CKPTPK),MODE=DUAL,DUPLEX=ON,LOGSIZE=4,
```

That's it. The response of 'CREATE' caused a new data set to be allocated and assigned as the CKPT1 data set. The NEWCKPT1 specification has been cleared. As stated in the HASP256 message, automatic forwarding in the event of an I/O error has been suspended because there is no NEWCKPT1 specification. At this point, if the movement of the checkpoint is to be permanent, the JES2 initialization decks should be updated and a new NEWCKPT1 established.

If you forget to update the initialization decks, then there are 2 ways JES2 will attempt to find the new checkpoints. During warm start, if there are existing JES2 MAS members active, an XCF message is sent to them asking where the checkpoint currently resides. The values from the initialization deck are updated with the response from the other members and warm start continues. If there are no other JES2 members active, then warm start will attempt to read the forwarding information in the header of the checkpoints specified in the init deck. Since each checkpoint has the names of both checkpoint data sets, JES2 uses this information to verify it has a consistent checkpoint specification. It is when this process fails that JES2 enters the initialization reconfiguration and asks the operator where the checkpoint is.



Suspending CKPT1

\$TCKPTDEF, RECONFIG=YES

```
*$HASP285 JES2 CHECKPOINT RECONFIGURATION STARTING
*$HASP233 REASON FOR JES2 CHECKPOINT RECONFIGURATION IS OPERATOR
REQUEST
*$HASP285 JES2 CHECKPOINT RECONFIGURATION STARTED - DRIVEN BY
MEMBER IBM1
*$HASP271 CHECKPOINT RECONFIGURATION OPTIONS

VALID RESPONSES ARE:

'1' - FORWARD CKPT1 TO NEWCKPT1
'2' - FORWARD CKPT2 TO NEWCKPT2
'5' - SUSPEND THE USE OF CKPT1
'6' - SUSPEND THE USE OF CKPT2
'CANCEL' - EXIT FROM RECONFIGURATION
CKPTDEF (NO OPERANDS) - DISPLAY MODIFIABLE
                        SPECIFICATIONS
CKPTDEF (WITH OPERANDS) - ALTER MODIFIABLE
                        SPECIFICATIONS
*14 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP271 FOR RELATED MSG)
```

Let's suspend the use of CKPT1. This would not normally be done using an operator reconfiguration. Normally this happens when an error is encountered and the best action is to suspend the use of the checkpoint with the problem. It is highly recommended that you always run with 2 checkpoint data sets.

This process starts the same as the forwarding case. We get the gateway message (HASP271) and reply option 5. This requests that CKPT1 be suspended.

CKPT1 Suspended



```
r 14,5
IEE600I REPLY TO 14 IS;5
$HASP280 JES2 CKPT1 DATA SET (SYS2.JESCKPT1 ON CKPTPK) IS NO LONGER
IN USE
$HASP255 JES2 CHECKPOINT RECONFIGURATION COMPLETE
```

\$dckptdef

```
$HASP829 CKPTDEF
$HASP829 CKPTDEF CKPT1=(DSNAME=SYS2.JESCKPT1,VOLSER=CKPTPK,
$HASP829 INUSE=NO),CKPT2=(DSNAME=SYS1.JESCKPT2,
$HASP829 VOLSER=CKPTPK,INUSE=YES,VOLATILE=NO),
$HASP829 NEWCKPT1=(DSNAME=,VOLSER=),
$HASP829 NEWCKPT2=(DSNAME=SYS2.JESCKPT2,VOLSER=CKPTPK),
$HASP829 MODE=DUAL,DUPLEX=ON,LOGSIZE=4,
```

And we are done! CKPT1 is no longer in use. As you can see this is one of the easiest things to do using CKPT reconfiguration. A \$D CKPTDEF shows that CKPT1 is no longer in use. (Note the INUSE=NO). The actual specification for CKPT1 is still there to simplify resuming the use of the data set.

Resuming CKPT1



```
$tckptdef,reconfig=yes
```

```
*$HASP285 JES2 CHECKPOINT RECONFIGURATION STARTING
*$HASP233 REASON FOR JES2 CHECKPOINT RECONFIGURATION IS OPERATOR
REQUEST
*$HASP285 JES2 CHECKPOINT RECONFIGURATION STARTED - DRIVEN BY
MEMBER IBM1
*$HASP271 CHECKPOINT RECONFIGURATION OPTIONS

VALID RESPONSES ARE:

'2' - FORWARD CKPT2 TO NEWCKPT2
'7' - UPDATE AND START USING CKPT1
'CANCEL' - EXIT FROM RECONFIGURATION
CKPTDEF (NO OPERANDS) - DISPLAY MODIFIABLE
                        SPECIFICATIONS
CKPTDEF (WITH OPERANDS) - ALTER MODIFIABLE
                        SPECIFICATIONS
*15 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP271 FOR RELATED MSG)
```

OK, now let's resume using CKPT1. Again, we start with the gateway message. Notice that there are less options on this message than in previous examples. Also, option '7' has appeared for the first time. This is the option we will use this time.

Confirm CKPT1 data set



r 15,7

IEE600I REPLY TO 15 IS:7

*\$HASP273 JES2 CKPT1 DATA SET WILL BE ASSIGNED TO
SYS2.JESCKPT1 ON CKPTPK

VALID RESPONSES ARE:

'CONT' - PROCEED WITH ASSIGNMENT

'CANCEL' - EXIT FROM RECONFIGURATION

CKPTDEF (NO OPERANDS) - DISPLAY MODIFIABLE
SPECIFICATIONS

CKPTDEF (WITH OPERANDS) - ALTER MODIFIABLE
SPECIFICATIONS

*16 \$HASP272 ENTER RESPONSE (ISSUE D R,MSG=\$HASP273 FOR RELATED MSG)

Again, before JES2 allocates a data set, it asks you to confirm that the data set is correct. In this case, let's say it is not correct. We want to resume CKPT1 to where it was when we started playing with reconfiguration. In our case, CKPT1=(DSN=SYS1.JESCKPT1, VOL=PRCKPT).

Specifying a new CKPT1



```

16,ckptdef ckpt1=(dsn=sys1.jesckpt1,vol=prckpt)
IEE600I REPLY TO 16 IS;CKPTDEF CKPT1=(DSN=SYS1.JESCKPT1,VOL=PRCKPT
*$HASP829 CKPTDEF CKPT1=(DSNAME=SYS1.JESCKPT1,VOLSER=PRCKPT),
      NEWCKPT1=(DSNAME=,VOLSER=),
      NEWCKPT2=(DSNAME=SYS2.JESCKPT2,VOLSER=CKPTPK),
      DUPLEX=ON
*$HASP273 JES2 CKPT1 DATA SET WILL BE ASSIGNED TO
      SYS1.JESCKPT1 ON PRCKPT

