

# JES2 Common Challenges



SHARE 95, Session 2665

July 26, 2000

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.

**Bob Jinkins**  
**JES2 Design/Development/Service**  
**Poughkeepsie, NY**  
 **jinkins@us.ibm.com**

# Agenda

---




## Learn from other people's misfortunes

- Plan ahead
- Solving everyday problems
- Hints, tips, tools, etc.
- First-aid for your system




# Plan ahead

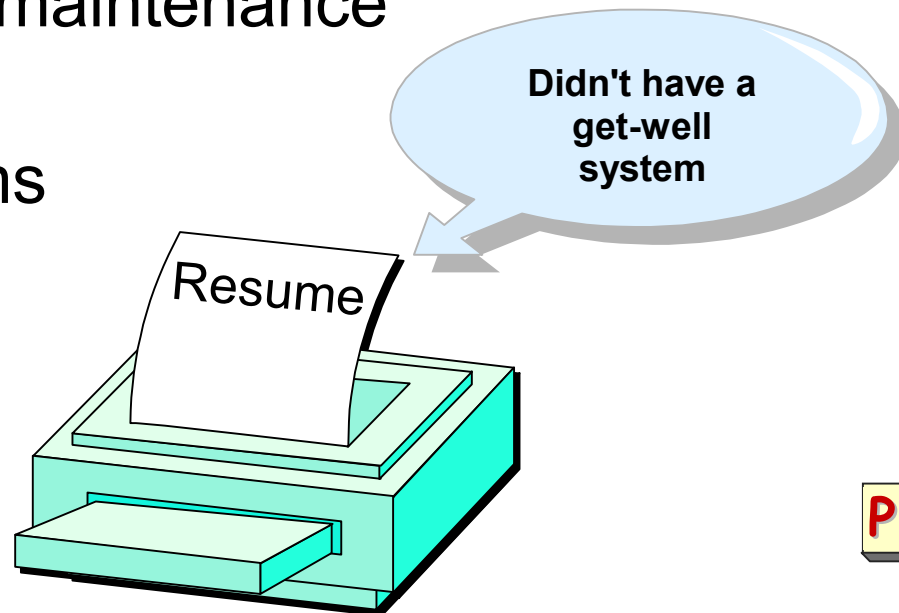
---




# Have a get-well system



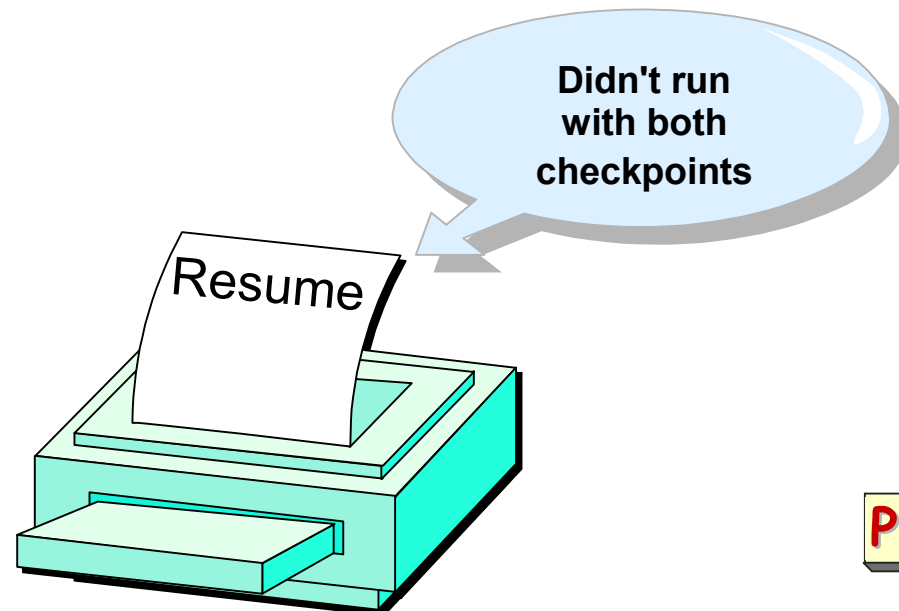
- Small separate system
  - Access to production DASD and catalogs
  - But doesn't require any production DASD
  - Access to IBM link and FTP capability desirable
- Use to
  - Apply emergency maintenance
  - Fix JES2 proc
  - Fix DASD problems
  - etc.



