



International Technical Support Organization

z/OS Planned Outage Avoidance SHARE Orlando 2011 Session 9730

www.ibm.com/redbooks

Preparer: Frank Kyne
Presenter: Marna Walle



Redbooks[®]

IBM ITSO - International Technical Support Organization

© 2011 IBM Corporation. All rights reserved.

Acknowledgements

This presentation is based partially on material in the *Planned Outage Avoidance Checklist* Redbook, written by:

- Paola Bari
- Brad Habbershaw
- Vivian Roberts
- Simon Truby
- Yvonne Zemotel

Plus invaluable help from Peter Relson, and updated with enhancements that have been introduced since that Redbook.

Intro

z/OS system programmers are presented with a challenge by their users:

"Every time you IPL, the downtime for my application is too long, so stop doing IPLs."

"Oh, by the way, when are you going to install that product for me?"

IBM helps you address this challenge in a number of ways....

Addressing this challenge

1) IPL whenever you want, for as long as you want

– Because there are situations where you **MUST** have planned IPLs:

- Upgrade to new z/OS Release or new processor
- Mass application of service

Parallel Sysplex data sharing and dynamic workload balancing.

End users typically don't know or care about *systems* - they care about their *applications*.

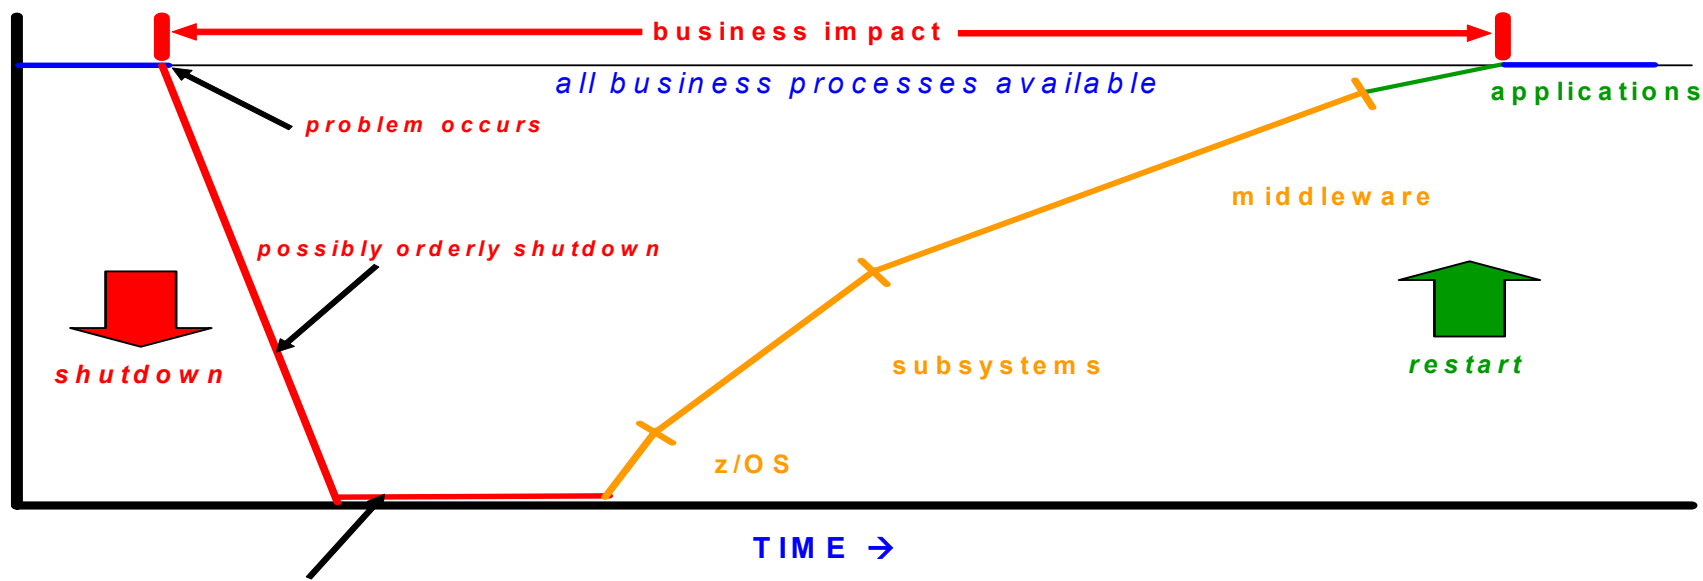
Exploiting data sharing and workload balancing lets you decouple application availability from the availability of a given system. So you can give your users the availability *they* need, and still have the system outages *you* need.