VALID RESPONSES ARE:

'CONT' - PROCEED WITH ASSIGNMENT
'CANCEL' - EXIT FROM RECONFIGURATION
CKPTDEF (NO OPERANDS) - DISPLAY MODIFIABLE
                        SPECIFICATIONS
CKPTDEF (WITH OPERANDS) - ALTER MODIFIABLE
                        SPECIFICATIONS
*17 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP273 FOR RELATED MSG)
  
```

Whenever the HASP273 message is displayed, JES2 is asking to confirm the data set to use for the checkpoint. At this point, you can use the file name specified (CKPT1, CKPT2, NEWCKPT1 or NEWCKPT2) to update the target data set specification. In this case we respecified our normal CKPT1 specifications.

Completed resuming CKPT1



17,cont

```
IEE600I REPLY TO 17 IS;CONT
$HASP280 JES2 CKPT1 DATA SET (SYS1.JESCKPT1 ON PRCKPT) IS NOW IN USE
$HASP255 JES2 CHECKPOINT RECONFIGURATION COMPLETE
```

\$dckptdef

```
$HASP829 CKPTDEF
$HASP829 CKPT1=(DSNAME=SYS1.JESCKPT1,VOLSER=PRCKPT,
$HASP829 INUSE=YES,VOLATILE=NO),
$HASP829 CKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=CKPTPK,
$HASP829 INUSE=YES,VOLATILE=NO),NEWCKPT1=(DSNAME=,
$HASP829 VOLSER=),NEWCKPT2=(DSNAME=SYS2.JESCKPT2,
$HASP829 VOLSER=CKTPK),MODE=DUAL,DUPLEX=ON,LOGSIZE=4,
```

This completes the resuming of CKPT1.

The resuming of CKPT1 to a new location has confused many customers over the years. Most never get past the gateway message, thinking that there is an option missing (resume to a new place). Others get stuck on the confirmation message, never thinking of specifying CKPT1 instead of NEWCKPT1. Seeing it on these charts makes it look easy. But the only way it will be easy for you, is if you get familiar with the process on a test JES2 system. When the system is down and there are managers staring over your shoulder watching you type, the VP is on the phone, this is not the time to become familiar with how reconfiguration works.

Error reconfiguration



```

$HASP291 MEMBER IBM1 -- JES2 CKPT2 DATA SET I/O ERROR
VOL=CKPTPK DSN=SYS2.JESCKPT2
*** CHECKPOINT DATA SET NOT DAMAGED BY THIS MEMBER ***
EXCP COMP CODE = 41   IOS COMP CODE   = 51
CSW UNIT STATUS = 00   CSW CHNL STATUS = 06
CSW RESIDUAL CT = 0000 IOB SENSE DATA = 0000
NO CURRENT CCW
*$HASP275 MEMBER IBM1 -- JES2 CKPT2 DATA SET - I/O ERROR - REASON CODE
K15
*$HASP233 REASON FOR JES2 CHECKPOINT RECONFIGURATION IS CKPT2 I/O
ERROR(S) ON 1 MEMBER(S)
*$HASP285 JES2 CHECKPOINT RECONFIGURATION STARTED - DRIVEN BY
MEMBER IBM1
*$HASP282 NEWCKPT2 DSNAME, VOLUME AND STRNAME ARE NULL

VALID RESPONSES ARE:

'TERM' - TERMINATE MEMBERS WITH I/O ERROR ON CKPT2
'DELETE' - DISCONTINUE USING CKPT2
CKPTDEF (NO OPERANDS) - DISPLAY MODIFIABLE
                        SPECIFICATIONS
CKPTDEF (WITH OPERANDS) - ALTER MODIFIABLE
                        SPECIFICATIONS
*24 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP282 FOR RELATED MSG)

```

Lets go through a simple device error on a checkpoint data set with no new checkpoint data set set up (bonus points to the person that identifies how I got an I/O error to test this form of reconfiguration). JES2 bypasses the gateway message (HASP271) and proceeds as if the checkpoint is being forwarded. Notice that "DELETE" is now a valid response. This will suspend the error checkpoint and is a fast way to continue normal operations in the event of a error. "DELETE" in only an option if you have 2 checkpoint data sets in use. Though "DELETE" is a simple way out of a reconfiguration, automatic forwarding is the best option.

Specifying a new CKPT2 data set



```
24ckptdef newckpt2=(dsn=sys1.jesckpt2,vol=spool1)
```

```
IEE600I REPLY TO 24 IS;CKPTDEF NEWCKPT2=(DSN=SYS1.JESCKPT2,VOL=SPO
*$HASP829 CKPTDEF NEWCKPT1=(DSNAME=,VOLSER=),
          NEWCKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=SPOOL1),
          DUPLEX=ON
*$HASP273 JES2 CKPT2 DATA SET WILL BE ASSIGNED TO NEWCKPT2
SYS1.JESCKPT2 ON SPOOL1
```

VALID RESPONSES ARE:

```
'CONT' - PROCEED WITH ASSIGNMENT
'TERM' - TERMINATE MEMBERS WITH I/O ERROR ON CKPT2
'DELETE' - DISCONTINUE USING CKPT2
CKPTDEF (NO OPERANDS) - DISPLAY MODIFIABLE
                        SPECIFICATIONS
CKPTDEF (WITH OPERANDS) - ALTER MODIFIABLE
                        SPECIFICATIONS
*25 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP273 FOR RELATED MSG)
```

Using CKPTDEF, we specify where we want CKPT2 to be placed (by specifying a value for NEWCKPT2). As with the operator reconfiguration, JES2 confirms the data set that will be used for the checkpoint.

Error Reconfig complete



25,cont

```
IEE600I REPLY TO 25 IS;CONT
IEF524I 1004, VOLUME CKPTPK PENDING OFFLINE
IOS000I 1004,**,SIM,**,**06,,CKPTPK,JES2
IOS000I 1004,**,SIM,**,**06,,CKPTPK,JES2
IOS000I 1004,**,SIM,**,**06,,CKPTPK,JES2
$HASP280 JES2 CKPT2 DATA SET (SYS1.JESCKPT2 ON SPOOL1) IS NOW IN USE
*$HASP256 FUTURE AUTOMATIC FORWARDING OF CKPT2 IS SUSPENDED UNTIL
NEWCKPT2 IS RESPECIFIED.
ISSUE $T CKPTDEF,NEWCKPT2=(...) TO RESPECIFY
$HASP255 JES2 CHECKPOINT RECONFIGURATION COMPLETE
```

Once you confirm the new data set specifications, JES2 completes the forwarding of the checkpoint data set. The I/O errors are a result of JES2 trying to write forwarding information in the old checkpoint data set. These can be ignored.

In case you haven't figured out how the error was introduced, the SIM in the IOS message should be a give away. What I did was place the checkpoint on a volume by itself and then did a V xxxx,OFFLINE,FORCE. IOS will BOX the device, failing all I/O to it. This way I can test I/O error processing and get my DASD back when I need it (using V xxxx,ONLINE,UNCOND).

Error RECONFIG with no operator intervention