# Run with both CKPT1 and CKPT2



- Don't let a glitch turn into cold start
- Insignificant performance impact
- Second checkpoint **not RESERVED**



# Allocate CKPT and SPOOL as unmovable


---



- Don't let DASD defrag utilities move checkpoint and SPOOL data sets
- **\$K24** (sometimes **\$K23**) if checkpoint is moved



# Use symbolics in JES2 PROC



- Use symbolics in **PROCxx** concatenations
- Use symbolics in **HASPPARM** concatenation
- Use symbolics for everything!
- Use **HASIPROC** in **SHASSAMP** as a model
- **Tip:** Keep a hardcopy by the console
- **Tip:** Test JES2 PROC changes by starting a second copy of your production JES2 ... it won't get too far



# Don't run out of BERTs



- Jobs will hang until BERTs are available
- Don't take JES2 down or you might have to **cold start**
- **OW44400** - **BERTWARN=0** after **\$ACTIVATE**
- **OW41547** - HASPHOPE not assembled
- **OW35410** - JES2 restart problems (OPEN)
- **Warning:** BERTNUM=2000 hardcoded in **SERVERPAC** prior to 4Q98

**\$T CKPTSPACE, BERTNUM=n, BERTWARN=80**





# Why IPL when you can hot start?

---



- Separate JES2 startup automation from IPL
- Need automation to
  - Start NJE/RJE
  - Start JES mode printers



# Be able to manually start JES2



- Be able to specify any JES2 start option
  - **CKPT1**
  - **CKPT2**
  - **RECONFIG**
  - **CONSOLE**
  - etc.
- Be able to respond to any startup message
- Don't put a **\$S** in your init parms
  - Example: Purge a job before it can take JES2 down
- **ZAPJOB** (new in release 10) doesn't help if you can't manually start JES2



# Keep JES2 parms in sync with checkpoint

---




- Some mismatches cause **HASP496** message
- **JOBCLASS** info rebuilt from init parm data if BERT errors for \$CATs
- Ensure **all** members are in sync
- Be ready for an unscheduled cold start



# Use reasonable WARN thresholds


---



- Don't use WARN=0 (no warning)
- Don't specify over 90%



# Install JES2 IPCS



- System programmer action required
- See *JES2 Migration Notebook* for instructions
- See *JES2 Diagnosis* for verification procedure
- Have an IPCS setup for every combination of MVS and JES2 releases in production
- Apply **OW33073** - JES2 IPCS enhancements


Example analysis:

```
*** WARNING: Work on one or more SPOOL volumes is not
selectable ($SPLEXST≠$SPLSLCT in checkpointed HCT)
```



# Rethink your migration testing policy



---



- Over reliance on in-house testing
- Some customers freeze all changes on a new release for weeks before migrating
- Potentially applicable **HIPERs** not applied
- Symptoms show up within hours of migration
- **Get very current on a new release before migrating**



# Solving everyday problems



# Why won't this job print?



- Job status
- Job held
- SPOOL volume not available
- No output group (maybe an archiver got it)
- Output on another device
- Device status
- Work selection mismatch
- **OUTDISP=HOLD** or **LEAVE**
- **HOLD=SYSTEM** or **OPER**
- Printer not eligible per RACF (**HASP186**)
- System draining (**\$P** or **\$PJES2**)





# Why won't this job print? *(continued)*


---



- Helpful commands
  - **\$Djob, LONG**
  - **\$DOjob**
  - **\$DPRTn**
  - **\$DSPOOL**
  - **\$S**



# Why won't this job execute?



- Job awaiting execution and not BERT locked?
- Job held
- Duplicate job name
- Affinity to systems where the job can't run
- Scheduling environment not available
- SPOOL volumes not available
- **\$PXEQ** in effect (**\$DMEMBER**, also **HASP222**)
- System draining (**\$P** or **\$PJES2**)
- Exit 14 or 49
- Independent mode (job/member must match)