Addressing this challenge

2) Reduce elapsed time for IPLs

IBM's Mean Time To Recovery reduction initiative.

Starting with z/OS 1.10, IBM has been specifically working to reduce the amount of time it takes to stop and restart z/OS and its major subsystems.



For Unplanned - gather diagnostic information

For Planned - implement planned changes

Addressing this challenge

3) Completely eliminate some IPLs

By continuing to reduce the number of "things" that require an IPL in order to implement or change them. This is the subject of this session.

- This session concentrates on enhancements in z/OS. But note that there are also many enhancements in the major subsystems that are NOT covered in here.

Planned outage avoidance

IBM has been steadily working on reducing the number of situations where a planned IPL is required:

- By providing the ability to *dynamically* change things about the system that used to require an IPL.
- By providing the ability to change subsystems (like TCP, VTAM, CICS, DB2, etc) without having to restart them:
 - e.g. a TCP or RRS outage is often viewed as equivalent to an IPL.
- By providing the ability to plan ahead to accommodate future *non-disruptive* growth - software and hardware support.
- By improving error recovery so that an IPL is no longer required to recover from certain errors.
- By changing the system so you can survive for longer between IPLs - improve handling of non-reusable address spaces, for example.

Planned outage avoidance

So, why don't you know about all these goodies??

- Because of the detailed nature of this work, the changes often do not grab the headlines in the announcement letters, so many people are unaware of these changes.
 - Some enhancements might not even be mentioned in the announcement.
- Due to workload, lack of staff, and the time required to roll new releases out to all images, many customers skip releases and may not be aware of features that came out in intervening releases.
- People simply don't have the time to do all the research:
 - We had 4 residents for 4 weeks to create the Planned Outage Avoidance Redbook, together with the help of the z/OS developers and ITSO sysprogs and we probably *still* missed some things...

Planned outage avoidance

Example: Installation of Netview for the first time. Requires:

- Adding modules to LPA
- Adding libraries to LNKLIST
- Adding libraries to APF
- Updating the Program Properties Table
- Updating the Subsystem Names Table
- Update REXX Environment Variables (IRXANCHR)
- Adjust RSVNONR value
- Add system symbols
- EMCS consoles

Question: How many of these items still require an IPL?

Planned outage avoidance

Example: You urgently need a fix, however the fix updates a load module that resides in LPA.

Question: Do you need an IPL with CLPA to pick up the service?
How would you know?

Planned outage avoidance

Example: You are in the process of migrating all your HFS files over to zFS and are coming near the end of the migration. To complete it, and remove the last HFS, you need to migrate the sysplex root HFS over to zFS.

Question: How disruptive is this? Do you need to stop just TCP and USS? Or the whole system? Or the whole sysplex?

Planned outage avoidance

Example: Even though z/OS supports dynamic reconfiguration ability, Standalone CFs do not. And when you define a CF link, you must specify the AID of the HCA. This means that you can't predefine CF links.

Enhancement: In z/OS 1.13, on the z/OS end of the link, you can specify an AID of *. So, as long as you fully configure the CF in advance, you can now subsequently add InfiniBand adaptors to the z/OS CPC, do a dynamic reconfig of just that box, and bring the new links online without PORing the standalone CF.

Planned outage avoidance

Many customers perform regular, planned IPLs - we call these "Therapeutic IPLs":

- To harden dynamic changes.
- To address storage creep/fragmentation.
 - You should track and trend and IPL based on actuals.
 - If you have a storage creep problem, **get the vendor to fix it!**
- To recover non-reusable ASIDs (IEF352I) or non-reusable LXs.
- For operator training (!!)
- To maintain the ability to do an IPL if you really do need it (users get used to the systems being unavailable every nth weekend).
- Because we have always done it this way!