```

$HASP291 MEMBER IBM1 -- JES2 CKPT2 DATA SET I/O ERROR
VOL=CKPTPK DSN=SYS1.JESCKPT2
*** CHECKPOINT DATA SET NOT DAMAGED BY THIS MEMBER ***
EXCP COMP CODE = 41   IOS COMP CODE = 51
CSW UNIT STATUS = 00   CSW CHNL STATUS = 06
CSW RESIDUAL CT = 0000 IOB SENSE DATA = 0000
NO CURRENT CCW
$HASP285 JES2 CHECKPOINT RECONFIGURATION STARTING
*$HASP275 MEMBER IBM1 -- JES2 CKPT2 DATA SET - I/O ERROR - REASON CODE
K10
*$HASP233 REASON FOR JES2 CHECKPOINT RECONFIGURATION IS CKPT2 I/O
ERROR(S) ON 1 MEMBER(S)
*$HASP285 JES2 CHECKPOINT RECONFIGURATION STARTED - DRIVEN BY
MEMBER IBM1
IOS000I 1004,**,SIM,**,**06,,,CKPTPK,JES2
IOS000I 1004,**,SIM,**,**06,,,CKPTPK,JES2
IOS000I 1004,**,SIM,**,**06,,,CKPTPK,JES2
$HASP280 JES2 CKPT2 DATA SET (SYS1.JESCKPT2 ON SPOOL1) IS NOW IN USE
*$HASP256 FUTURE AUTOMATIC FORWARDING OF CKPT2 IS SUSPENDED UNTIL
NEWCKPT2 IS RESPECIFIED.
ISSUE $T CKPTDEF,NEWCKPT2=(...) TO RESPECIFY
$HASP255 JES2 CHECKPOINT RECONFIGURATION COMPLETE
  
```

This is the same error as the previous case. However, this time NEWCKPT2 was specified along with OPVERIFY=NO. This time, there are no WTORs, everything is processed with no operator intervention. Some of these messages had to be retrieved from SYSLOG since they never made it to my screen. The highlighted HASP256 message is the best way to know that something happened. It is suggestion to have automation send notification to the system programmer (page, e-mail) that the HASP256 was issued. Also, have automation log the checkpoint specifications after the HASP255 message is issued to a location that can be retrieved when needed (even when JES2 is down).

What happens if you say TERM?



```
*$HASP275 MEMBER IBM1 -- JES2 CKPT2 DATA SET - I/O ERROR - REASON CODE
K15
*$HASP233 REASON FOR JES2 CHECKPOINT RECONFIGURATION IS CKPT2 I/O
ERROR(S) ON 1 MEMBER(S)
*$HASP285 JES2 CHECKPOINT RECONFIGURATION STARTED - DRIVEN BY
MEMBER IBM1
*$HASP282 NEWCKPT2 DSNAME, VOLUME AND STRNAME ARE NULL

VALID RESPONSES ARE:

'TERM' - TERMINATE MEMBERS WITH I/O ERROR ON CKPT2
'DELETE' - DISCONTINUE USING CKPT2
CKPTDEF (NO OPERANDS) - DISPLAY MODIFIABLE
                        SPECIFICATIONS
CKPTDEF (WITH OPERANDS) - ALTER MODIFIABLE
                        SPECIFICATIONS
*15 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP282 FOR RELATED MSG)
```

Lets go through the same error and this time, lets just give up and say 'TERM' to exit the reconfiguration. Everything starts the same way.

What happens if you say TERM?



15,term

```
IEE600I REPLY TO 15 IS;TERM
```

```
*$HASP281 ARE YOU SURE THAT YOU WANT TO END THE CHECKPOINT
RECONFIGURATION BY CAUSING THE 1 MEMBER(S) WITH I/O ERRORS
TO ABEND? THIS COULD RESULT IN A LOSS OF THE JOB QUEUE AND
THE OUTPUT QUEUE REQUIRING A COLD START.
```

VALID RESPONSES ARE:

```
'TERM' - CONFIRM THAT 1 MEMBERS SHOULD BE TERMINATED
OTHER - CONTINUE RECONFIGURATION (RE-ISSUE LAST WTOR)
*16 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP281 FOR RELATED MSG)
```

16,term

```
IEE600I REPLY TO 16 IS;TERM
```

```
*17 $HASP279 MEMBER IBM1 -- SPECIFY EMERGENCY CKPT VOLUME SERIAL OR
'*NONE*'
```

17*none*

```
IEE600I REPLY TO 17 IS;*NONE*
```

```
*$HASP095 JES2 CATASTROPHIC ERROR. CODE = $K15
```

s jes2,parm=(warm,noreq)

```
IEF403I JES2 - STARTED - TIME=17.50.07
$HASP490 HOT START DENIED -- RE-IPL REQUIRED
$HASP085 JES2 TERMINATION COMPLETE
```

When you reply 'TERM' JES2 tries to let you know that this may not be a smart thing to do. This is to help those who did not attend this presentation know that they are throwing out data and that they may have to cold start. Any reply other than 'TERM' will put them back where they were when they said 'TERM'.

In this case, I was stubborn, and just replied 'TERM'

This gets us the EMERGENCY CKPT message. This message comes out in 2 cases, when you TERM out of an error reconfiguration or when there is an error in the reconfiguration code itself. An emergency ckpt is a copy of the current checkpoint written out (to DASD) that JES2 can use to restart. The data set does not have to exist. The name is hard coded as SYS1.EMERGENCY.HASPCKPT. You reply a volume with space and we create the data set on the volume. Over the years, this has saved a number of installation that terminated out of an error reconfiguration, created one of these data sets (just in case) and then found their CKPT1 and CKPT2 data sets unusable (often there was no CKPT2). If you ever get this message, it is worth the time to create an emergency checkpoint, just in case.

The emergency checkpoint can be pre-allocated however, if it is too small, then JES2 will be unable to use it for the checkpoint. Because of this, it is recommended that you allow JES2 to create the data set. Once the emergency data set is no longer needed, it should be deleted.

I chose not to and JES2 came down with a \$K15. At this point, hot starts are denied (since some checkpoint data has been lost).

My CKPT is too Small, What Now?



- Try to increase number of JOBS and get:

```
$HASP003 RC=(74),OUTDEF - CURRENT CHECKPOINT DATA SETS ARE TOO  
$HASP003 SMALL
```

- Multiple ways to address this
 - ▶ RECONFIG to move to larger data sets
 - ▶ Resize CKPTs while JES2 is not looking
- How to do it depends on where CKPT is
- Only use these when the system is quiet

As time passes, what used to seem big suddenly seems small (Who needs more than 16 meg of storage, or 2 gig of storage). Same is true of the checkpoint data sets. It may be a \$ACTIVATE, or an increase in the number of SPOOL volumes in the system, or just installing a new release. Someday, you will discover your checkpoints are too small. The traditional way to increase the checkpoint size is to use the reconfiguration process to move the checkpoint to a new place (data set or CF) that is larger than the current specification. This works, but what if you don't want to move the checkpoint, just make it bigger without moving it. There are some creative ways that can be used that do not require the checkpoint to be moved and involve little to no use of the reconfiguration process.

The method you use will depend on your availability requirements and where you have placed your checkpoint.

These procedures involve running with one checkpoint for some period of time. This exposes you to a loss of all CKPT data if there is a major crash at this time. Consider doing this at a time when the system is fairly quiet.

Increasing a CF checkpoint



- CF is the easiest and least error prone to expand
 - ▶ Create a policy with a larger size for your CF
 - ▶ Activate that policy
 - Get "1 POLICY CHANGE(S) PENDING."
 - ▶ Now need to disconnect JES2 from structure
 - \$PJES2,ABEND - or -
 - RECONFIG to suspend using the checkpoint
 - ▶ SETXCF START,REBUILD,STRNAME=xxx
 - ▶ Reconnect to structure (hot start or reconfig)

Increasing the size of a checkpoint on a CF is the easiest and least error prone to do. Most of the work is done by MVS. The trick is to do this when JES2 is not looking (not connected to the structure).

- ▶ First, update the CFRM policy to increase the size of the checkpoint structure. Do not forget to activate the new policy (common error). When the policy is activated, you should see at least one policy change is pending (if you do a DISPLAY XCF,STRUCTURE,STRNAME=ckpt_structure it should indicate a "POLICY CHANGE PENDING - CHANGE" in the output).
- ▶ Next you have to get JES2 to disconnect from the structure. This can be done by either ABENDING JES2 (\$PJES2,ABEND) or using the reconfiguration dialog to suspend the use of the checkpoint structure. Either option could impact the performance of your system. Ensure you pick a time when this impact is not going to be a problem
- ▶ Rebuild the structure using the SETXCF START,REBUILD,STRNAME=ckpt_structure command. This will allow MVS to update the size of the structure. At this point, a DISPLAY XCF of the structure should not show any pending changes.
- ▶ Reconnect to the structure by either hot starting JES2 or using the reconfiguration dialog to resume using the checkpoint data set. At this point, a \$D CKPTSPACE should reflect the larger checkpoint.

Increasing a DASD checkpoint



- DASD CKPTs are a bit more of a challenge
- Idea same as CF checkpoints
 - ▶ Stop using the data set
 - ▶ Make data set bigger (delete and reallocate)
 - ▶ Resume using the data set.
- Make sure when you are manipulating the data set, it is the correct data set

Increasing the size of a DASD checkpoint is not that much harder than a CKPT on CF but it is more error prone. The process is similar to a CF checkpoint.

- ▶ Stop JES2 from using the data set
- ▶ Make the data set larger
- ▶ Resume using the data set

The most important aspect is to make sure you are working with the correct data set name and volume. The biggest mistake people make is manipulating the wrong checkpoint.

Increasing a DASD checkpoint (cont...)



- Case 1, CKPT1 on CF, CKPT2 on DASD
- Want to increase size of CKPT2 data set
 - ▶ Since we are MODE=DUPLEX we can just stop duplexing (ensure command completes)
 - *ROUTE ALL \$T CKPTDEF,DUPLEX=OFF
 - ▶ Delete current CKPT2 data set and reallocate
 - Verify you are deleting correct data set
 - Ask coworker to verify DSN and VOLSER
 - When you are sure, check again
 - ▶ Resume duplexing
 - *ROUTE ALL \$T CKPTDEF,DUPLEX=ON

The first case is a CKPT1 on CF and CKPT2 on DASD. Actually, this works any time you are increasing CKPT2 and

- ▶ there is a CKPT1 data set and it is INUSE=YES
- ▶ MODE=DUPLEX is set

First stop using CKPT2 by setting DUPLEX=OFF on all members. This unallocates CKPT2 on all members. Make sure you wait for the command to complete on all members! Are you the kind of person that wants to verify this? Well you can set DUPLEX=ON on just one member. You should see the following messages:

```
$HASP266 JES2 CKPT2 DATA SET IS BEING FORMATTED
```

```
$HASP267 JES2 CKPT2 DATA SET HAS BEEN SUCCESSFULLY FORMATTED
```

These messages are issued whenever all members are not duplexing and then one starts. If you set DUPLEX=ON on a member, reissue the DUPLEX=OFF command.

Now that no-one is using CKPT2, it can be safely deleted and reallocated with a larger size. However before you delete anything, verify that you are deleting the correct data set. Issue a \$DCKPTDEF,CKPT2 to display the current CKPT2 data set name AND volser. Double check it. Ask a coworker to double check it (side note, if it is wrong, blame him, he checked it). Once you are sure it is right, check it one more time.

This may seem a bit excessive and humorous, but I there have been many stories where experienced system programmers have gotten it wrong and made a real mess. Take the time to get it right!

Once a new data set has been created, DO NOT FORGET to set DUPLEX=ON on all members. Again, another easy thing to forget.

Increasing a DASD checkpoint (cont...)



- Case 2, CKPT1 and CKPT2 on DASD
- Want to increase size of CKPT1 data set
 - ▶ Use RECONFIG to suspend CKPT1
 - ▶ Delete current CKPT1 data set and reallocate
 - Verify you are deleting correct data set
 - Ask coworker to verify DSN and VOLSER
 - When you are sure, check again
 - ▶ Use RECONFIG to resume using new CKPT1

This is the general case of resizing a DASD CKPT no matter where it is or the mode you are in. It requires using the form of the reconfiguration that suspends the use of a CKPT and then the form that resumes the use of the CKPT. Recall that these are 2 of the simpler forms of the reconfiguration dialog.

First stop using CKPT1 by entering the reconfiguration dialog and selecting option 5 "SUSPEND THE USE OF CKPT1".

Now that no-one is using CKPT1, it can be safely deleted and reallocated with a larger size. However before you delete anything, remember to verify that you are deleting the correct data set. Issue a `$DCKPTDEF,CKPT1` to display the current CKPT1 data set name AND volser. Double check it.

Once a new data set has been created, DO NOT FORGET to resume using CKPT1 by entering the reconfiguration dialog and selection option 7 "UPDATE AND START USING CKPT1".

At this point, a `$DCKPTSPACE` should show the updated size.

Initialization Reconfiguration



2 flavors of reconfiguration

- Automatically when JES2 is confused
 - ▶ JES2 cannot find the checkpoint

- Specify PARM=RECONFIG
 - ▶ Allows parameters to be overridden
 - ▶ Only valid on all-member-warm start
 - ▶ JES2 will confirm new values
 - ▶ What you specify is what you get
 - ▶ If changing CKPT specifications, you should specify which data set to read from

There are 2 processes commonly referred to as the initialization reconfiguration. The first is an automatic process that is entered when JES2 cannot find the checkpoint data sets and the second a manual process that actually allows you to change where the checkpoints are located.

The automatic process is often entered unexpectedly when there is an error. It can be very confusing if you are not familiar with the process or are not aware what JES2 is trying to do.

The manual process allows installations to alter the checkpoint configuration during a JES2 all-member warm start. Specifying PARM=RECONFIG instructs JES2 to override the checkpoint parameters it reads from the checkpoint and use the values from initialization instead. JES2 will confirm the checkpoint configuration before it uses them. RECONFIG is a very powerful option. If not used properly, this can damage the CKPT data sets and require a cold start.

At any time, if this is an all-member warm start, the automatic process can become a reconfiguration process.

Automatic Initialization Reconfiguration



- Reconfiguration can be entered because initialization cannot locate CKPTs
- JES2 will not restart until it has allocated the same checkpoints it had when JES2 came down
 - ▶ This is the key area of confusion!
- Also entered if error reading checkpoints
- Can be transformed into a PARM=RECONFIG form
- Note: assumes APAR OW56648 applied

This form of the initialization reconfiguration is entered automatically by JES2 initialization. It is a cry for help from the initialization process. The goal of JES2 initialization is to get back to the CKPT data sets that were in use when JES2 came down. How does JES2 know where the checkpoint was when it came down? There is a data area in EACH checkpoint data set that contains the specifications for CKPT1 and CKPT2. There is also the instorage copy that JES2 is using at the time (started with what was specified at initialization, updated with what we get from other members). JES2 will not start unless all 3 copies of the checkpoint specifications agree.

If we are in this form of reconfiguration, then for any one of a number of reasons, JES2 cannot figure out how to get a consistent checkpoint specification. Or it could be that the reply an operator gave to a WTOR lead initialization to suspect the checkpoint data sets are not correct. Or perhaps the forwarding information is missing or cannot be read from a data set. In any case, this form of the initialization reconfiguration is NOT trying to change anything, just put things back the way they were before.

Note: All examples assume that APAR OW56648 is applied.

Error starting JES2



```

$HASP536 MEMBER IBM1 -- CONNECT TO STRNAME SPOOLCKPT1 FOR CKPT1
FAILED, RC=00000000 RSN=00000000
ERRCODE=07, STRUCTURE NOT ALREADY USED
*$HASP284 JES2 INITIALIZATION CHECKPOINT DIALOG STARTED
*$HASP277 JES2 CAN NOT FIND OR USE THE CHECKPOINT DATA SET(S)
THAT WERE IN USE WHEN THE MAS WAS LAST ACTIVE BECAUSE CKPT1
STRUCTURE NOT ALREADY USED

CURRENT CHECKPOINT VALUES:
CKPT1=(STRNAME=SPOOLCKPT1, INUSE=YES),
CKPT2=(DSNAME=SYS1.JESCKPT2, VOLSER=CKTPK, INUSE=YES)

VALID RESPONSES ARE:
CKPT1=... - UPDATE CURRENT CHECKPOINT SPECIFICATION WITH
CKPT2=... THE VALUES USED WHEN JES2 MAS WAS LAST ACTIVE
'RECONFIG' - THE CHECKPOINT DATA SET(S) THAT WERE IN USE
            WHEN JES2 WAS LAST ACTIVE ARE NO LONGER
            AVAILABLE (ALL-MEMBER WARM START ONLY)
'CONT'    - ATTEMPT INITIALIZATION WITH THE VALUES LISTED
'TERM'    - TERMINATE JES2 INITIALIZATION ON THIS MEMBER
*06 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP277 FOR RELATED MSG)

```

Setup is my system was up and running with CKPT1 on a CF (structure name of SPOOLCKPT1) and CKPT2 is on DASD. JES2 was shutdown and so was MVS for a maintenance window. I IPLed and am warm starting JES2. Initialization seems to be having a problem reading data from the CF. The problem is, while the system was down, I applied maintenance to my CF and did a POR (Power on reset). Since I did not move my CKPT off the coupling facility, I lost all my data. But I still have CKPT2 on DASD. At this point, JES2 has failed to read CKPT1 but has read CKPT2. CKPT2 agrees that there is a CKPT1 but JES2 cannot validate that CKPT1 is a checkpoint. The operator is being asked if everything is OK. This is why it is important to keep track of where the checkpoint data set is supposed to be.

Error starting JES2

(cont...)



36,cont

```

IEE600I REPLY TO 36 IS:CONT
$HASP290 MEMBER IBM1 -- JES2 CKPT1 IXLLIST READ_LIST REQUEST FAILURE
*** CHECKPOINT DATA SET NOT DAMAGED BY THIS MEMBER ***
RETURN CODE      = 00000008
REASON CODE     = 0C1C0825
RECORD          = CHECK
*$HASP460 UNABLE TO CONFIRM THAT CKPT1 IS A VALID CHECKPOINT
DATA SET DUE TO AN I/O ERROR READING THE LOCK RECORD.
VERIFY THE SPECIFIED CHECKPOINT DATA SETS ARE CORRECT:

VALUES FROM CKPTDEF
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES),
CKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=CKPTPK,INUSE=YES)

VALUES JES2 WILL USE
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES),
CKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=CKPTPK,INUSE=YES)