# Why won't this job execute? *(continued)*

---



- **JOBCLASS(x)**
  - **QHELD=YES**
  - **XEQCOUNT=MAX=**limit exceeded
  - **MODE=JES** and an initiator is not available
  - **MODE=WLM** and WLM not in **GOAL** mode



# Why won't this job execute? *(continued)*


---



- Display commands
  - **\$Djob,DELAY**
  - **\$Djob,LONG**
  - **\$DJOBCLASS(x)**
  - **\$DMEMBER**
  - **\$DI** Watch out for **INELIGIBLE\_CLASS=(x-rsn,...)**
  - **\$DSPOOL**
  - **\$S**
  - **D WLM,...**
  - **\$DPERFDATA(SAMPDATA)**




# Why won't this job execute? *(continued)*



- Delays reported by **\$Djob,DELAY**
  - **HOLD** - job held via \$HJ, \$HA, TYPERRUN=HOLD, duplicate job name, or JOBCLASS QHOLD=YES
  - **SYSAFF** - member job has SYSAFF to not available
  - **SCHENV** - scheduling environment not available
  - **MEMBER\_STATUS** - SCHENV + SYSAFF + BOSS + \$P + \$PXEQ
  - **LIMIT** - JOBCLASS limit prevents job from starting
  - **LOCKED** - BERT lock is held somewhere
  - **BUSY\_ON\_DEVICE** - job is busy on a device
  - **SPOOLS\_NOT\_AVAILABLE**
  - **YES** - any of the above - filter only
  - **NO** - none of the above - filter only




# Why won't this job execute? *(continued)*



- Delays **NOT** reported by **\$Djob,DELAY**
  - No WLM initiators to select service class
    - ▶ Eventually, service class will no longer be meeting its goals and WLM may decide to start another init
  - Jobs are in **MODE=WLM** job classes but WLM is not in GOAL mode
  - Job class intended to be WLM but inadvertently set to **MODE=JES**
  - No JES2 initiators in job class on any system job can run on
  - All initiators are currently busy




# High CPU in JES2



- **PERFDATA** (undocumented) collects JES2 main-task data
- In a period when JES2 CPU is high
  - **\$TPERFDATA(\*),RESET** and **D A,JES2**
  - Wait 10 or 20 minutes
  - **D A,JES2** and **\$DPERFDATA**
    - Use **\$DPERFDATA,L=Z** prior to release 7
- Don't be alarmed by CPU% numbers, the PCE types add up to 100%
- Use **D A,JES2** to get CPU and elapsed time for the address space



# High CPU in JES2 (continued)



## ■ Example PERFDATA for a specific PCE type

```


$DPERFDATA(PCESTAT),PCENAME=STAC
$HASP660 $DPERFDATA(PCESTAT),PCENA
$HASP660 PCE PERFORMANCE STATISTICS - INTERVAL=2:07:19.379957,
$HASP660 CPU=30:05.945181,
$HASP660 PCENAME=STAC,TIME=51:41.505259,CPU=22:54.777037,CPU%=
$HASP660 76.12,QSUSE_TIME=50:24.967096,IOCOUNT=29,CKPT_COUNT=
$HASP660 601,
$HASP660     WAIT=IO,MOD=HASPNUC,SEQ=17000000
$HASP660     COUNT=29,AVGWAIT=0.018144,
$HASP660     POST=IO,COUNT=29,AVGWAIT=0.018144,
$HASP660     WAIT=CKPT,MOD=HASPNUC,SEQ=28410000
$HASP660     COUNT=1686,AVGWAIT=0.069127,CMOD=HASPSTAC,CSEQ=
$HASP660 13100000,
$HASP660     POST=RESOURCE,COUNT=1686,AVGWAIT=0.069127,
$HASP660     WAIT=POST,XECB,MOD=HASPNUC,SEQ=97330800
$HASP660     COUNT=29,AVGWAIT=0.007361,
$HASP660     POST=XECB,COUNT=29,AVGWAIT=0.007361,
$HASP660     WAIT=STAC,INHIBIT=NO,MOD=HASPSTAC,SEQ=09900000
$HASP660     COUNT=11226361,AVGWAIT=0.006517,
$HASP660     POST=RESOURCE,COUNT=11226361,AVGWAIT=0.006517