Planned outage avoidance

z/OS "things" that can be changed dynamically:

- APF List (SET PROG, SETPROG)
- LNKLIST (SET PROG, SETPROG)
- LPA - Add, delete, and (sometimes) update modules
- Exits (SET PROG, SETPROG)
- Subsystems (SETSSI ADD,S=ssn)
- System Symbols (IEASYMUP in z/OS 1.6)
- APPC settings (SETQAPPC command)
- Local time (SET CLOCK or SET TIMEZONE)
- Number of page data sets (PAGEADD/PAGEDEL)
- PPT (SCHEDxx - SET SCH=xx)
- RACF Started Class, CDT, database templates
- SVCs

Planned outage avoidance

z/OS "things" that can be changed dynamically:

- Params in DEVSUP member (SET DEVSUP=xx)
- Params in ALLOCxx member (SETALLOC)
- Many GRS-related attributes
- JES2 (exits), JES3 (TCP NJE nodes and spool volumes)
- Sysplex root - switch to new data set, define backup data set
- Change GRS mode from GR SRNL=EXCLUDE to full RNLs without a sysplex IPL
- IGGPRE00 and IGGPOST0 exits
- Turn RRS archive logging on and off dynamically
- CSVLLIX1 and CSVLLIX2 LLA exits
- Many more - see Planned Outage Avoidance Redbook and other Redbooks for the excruciating details!

Planned outage avoidance - Dynamic APF

Dynamic APF

- Dynamic APF is usually updated when installing a new product which requires additional data sets to be APF-authorized.
- There is no limit on the number of Dynamic APF libraries, however there IS a limit of 255 Static APF libraries.
- Remember that libraries no longer need to be APF-authorized to be in LPALST, however modules retrieved from LPA are treated as if they ARE APF-authorized - so protect them accordingly.
- Symbolic Alias Facility cannot be used with libraries that are added to APF at IPL time as CAS is not initialized when the APF list is processed - so you must specify fully qualified name.
- Comments:
 - By now, everyone should be using dynamic APF (PROGxx rather than IEAAPFxx).

Planned outage avoidance - Dynamic LNKLST

Changes to LNKLST are typically used when:

- Adding a new product that requires additional data set(s) in the LNKLST:
 - It is hoped that such new data sets are (for the most part) to be used by jobs that have not yet started.
- You need to compress or do some other processing on a data set that is part of LNKLST:
 - See "Changing the current LNKLST set" section in z/OS Init & Tuning Reference for procedures to remove ENQs, compress, and delete libraries from LNKLST.
- Remember that you are limited to 255 extents for LNKLST data sets:
 - Each *extent* of a PDS counts as 1.
 - A PDSE *data set* counts as 1, regardless of how many extents it actually consists of.

Planned outage avoidance - Dynamic LNKLST

LNKLST sets:

- Jobs or address spaces continue to use their current LNKLST set until the job ends or the LNKLST set for the job or address space is updated using the `SETPROG LNKLST,UPDATE,JOB=aaaaa` option.
 - Remember that `SETPROG LNKLST,ADD` is NOT the same as `SETPROG LNKLST,UPDATE`.
 - The `UPDATE` option should be used with *extreme* care.
 - ▶ If the new library will only be used by address spaces that have not started yet, this should not be an issue as you do NOT need to use the `UPDATE` command
 - ▶ If you want to use a newly LNKLSTed library with a subsystem added by Dynamic SSI (`initrtn=`), you **MUST** do an `UPDATE` because `MASTER` must see the new LNKLST in order for the dynamic subsystem add to work - suggest `UPDATE,JOB=*MASTER*` in this case to minimize the risk
 - There is no situation where it is 100% safe to issue the `UPDATE` command

Planned outage avoidance - Dynamic LPA

Dynamic LPA is typically used when:

- Installing a new product that needs things in LPA.
- When a product has items that it needs in common storage that must reside in a PDSE.
- Use the SETPROG LPA command to *replace* modules only where the owning product specifically states that Dynamic LPA is supported.

Otherwise, replacement could result in partial updates:

- If the owning product has already saved the module address, the system will NOT conduct an LPA search and will NOT find the updated module.
- Also, the addresses of all modules that are accessed via a program call (PC) instruction are stored in the PC table. That table is NOT updated by the SETPROG LPA command. Therefore, these modules cannot be replaced using the SETPROG LPA command. You must IPL for the updates to take effect.