DATA SET JES2 COULD NOT CONFIRM
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES)
*37 $HASP417 ARE THE VALUES JES2 WILL USE CORRECT? ('Y' OR 'N')

```

The operator looked at the log book (or wherever you have the current checkpoint locations) and agrees that this is where the checkpoint data sets should be. He replies continue. JES2 will try one more time to read the data from the damaged checkpoint. JES2 reports the failure and asks one more time, are these the correct data sets. At this point it lists what was specified in the init deck, what JES2 will use, and what data set had a problem.

Error starting JES2 (cont...)



37y

```
IEE600I REPLY TO 37 IS;Y
$HASP478 INITIAL CHECKPOINT READ IS FROM CKPT2
(SYS1.JESCKPT2 ON CKTPK)
LAST WRITTEN FRIDAY, 1 AUG 2003 AT 16:25:01 (GMT)
*$HASP493 JES2 ALL-MEMBER WARM START IS IN PROGRESS
$HASP537 THE CURRENT CHECKPOINT USES 322 4K RECORDS
$HASP492 JES2 ALL-MEMBER WARM START HAS COMPLETED
```

Finally, after the operator replied that the data sets were correct, JES2 will initialize.

Error starting JES2

Case 2



```
$HASP536 MEMBER IBM1 -- CONNECT TO STRNAME SPOOLCKPT1 FOR CKPT1
FAILED, RC=00000000 RSN=00000000
ERRCODE=07, STRUCTURE NOT ALREADY USED
*$HASP284 JES2 INITIALIZATION CHECKPOINT DIALOG STARTED
*$HASP277 JES2 CAN NOT FIND OR USE THE CHECKPOINT DATA SET(S)
THAT WERE IN USE WHEN THE MAS WAS LAST ACTIVE BECAUSE CKPT1
STRUCTURE NOT ALREADY USED

CURRENT CHECKPOINT VALUES:
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES),
CKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=CKPTPK,INUSE=YES)

VALID RESPONSES ARE:
CKPT1=... - UPDATE CURRENT CHECKPOINT SPECIFICATION WITH
CKPT2=...   THE VALUES USED WHEN JES2 MAS WAS LAST ACTIVE
'RECONFIG' - THE CHECKPOINT DATA SET(S) THAT WERE IN USE
            WHEN JES2 WAS LAST ACTIVE ARE NO LONGER
            AVAILABLE (ALL-MEMBER WARM START ONLY)
'CONT'    - ATTEMPT INITIALIZATION WITH THE VALUES LISTED
'TERM'    - TERMINATE JES2 INITIALIZATION ON THIS MEMBER
*59 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP277 FOR RELATED MSG)
```

In case 2 we start JES2 under the same conditions as the previous case.

Error starting JES2

Case 2 (cont...)



59,cont

```
IEE600I REPLY TO 59 IS:CONT
$HASP290 MEMBER IBM1 -- JES2 CKPT1 IXLLIST READ_LIST REQUEST FAILURE
*** CHECKPOINT DATA SET NOT DAMAGED BY THIS MEMBER ***
RETURN CODE   = 00000008
REASON CODE   = 0C1C0825
RECORD        = CHECK
*$HASP460 UNABLE TO CONFIRM THAT CKPT1 IS A VALID CHECKPOINT
DATA SET DUE TO AN I/O ERROR READING THE LOCK RECORD.
VERIFY THE SPECIFIED CHECKPOINT DATA SETS ARE CORRECT:

VALUES FROM CKPTDEF
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES),
CKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=CKPTPK,INUSE=YES)

VALUES JES2 WILL USE
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES),
CKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=CKPTPK,INUSE=YES)

DATA SET JES2 COULD NOT CONFIRM
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES)
*60 $HASP417 ARE THE VALUES JES2 WILL USE CORRECT? ('Y' OR 'N')
```

Still looks the same as the previous case. However, I forgot to tell the operator something.....

Error starting JES2 Case 2 (cont...)



```
60y
IEE600I REPLY TO 60 IS;Y
$HASP478 INITIAL CHECKPOINT READ IS FROM CKPT2
(SYS1.JESCKPT2 ON CKPTPK)
LAST WRITTEN FRIDAY, 1 AUG 2003 AT 17:33:03 (GMT)
*$HASP050 JES2 RESOURCE SHORTAGE OF TGS - 94% UTILIZATION REACHED
$HASP406 RACF      WAS EXECUTING
$HASP096 DISASTROUS ERROR AT SYMBOL CBIMPL4 IN CSECT HASPNUC,
MTTR=0001A401,UNIT=0982,VOLSER=SPOOL1
$HASP096 DISASTROUS ERROR AT SYMBOL CBIMPL4 IN CSECT HASPNUC,
MTTR=0001AC06,UNIT=0982,VOLSER=SPOOL1
$HASP073 RECOVERY SUCCESSFUL - NORMAL PROCESSING RESUMES
$HASP096 DISASTROUS ERROR AT SYMBOL CBIMPL4 IN CSECT HASPNUC,
MTTR=0001AC06,UNIT=0982,VOLSER=SPOOL1
$HASP073 RECOVERY SUCCESSFUL - NORMAL PROCESSING RESUMES
$HASP096 DISASTROUS ERROR AT SYMBOL CBIMPL4 IN CSECT HASPNUC,
MTTR=0001B406,UNIT=0982,VOLSER=SPOOL1
RECOVERY SUCCESSFUL - NORMAL PROCESSING RESUMES
$HASP455 JOBT07   IS PURGED, CONTROL BLOCK READ ERROR
```

I forgot to mention that last night we had some problems with a bunch of DASD and had to move the CKPT2 data set temporarily. I was on the phone with operations, it was late, they just wanted to get some jobs run before they shut down the system for a maintenance window. They didn't mention that they were going to restart the coupling facilities. And they did fix the DASD problems. Now they are getting hundreds of errors and I am on the phone instead of sleeping in on Sunday morning.

Error starting JES2 Case 2 (cont...)