```





# High CPU in JES2 (continued)





- PERFDATA analysis result
  - The **STAC** PCE processes **CANCEL** and **STATUS** SSI requests queued from other address spaces
  - Application(s) made **11,226,361 requests** in a little over 2 hours
  - **1,485 requests per second**
  - JES2 is a victim
  - Look for applications that make SSI requests for function codes 2 (**CANCEL**) or 3 (**STATUS**)




# Hints, tips, tools, etc.

---




# Specify enough buffers, etc.



- JES2 is good about not backing storage when it's not in use
- Usually take defaults except as noted below
- **BUFDEF EXTBUF=LIMIT=** usually take default, but not less than 1000
- **CONDEF BUFNUM=**
  - At least 1000
  - **BUFNUM=9999, BUFWARN=50** is better
- **INTRDR RDINUM=** at least 100
  - Warm start (usually an IPL) to increase
- **SMFDEF BUFNUM=** see book, but at least 100


# CKPT reconfiguration secrets



- Don't reply **TERM** in a reconfiguration for a checkpoint error. You'll have to IPL and you might have to **cold start**
- The key to understanding initialization dialogs
  - By default, JES2 tries to find the checkpoints it used when it was last active. You tell JES2 where to start and JES2 looks for a consistent set of checkpoints
  - An **explicit** specification of **RECONFIG** on an **All-Member Warm Start** lets you tell JES2 that you know better and to use the configuration you say to use (Specify the **READFROM** checkpoint)


See session 2664, *[JES2 Checkpoint Reconfiguration Dialog](#)*

# DEBUG, a neglected test tool



- **DEBUG CKPT=YES,VERSION=YES**
  - Test with it at least some of the time
  - Less of a performance pig after SP5.2.0
  - Catches hard to diagnose bugs that hit intermittently
- **DEBUG STORAGE=YES**
  - Detects use of storage beyond end of **\$GETWORK** areas (**\$GW3**)
  - Use it in production
    - 6 to 11 additional instructions per **\$RETNWORK**
- **DEBUG SYMREC=YES**
  - Issues **HASP805** when an interesting symptom record is written

# Don't use **DEBUG=NO**



- Default for most **DEBUG** options is **NO**
- **BERT=YES** is the exception
- *First Failure Data Capture* for correctable BERT errors (**\$BR3**)
- Don't say **DEBUG=NO** in your init parms
- Switch to **BERT=NO** if too many **\$BR3** errors

# Look in EREP after JES2 testing

---



- Search detail software records for **SC1BH**
- Some records report specific errors caused by exits
- Look at the **JES2 Diagnosis** book to determine whether a record is interesting