Planned outage avoidance - Dynamic LPA

```
SETPROG LPA,ADD,MODNAME=(m1,...,mN)|MASK=mmm*,  
DSNAME=d
```

- Add the named module(s) to LPA from the named data set.
- Important to add, in the same operation, module and all its aliases.
 - Note that starting with z/OS 1.12, all aliases will AUTOMATICALLY be added by the ADD command (as long as you specify ADDALIAS).
 - Prior to z/OS 1.12, you had to explicitly specify each alias.
- Can do an ADD for a module that is already in LPA:
 - Modules added to the system by dynamic LPA processing are placed into CSA or ECSA storage. Therefore, it is important to ensure that the system CSA and ECSA sizes are adequately defined to handle the additional consumption of CSA storage resulting from the issuance of the dynamic LPA request. Further protection can be gained through the use of the CSAMIN parameter.

Planned outage avoidance - Dynamic LPA

LPA and PDSE....

It is not possible to load modules from a PDSE into LPA as part of IPL.

- Yet there are products that require modules from PDSEs to reside in LPA.

z/OS 1.12 added the LPA Deferred Wait capability:

- LPA ADD statements found in PROGxx member processed during IPL will be deferred and processed later in the IPL.
 - Prior to 1.12, such statements were ignored.
- To exploit this, application code must be changed to specify that requests for modules from LPA should not be processed until all deferred requests have been processed.

Planned outage avoidance - Dynamic LPA

Things not available:

- Safe Delete from LPA

Planned outage avoidance SVCs

Adding/updating SVCs

- There are two aspects to installing a new SVC:
 - Must get the SVC load module into the system
 - Must update the SVC table
- Dynamic LPA should be usable for the former if the module is in LPA (and could be used if it's in the nucleus if the CSECT has no external references, specifying to page-fix the LPA module).
- SVCUPDTE can be used to update the SVC table.
- If updating an existing SVC, you must use Dynamic LPA to load the new code, and SVCUPDTE to update the SVC table with the new address.
- Documented in Auth Assembler Services Guide and Auth Assembler Services Reference.
- Some changes in this area in z/OS 1.12 (easier, can update the new address in the SVC table, but not attributes)

Planned outage avoidance SVCs

Sample program to invoke SVCUPDTE

```

SVCUP  CSECT
SVCUP  AMODE 24
SVCUP  RMODE 24
        USING SVCUP,12
*
        SAVE (14,12),,&SYSDATE-&SYSTIME-SVDCDC
LR      12,15          BASE.
LR      14,13          SAVE HI-SAVE PTR.
LA      13,SAVE        POINT AT LO-SAVE.
ST      14,4(,13)     CHAIN LO- &..
ST      13,8(,14)     .HI-SAVE.
L       2,X'10'(,0)   POINT AT CVT.
TM      X'74'(2),X'80' TEST, WHETHER..
BZ      SVCUPA        .NOT RUNNUNG MVS/XA.
LA      2,SVCUPA      SET..
O       2,HBITOW      .31-BIT..
BSM     0,2          ADDR MODE.
SVCUPA DS      0H
*
        MODESET KEY=ZERO,MODE=SUP
*
        SVCUPDTE 216,REPLACE,TYPE=3,EPNAME=DFHCSVC
*
        MODESET KEY=NZERO,MODE=PROB
*
        SR      15,15
        L       13,4(,13)
        RETURN (14,12),RC=(15)
        DS      0F
HBITOW  DC      X'80000000'
SAVE    DC      18F'0'
*
* THE FOLLOWING MACROS MUST BE INCLUDED IN THE SOURCE PROGRAM *
* CVT      - TO MAP THE FIELD CVTNUCLU *
* IHAPSA   - TO SUPPLY CVT BASE *
* NUCLKUP- TO FIND THE SVC UPDATE SERVICE ENTRY POINT (IEAVESTU)*
        PRINT NOGEN
        IHAPSA
        IKJTCTB
        CVT    DSECT=YES
        END    SVCUP

```


Planned outage avoidance - Dynamic Exits

Dynamic Exits

- The Dynamic Exits facility (added in MVS 5.1!) lets you associate multiple exit routines with an exit point AND lets you update them dynamically.