```
*$HASP863 MORE THAN 10 SPOOL I/O ERRORS OCCURRED. SHOULD JES2
CONTINUE THE WARM START?
```

```
VALID RESPONSES ARE:
```

```
'CONTINUE' - CONTINUE THE WARM START
'ABEND'     - TERMINATE JES2 WITH $V02 ABEND
*61 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP863 FOR RELATED MSG)
61,abend
IEE600I REPLY TO 61 IS;ABEND
*$HASP095 JES2 CATASTROPHIC ERROR. CODE = $V02
$HASP088 JES2 ABEND ANALYSIS
$HASP088 -----
$HASP088 FMID      = HJE7707      LOAD MODULE = HASJES20
$HASP088 SUBSYS   = JES2   z/OS 1.4
$HASP088 DATE     = 2003.213      TIME      = 14.03.52
$HASP088 DESC    = TERMINATION DUE TO EXCESSIVE SPOOL ERRORS
```

OH THANK GOODNESS I MIGRATED TO z/OS 1.4!

The message is telling the operators they messed up. The checkpoints haven't been damaged and all they need to do is reply ABEND and JES2 shuts down without writing the checkpoints. I can re-start and not loose all the jobs in the system. 5 minutes and I am back to sleep.

Not running z/OS 1.4, all is not lost. Since JES2 knows nothing about the correct checkpoint its contents are still intact. However the longer the system runs with the wrong checkpoint, the more jobs that will be lost as their SPOOL control blocks are overwritten. If you quickly stop the system and a restart with the correct checkpoint should recover most of the jobs.

Error starting JES2 Case 2 (cont...)



S JES2.....

```
$HASP290 MEMBER IBM1 -- JES2 CKPT1 IXLLIST READ_LIST REQUEST FAILURE
*** CHECKPOINT DATA SET NOT DAMAGED BY THIS MEMBER ***
RETURN CODE   = 00000008
REASON CODE   = 0C1C0825
RECORD        = CHECK
*$HASP460 UNABLE TO CONFIRM THAT CKPT1 IS A VALID CHECKPOINT
DATA SET DUE TO AN I/O ERROR READING THE LOCK RECORD.
VERIFY THE SPECIFIED CHECKPOINT DATA SETS ARE CORRECT:

VALUES FROM CKPTDEF
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES),
CKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=CKPTPK,INUSE=YES)

VALUES JES2 WILL USE
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES),
CKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=CKPTPK,INUSE=YES)

DATA SET JES2 COULD NOT CONFIRM
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES)
*69 $HASP417 ARE THE VALUES JES2 WILL USE CORRECT? ('Y' OR 'N')
```

Tell operations to restart JES2. They do and are right back where we started (almost, the reason we are here is slightly different).

Error starting JES2

Case 2 (cont...)



69,n

```

IEE600I REPLY TO 69 IS;N
*$HASP284 JES2 INITIALIZATION CHECKPOINT DIALOG STARTED
*$HASP277 JES2 CAN NOT FIND OR USE THE CHECKPOINT DATA SET(S)
THAT WERE IN USE WHEN THE MAS WAS LAST ACTIVE BECAUSE OF
REPLY OF 'N' TO HASP417 MESSAGE

CURRENT CHECKPOINT VALUES:
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES),
CKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=CKPTPK,INUSE=YES)

VALID RESPONSES ARE:
CKPT1=... - UPDATE CURRENT CHECKPOINT SPECIFICATION WITH
CKPT2=... THE VALUES USED WHEN JES2 MAS WAS LAST ACTIVE
'RECONFIG' - THE CHECKPOINT DATA SET(S) THAT WERE IN USE
            WHEN JES2 WAS LAST ACTIVE ARE NO LONGER
            AVAILABLE (ALL-MEMBER WARM START ONLY)
'CONT'    - ATTEMPT INITIALIZATION WITH THE VALUES LISTED
'TERM'    - TERMINATE JES2 INITIALIZATION ON THIS MEMBER
*70 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP277 FOR RELATED MSG)

```

This time, I tell the operator to reply 'N' because the CKPT2 specification is wrong. 'CKPTPK' is the bad pack that we moved off of last night. I could have updated the init deck but I would have had to get out of bed and go down to the terminal. Good thing I never go to sleep without a JES2 messages manual under my pillow. Besides, TSO is not up and it takes too long to IPL the emergency system.

Error starting JES2 Case 2 (cont...)



```
70,ckpt2=vol=spool1
```

```
IEE600I REPLY TO 70 IS;CKPT2=VOL=SPOOL1  
*$HASP277 JES2 CAN NOT FIND OR USE THE CHECKPOINT DATA SET(S)  
THAT WERE IN USE WHEN THE MAS WAS LAST ACTIVE BECAUSE OF  
REPLY OF 'N' TO HASP417 MESSAGE
```

```
CURRENT CHECKPOINT VALUES:
```

```
CKPT1=(STRNAME=SPOOLCKPT1,INUSE=YES),  
CKPT2=(DSNAME=SYS1.JESCKPT2,VOLSER=SPOOL1,INUSE=YES)
```

```
VALID RESPONSES ARE:
```

```
CKPT1=... - UPDATE CURRENT CHECKPOINT SPECIFICATION WITH  
CKPT2=... THE VALUES USED WHEN JES2 MAS WAS LAST ACTIVE  
'RECONFIG' - THE CHECKPOINT DATA SET(S) THAT WERE IN USE  
WHEN JES2 WAS LAST ACTIVE ARE NO LONGER  
AVAILABLE (ALL-MEMBER WARM START ONLY)  
'CONT' - ATTEMPT INITIALIZATION WITH THE VALUES LISTED  
'TERM' - TERMINATE JES2 INITIALIZATION ON THIS MEMBER  
*71 $HASP272 ENTER RESPONSE (ISSUE D R,MSG=$HASP277 FOR RELATED MSG)
```

OK, just tell them to update the CKPT2 volser and everything looks good.

Error starting JES2

Case 2 (cont...)



71, cont

```
IEE600I REPLY TO 71 IS:CONT
$HASP290 MEMBER IBM1 -- JES2 CKPT1 IXLLIST READ_LIST REQUEST FAILURE
*** CHECKPOINT DATA SET NOT DAMAGED BY THIS MEMBER ***
RETURN CODE   = 00000008
REASON CODE   = 0C1C0825
RECORD        = CHECK
*$HASP460 UNABLE TO CONFIRM THAT CKPT1 IS A VALID CHECKPOINT
DATA SET DUE TO AN I/O ERROR READING THE LOCK RECORD.
VERIFY THE SPECIFIED CHECKPOINT DATA SETS ARE CORRECT:

VALUES FROM CKPTDEF
CKPT1=(STRNAME=SPOOLCKPT1, INUSE=YES),
CKPT2=(DSNAME=SYS1.JESCKPT2, VOLSER=CKPTPK, INUSE=YES)

VALUES JES2 WILL USE
CKPT1=(STRNAME=SPOOLCKPT1, INUSE=YES),
CKPT2=(DSNAME=SYS1.JESCKPT2, VOLSER=SPOOL1, INUSE=YES)