# Specify a large enough JES2 REGION



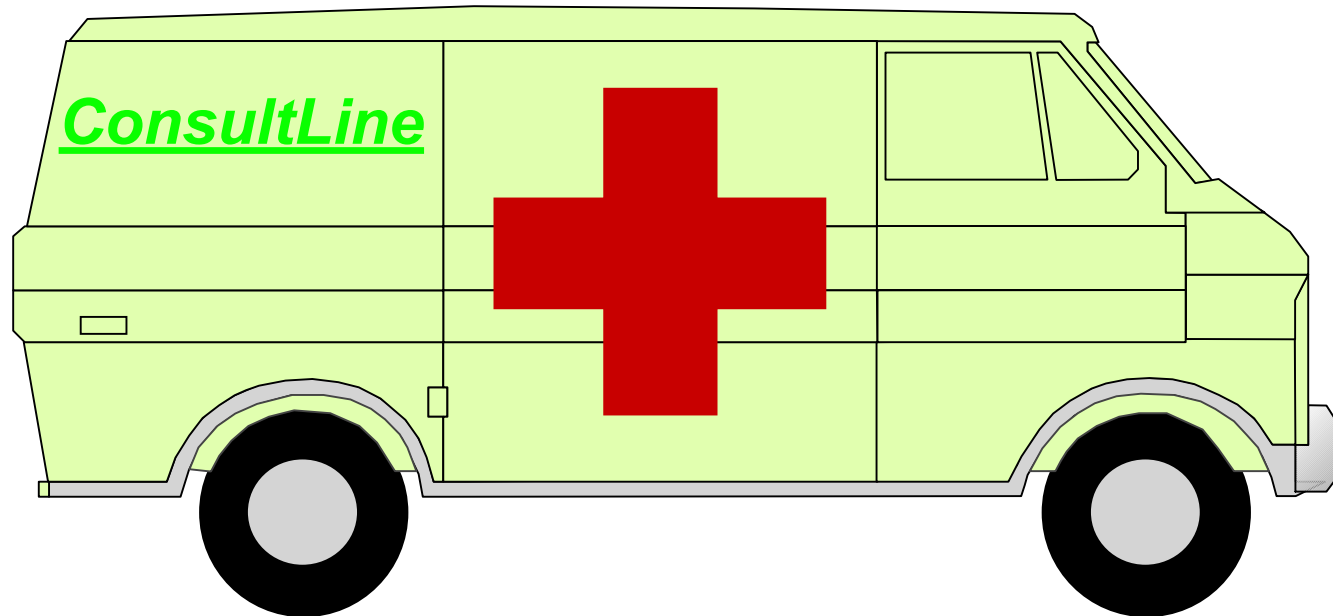

- JES2 storage increases after JES2 initialization
- Specify at least **REGION=200M**
- **0M** is better
  - But, have enough auxiliary storage
  - Defect in JES2 or an exit can use up entire region
  - **WAIT03E** if page data sets too small
  - Have seen **WAIT03E** caused by IEFUJV exit






# First-aid for your system

---




# \$S can cure the common cold



- **REQ** option forced for some initialization errors
- Operator didn't enter **\$S** after **HASP400**
- Operator accidentally entered **\$P**
- **\$PJES2** entered prematurely
  - Next **\$PJES2** doesn't work
  - Enter **\$S** before retrying **\$PJES2**
- Missing or lost POST defect
  - Job, device, etc. appears hung
  - More likely noticed on a lightly loaded JES
  - **\$S** forces POST and activity takes off



# Accidentally deleted SPOOL volume



- Don't take any members down

- **\$DSPPOOL(SPOOL1),UNITDATA**

```
$HASP893 VOLUME(SPOOL1) UNITDATA=(EXTENT=00,TRKRANGE=(002D,
$HASP893 826D),RECMAX=12,TRKPERCYL=15)
$HASP646 48.1904 PERCENT SPOOL UTILIZATION
```


- ▶ Start track X'**002D**' = **45**
  - ▶ Last track X'**826D**' = **33,389**
  - ▶ Extent size = **33,389 - 45 + 1 = 33,345**
  - ▶ If the JES isn't up, get info from AMASPZAP dump of ckpt from get-well system
- Must reallocate in exactly the same spot

```
//STEP1 EXEC PGM=IEFBR14
//DD1 DD DSN=SYS1.HASPACE,DISP=(,KEEP),UNIT=3390,
// VOL=SER=SPOOL1,SPACE=(ABSTR,(33345,45))
```

- Forget it if RACF says *erase-on-scratch*




# Consoles not processing commands



- Most MVS commands don't respond
- Related to WTO buffer shortage?  
IEA404A SEVERE WTO BUFFER SHORTAGE - 100% FULL
- Likely JES2 involvement?  
IEE767A SEVERE BUFFER SHORTAGE FOR SYSTEM LOG - 100% FULL
- JES2 hung terminating or (re)starting?
- Issue **K M,LOGLIM=0**, possibly many times
  - For example, if **K M,REF** shows LOGLIM=001000, you will have to issue the command over 40 times
  - **IEE767A** disappears and reappears
  - **TIP:** Specify **INIT LOGLIM(50000)** in **CONSOLxx**
- When **IEE767A** is gone and **D C,B** shows no backlog, get JES2 restarted in a hurry



# Started JES2 with wrong checkpoint



- Can only happen on All-Member Warm Start
- Both good and bad CKPTs use same SPOOL
- Thousands of errors while starting
- Few if any jobs left after start completed
- If CKPT you should have used is still available
  - Take JES2 down very quickly
  - SPOOL damage mostly limited to track groups used by jobs created **after** starting with the wrong CKPT (for example, track groups used for latest SYSLOG)
  - Re-IPL and start with the correct checkpoint
  - Expect some CBIMPL4 and other errors for a while
- Less likely in Rel 10 and after OW43697



# Questions ?