- The following IBM exits support dynamic exits:

SMF RACF
 DynAlloc LLA
 USS DADSM
 GRS O/C/EOV
 Dump JES2
 HealthChecker
 Others

EXIT	DEF	EXIT	DEF	EXIT	DEF
SYS.IEFACTRT	E	SYSSTC.IEFACTRT	E	SYSSTC.IEFU29	E
SYSSTC.IEFU29L	I	CSVDYLPA	E	CSVDYNEX	E
HZSADDCHECK	E	IEASDUMP.QUERY	E	IEASDUMP.GLOBAL	E
IEASDUMP.LOCAL	E	IEASDUMP.SERVER	E	IEASDUMP.POSTDMP	E
IXC_ELEM_RESTART	E	IXC_WORK_RESTART	E	ISGNQXIT	E
ISGNQXITFAST	E	ISGCNFXITSYSTEM	E	ISGCNFXITSYSPLEX	E
ISGNQXITBATCH	E	ISGNQXITQUEUED1	E	ISGNQXITQUEUED2	E
ISGENDOFLQCB	E	ISGNQXITPREBATCH	E	ISGNQXITBATCHCND	E
ISGDGRSRES	E	CNZ_MSGTOSYSLOG	E	IEHINITT_EXIT	E
REKEY_EXIT	E	IEF_ALLC_OFFLN	E	IEF_SPEC_WAIT	E
IEF_VOLUME_ENQ	E	IEF_VOLUME_MNT	E	IEFDB401	E
IEF_ALLC_MOD	E	IEF_ALLC_EVENT	E	CEE_ABEND_EXIT	E
CNZ_WTOMDBEXIT	E	IEFJFRQ	E	SYSTSO.IEFUSI	E
SYSJES2.IEFUJI	E	SYSJES2.IEFACTRT	E	SYSJES2.IEFUSI	E
SYSSTC.IEFUJI	E	SYSSTC.IEFUSI	E	SYSSTC.IEFUSO	E
SYSSTC.IEFUJP	E	SYSSTC.IEFU85	E	SYSSTC.IEFU84	E
SYSSTC.IEFU83	E	SYS.IEFUAV	E	SYS.IEFUTL	E
SYS.IEFUSO	E	SYS.IEFUJP	E	SYS.IEFUSI	E
SYS.IEFUJV	E	SYS.IEFU85	E	SYS.IEFU84	E
SYS.IEFU83	E	IRREVVX01	E	IRRVAF01	E
IGDACSDX	E	BPX_PREPROC_INIT	E	BPX_POSPROC_INIT	E
BPX_IMAGE_INIT	E	BPX_PREPROC_TERM	E	CSVLLIX1	E
CSVLLIX2	E	HASP.\$EXIT0	E		

Planned outage avoidance - Dynamic SSI

Dynamic SSI:

- Dynamic SSI is used to define new subsystem interfaces (by operator command) without requiring an IPL.
- Note that if you try to ADD a subsystem whose initialization routine comes from a library that was dynamically added to LNKLST, you **MUST** issue an UPDATE JOB(*MASTER*).
- You can activate and deactivate dynamically added subsystems, but you can't delete them, nor change the attributes you specified on the SETSSI command (so make sure you get it right the first time!).
- There is no SET SSN command to get the system to read a new or updated IEFSSNxx member.

Planned outage avoidance - Page data sets

It is possible to add new LOCAL page data sets to react to an AUX shortage.

You can also remove a LOCAL page data set, for example if you need to free up the volume:


- Make sure ESQA is large enough - need x'500' bytes per used cylinder in the data set being PAGEDELED during the entire process.

Make sure that PAGTOTL is large enough - you can add or delete page data sets dynamically, but an IPL is required to change this value.

You **CANNOT** add or delete COMMON or PLPA data sets.

TSA includes automation to PAGEADD on AUX shortage msgs...

- Recommend to keep some spare pre-formatted data sets



ASM_PAGE_ADD HealthCheck



Planned outage avoidance - PPT

It is possible to add or change Program Properties Table entries dynamically:

- Create or update the SCHEDxx member with your definitions.
- z/OS comes with a number of system entries that are automatically merged with the SCHEDxx entries at IPL or when you issue SET SCH=xx
 - If there is a clash, SCHEDxx specification overrides.