DATA SET JES2 COULD NOT CONFIRM
CKPT1=(STRNAME=SPOOLCKPT1, INUSE=YES)
*72 $HASP417 ARE THE VALUES JES2 WILL USE CORRECT? ('Y' OR 'N')
```

Argh, JES2 is asking me to again confirm the checkpoint specifications. Why doesn't JES2 just believe what I said. Oh yea, things did get messed up last time. Better have them read the values one more time just to be safe.

Error starting JES2 Case 2 (cont...)



72y

```
IEE600I REPLY TO 72 IS;Y
$HASP478 INITIAL CHECKPOINT READ IS FROM CKPT2
(SYS1.JESCKPT2 ON SPOOL1)
LAST WRITTEN FRIDAY, 1 AUG 2003 AT 17:35:27 (GMT)
*$HASP493 JES2 ALL-MEMBER WARM START IS IN PROGRESS
$HASP537 THE CURRENT CHECKPOINT USES 171 4K RECORDS
$HASP492 JES2 ALL-MEMBER WARM START HAS COMPLETED
```

Everything looks good, a final reply and everything takes off. No more errors and I can get some sleep.

I sure am glad I went to that reconfig presentation at SHARE!

System managed rebuild



- z/OS function to move structures from one coupling facility to another
- Operator initiated (SETXCF START,REBUILD)
- Only works if there is no existing problem
 - ▶ If connectivity is lost, cannot use
- Only moves to another CF (not to DASD)
- For normal maintenance, makes JES2 consistent with other CF exploiters

System managed rebuild is not intended as a replacement for CKPT reconfiguration. It is an attempt to make JES2's usage of CF more consistent (operationally) with other components. System managed rebuild addresses the problem of taking a CF off-line for maintenance. It is mostly MVS code (with some JES2 code) that allows the data in a structure to be moved to another coupling facility. It only works if there are no problems with the current coupling facility. There do not have to be any active users of the structure to use the service (JES2 can be down and MVS can move the CKPT). The command to use the service are:

```
SETXCF START,REBUILD,STRNAME=nnnnnnnn or
```

```
SETXCF START,REBUILD,CF=nnnnnnnn
```

While the structure is being moved, JES2 cannot access the checkpoint. No jobs are started, no new SYSOUT can start printing, etc.

Using System-managed Rebuild



```
setxcf start,rebuild,strname=spoolckpt1
IXC570I SYSTEM-MANAGED REBUILD STARTED FOR STRUCTURE
SPOOLCKPT1 IN COUPLING FACILITY JESSES1
  PHYSICAL STRUCTURE VERSION: B9CE8C14 EAC11DA5
  LOGICAL STRUCTURE VERSION: B9CE8C14 EAC11DA5
  START REASON: OPERATOR-INITIATED
  AUTO VERSION: B9CE9756 214E3186
IXC367I THE SETXCF START REBUILD REQUEST FOR STRUCTURE
SPOOLCKPT1 WAS ACCEPTED.
IXC578I SYSTEM-MANAGED REBUILD SUCCESSFULLY ALLOCATED
STRUCTURE SPOOLCKPT1.
OLD COUPLING FACILITY: JESSES1
OLD PHYSICAL STRUCTURE VERSION: B9CE8C14 EAC11DA5
NEW COUPLING FACILITY: TESTCF
NEW PHYSICAL STRUCTURE VERSION: B9CE9757 3FF3E309
LOGICAL STRUCTURE VERSION: B9CE8C14 EAC11DA5
  AUTO VERSION: B9CE9756 214E3186
IXC577I SYSTEM-MANAGED REBUILD HAS
BEEN COMPLETED FOR STRUCTURE SPOOLCKPT1
STRUCTURE NOW IN COUPLING FACILITY TESTCF
  PHYSICAL STRUCTURE VERSION: B9CE9757 3FF3E309
  LOGICAL STRUCTURE VERSION: B9CE8C14 EAC11DA5
  AUTO VERSION: B9CE9756 214E3186
```

This is the series of messages you can expect to see. Notice, there are no JES2 messages about the process. MVS handles the move. The MVS messages do indicate where the structure was moved from/to. For more information on this, consult the related SYSPLEX documentation.



Other Hits and Tips

- Read messages carefully and if needed the descriptions in the books
- Specify CKPTn start option or READFROM= when altering CKPT specifications with RECONFIG start parm
 - ▶ READFROM=HIGHEST is NOT the best choice. You know which data set to read from
- Ensure operations knows what to do in the event of a reconfiguration
 - ▶ Preset NEWCKPTn and use OPVERIFY=NO

Maybe this sounds too simple but careful reading of the messages and the descriptions in the book often solves most problems with reconfiguration. There are some subtle differences in message texts that can help clarify what is going on.

Letting READFROM= default to HIGHEST has resulted in COLD starts in the past. At this point in processing, JES2 logic cannot always decide the correct data set to read from. The right garbage has fooled the logic in the past. We have taken action to prevent this but if you are changing the checkpoint specification with RECONFIG and know that only one data set has the latest checkpoint data, give that information to JES2.

If you think RECONFIG is complicated now, ask someone about how it worked prior to SP 5.1.0. But don't let the relative simplicity lull you into thinking you do not need to practice this on a test system. The more familiar you are with the process, the easier it is to deal with if you have a problem. This includes the system programmers and the operations staff.

Use OPVERIFY=NO with preset NEWCKPTn's (that was pre-allocated large enough to hold a checkpoint). This will reduce the impact of an error and prevent the middle of the night call. Automate on the HASP256 message to send a message to someone when it is issued. This way you know there has been a reconfiguration and you need to reset NEWCKPTn.

Other Hits and Tips (cont...)



- You can forward a data set to itself.
 - ▶ When you get the confirmation message, set NEWCKPTn to the same value as CKPTn
 - ▶ This rebuilds the data structures, writes out the current checkpoint data, and you continue using the same data set
 - ▶ Useful if the cause of the problem was fixed and you don't want to move to a new data set
- Keep a list of where you can place a CKPT available for operations
- Keep track of where you last placed your checkpoint

There have been times when an error was only temporary. Someone pulled the wrong cable, configured the wrong DASD off a system, etc. For these cases, if you did not automatically forward your checkpoint, you can fix the problem and specify NEWCKPT1 the same as CKPT1 and forward CKPT1. This may not be the best option but it is good to know that it is an option.

Keeping a list of places where you can safely place your checkpoint in the event of an error can greatly reduce recovery time. This is especially true if the failure involves multiple devices and not just the one the checkpoint is currently using. This same list can also be used to keep track of where the checkpoint is currently located. It does no good to move the checkpoint, save the queues, and then loose them on the next IPL because no one knows where the checkpoint went.

Questions?