Note that only PPT entries are read from SCHEDxx when you do a SET SCH - other entries, like the size of the master trace table, are NOT processed.

Planned outage avoidance - RRS log stream

Many customers don't define an Archive log stream for RRS because of the volume of data it contains and the fact that it is rarely used.

Prior to z/OS 1.10, you could run RRS without defining the Archive log stream, however enabling the log stream (maybe for IBM Service) required restarting RRS.

z/OS 1.10 added the `SETRRS ARCHIVELOGGING,ENABLE|DISABLE` command to let you turn the log stream on and off without restarting RRS.

- For more info, see [z/OS 1.10 Implementation Guide \(Redbook\)](#)

Planned outage avoidance - RACF exit

Sample RACF password phrase exit (ICHPWX11) provided by z/OS 1.9 added ability to call System Rexx.

- This lets you code the exit function in REXX (easier).
- Exit code can now be updated without requiring a system IPL.

Planned outage avoidance JES2

JES2 changes to avoid IPLs/cold starts

- The majority of the JES2 Init statements and parameters can be modified by commands and/or a single member JES2 Hot Start. Very few parms left that require a cold start to change or increase, and a small number more that require a cold start to decrease.
 - See section "*JES2 Initialization Statement and Parameter Summary Tables*" in JES2 Init & Tuning Reference, SA22-7533.
 - ▶ Basically, just changes to OWNNODE or some changes to SPOOLDEF require a Cold start
 - Make sure JES2 Parms are updated to reflect any dynamic changes (those made by command) - some parms can be *increased* dynamically, but require a cold start to *decrease* - not updating parms to match a dynamic change could result in JES2 looking for a cold start at the next IPL.

Planned outage avoidance JES2

JES2 Dynamic proclib support

- Dynamic PROCLIB concatenations can be defined in JES2PARM using PROCLIB(xxxxxx) statement rather than statically with PROCxx DD statements
- PROCLIB concatenations defined in this way can be dynamically changed (\$T PROCLIB), deleted (\$DEL PROCLIB), added (\$ADD PROCLIB), and displayed (\$D PROCLIB)
- ALSO, JES2 can be told to ignore damaged/missing data sets in the concatenation (use the UNCONDITIONAL keyword)
- NO MORE MAS-wide RESTARTS TO CHANGE PROCLIBS!
- Added way back in z/OS 1.2.

Planned outage avoidance JES2

JES2 dynamic proclib support

- Old (in JES2 JCL):

```
//PROC01 DD DSN=USER.PROCLIB1,VOL=SER=J2COM1,UNIT=3390
//      DD DSN=USER.PROCLIB2,VOL=SER=J2COM1,UNIT=3390
//      DD DSN=SYS1.PROCLIB
```

- New (in JES2PARM member):

```
PROCLIB (PROC01) DD (1) = (DSN=USER.PROCLIB1,VOLSER=J2COM1,UNIT=3390),
                  DD (2) = (DSN=USER.PROCLIB2,VOLSER=J2COM1,UNIT=3390),
                  DD (3) = (DSN=SYS1.PROCLIB)
```

Planned outage avoidance JES2


JES2 changes to avoid stopping/restarting long running tasks

- Some long running started tasks continually send output to the spool. Eventually, this could start filling the spool, but output can't be deleted until the spool file is closed.
- So, these STCs are restarted for no other reason than to free up the the spool files....
- The JOBCLASS(STC) JESLOG SPIN=spinvalue statement should be used in this case:
 - Causes spool files to be closed and reopened based on time or volume of output produced.
 - See new UNALLOC keyword in z/OS 1.13 (similar to SYSLOG, now for all)
 - Alternative is to use SEGMENT=xxx on SYSOUT DD stmt.

Planned outage avoidance PDSE

DFSMSdfp recovery enhancements:

- In the past, it was necessary to re-IPL a system or systems to resolve a hang condition, deadlock condition, or storage problem in the PDSE address space. Now, DFSMSdfp will *optionally* use two PDSE address spaces, SMSPDSE and SMSPDSE1. The default will continue to be to only use SMSPDSE:
 - SMSPDSE1 is a restartable address space that provides connections to, and processes requests for, those PDSE data sets that are not part of the global connections associated with SMSPDSE (LNKLST).
 - To create the SMSPDSE1 address space in a sysplex environment, set the IGDSMSxx member as follows:
 - ▶ PDSESHARING(EXTENDED)
 - ▶ PDSE_RESTARTABLE_AS(YES)
 - IPL is recommended for this to be set up initially.



PDSE_SMSPDSE1 HealthCheck



Planned outage avoidance - PDSE

Some customers have encountered problems with hung latches related to PDSE processing - may be accompanied by messages such as:

- IGW038A POSSIBLE PDSE PROBLEM(S). (SMSPDSE or SMSPDSE1)
- RECOMMEND ISSUING V SMS,PDSE,ANALYSIS

The V SMS,PDSE,ANALYSIS command should identify any hung latches (which only have a single-system scope).

One way to release the latches is to do an IPL, but a nicer alternative is to use the V SMS,PDSE,FREELATCH command:

- Also recommend taking a dump of the *MASTER*, SMSPDSE, and SMSPDSE1 address spaces and save a formatted logrec report for later analysis.

VSAM record management trace

Prior to z/OS 1.12, VSAM record management trace could only be enabled by specifying the TRACE parameter on the DD card.

- Meant that trace couldn't be turned on and off without restarting job/started task.

z/OS 1.12 introduces a new Parmlib member (IDAVDTxx) and a new started task (IDAVDT)

- When IDAVDT is started, you can use MODIFY commands to control VSAM record management trace, with no JCL changes required in jobs or started tasks.

Planned outage avoidance System Logger

System Logger


- Ability to update most log stream attributes while the log stream is connected. LOGR CDS format level must be HBB7705 or later. If this is not the case, it is not possible to change log stream attributes if the log stream is still connected
- Offload hang detect:
 - Msgs IXG310I, IXG311I, IXG312E if offload processing is hung.
- Ability to force disconnection or deletion of a log stream:
 - SETLOGR FORCE,DEL|DISC,LSN=log_stream_name command.
 - Can be used to force disconnection of a log stream, or force deletion of the log stream from the LOGR CDS:
 - ▶ Previously forcing a disconnect might require a restart of the System Logger address space - tantamount to an IPL....

Planned outage avoidance Reusable ASIDs

There is a finite supply of address spaces and linkage indexes. Terminating an address space that uses cross-memory can result in that address space becoming non-reuseable.


z/OS 1.9 added a new REUSEASID keyword in DIAGxx member.

- Indicates to enable new support in z/OS to enable some of these address spaces to be reused.
- New option on START command (REUSEASID=YES), indicating that the task should use a "non-reuseable" address space.
 - Requires support in the product being started. Initial support was in LLA, DLF, and VLF



IEA_ASIDS HealthCheck

IEA_LXS HealthCheck



Planned outage avoidance - IEASYMUP

ITSO used to provide a program called SYMUPDTE to dynamically update System Symbols

- SYMUPDTE can be used to update existing symbols without an IPL. It can also add new symbols without an IPL.
- Be sure to pull the documentation to understand the limitations of using this program.

This program is now delivered in SYS1.SAMPLIB as IEASYMUP, however no JCL or documentation is provided:

- Note that IEASYMUP REQUIRES a RACF profile.

The latest documentation for IEASYMUP is in the Planned Outage Avoidance Redbook.

Planned outage avoidance - IEASYMUP

JCL to link IEASYMUP:

```
//KYNEFL JOB (0,0), 'LINK SYM', CLASS=A, MSGCLASS=X, NOTIFY=KYNEF
//S2 EXEC PGM=IEWL,
//  PARM='XREF, NCAL, LIST, RENT, LET, AC=1'
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD UNIT=SYSDA, SPACE=(TRK, (20, 10))
//SYSLMOD DD DSN=KYNEF.IEASYMUP.LOADLIB, DISP=(, CATLG),
//          SPACE=(CYL, (1, 1, 5)), RECFM=U, LRECL=0, BLKSIZE=6144,
//          UNIT=SYSDA
//SYSOBJS DD DSN=SYS1.SAMPLIB, DISP=SHR
//SYSLIN  DD *
           INCLUDE SYSOBJS(IEASYMUP)
           NAME IEASYMUP (R)
```

RACF

DEFINE IEASYMUP.* profile in FACILITY class, grant UPDATE access to sysprog and refresh RACF (SETR RACLIST(FACILITY) REFR

JCL to run IEASYMUP:

```
//KYNEFR JOB (0,0), 'TEST SYMUPDTE', CLASS=A, MSGCLASS=X
//SYMUPDTE EXEC PGM=IEASYMUP, PARM='TESTFK=TEST1'
//STEPLIB DD DSN=KYNEF.IEASYMUP.LOADLIB, DISP=SHR <==== APF library
```

Planned outage avoidance Misc


Tools and Miscellaneous Features

- Image Focus from New Era Software monitors changes between IPLs to ensure parms used to IPL reflect the current configuration and are syntactically correct.
- OS/390 2.10 added ability to rename an ENQed duplicate data set - see STGADMIN.DPDSRN RACF profile.
- Check Parmlib member syntaxes prior to IPL:
 - Use the SPPINST exec (provided in SAMPLIB) to check for valid LOADxx and GRSRNLxx member syntax.
 - Use CEEPRMCC program or CEEPRMCK Clist to check LE members
 - PARMLIB TSO command to check IKJTSOxx members
 - For BPXPRMxx use SETOMVS SYNTAXCHECK command
- Ability to activate USS service dynamically:
 - F OMVS,ACTIVATE=SERVICE
 - But consider carefully the systems management aspects.


Planned outage avoidance Misc

Tools and Miscellaneous Features

- Program to delete EMCS consoles without a sysplex IPL
 - IEARELEC provided in SYS1.SAMPLIB (See APAR OA06857)
 - ▶ //JOBA JOB ...
 - ▶ //sss EXEC PGM=IEARELEC,PARM='CONSNAME(console01)'
 - IPL Times - 100K cons 605 secs, 33K 187 secs, 16 122 secs



CNZ_EMCS_Inactive_Consoles_HealthCheck



- MQ supports ability to update "early code" without an IPL
- DB2 V9 added the ability to update early code without an IPL using the REFRESH DB2 command.

Planned outage avoidance

Summary

- There are **MANY** features in the system and subsystems that can help reduce the number of planned outages.
- **DON'T** be controlled by "we've always done it this way"...
- Dynamic change capability **INCREASES** the need for strong system management.
- Data sharing and dynamic workload balancing contribute more to flexibility than **ANY** of the dynamic change capabilities:
 - It provides the flexibility to select more or less frequent IPLs, taking application impacts out of the picture.
- You need to consider all the risks and advantages of frequent vs infrequent IPLs, and all the capabilities, to determine a frequency that is right for you.

Planned outage avoidance

Further information:

- Paper by Peter Relson: "The Dynamics of z/OS Operating System"
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1876>
- SHARE Session 9703 by Peter Relson
- ITSO RedBook *Planned Outage Avoidance Checklist*, SG24-7328
- ITSO Redbook *System z Mean Time to Recovery Best Practices*, SG24-7816
- ITSO RedBooks - *z/OS 1.x Implementation* - by Paul Rogers, one for each z/OS release.
- *z/OS Summary of Message and Interface Changes*, SA22-7505
- *z/OS Installation and Tuning Reference*, SA22-7592
- *z/OS Introduction and Release Guide*, GA22-7504
- *z/OS Migration*, GA22-7499
- *z/OS System Commands*, SA22-7627

Post-Redbook Outage Avoidance items

The following enhancements were delivered AFTER the Planned Outage Avoidance Redbook was written:

z/OS 1.8 Enhancements

- Ability to move GRS Contention Notification System without an IPL
- Ability to dynamically change size of SMSPDSE1 hiperspace
- New SET DEVSUP=xx command to dynamically activate changes to DEVSUP member
- Dynamically add TCP NJE nodes to JES3

Post-Redbook Outage Avoidance items

z/OS 1.9 Enhancements

- Sample RACF ICHPWX11 (password phrase) exit updated to call System Rexx - allowing you to update the function of the exit without an IPL
- New Healthchecks to monitor for pending shortages of linkage indexes and non-reusable address spaces
- SETPROG LNKLIST command enhanced to make it more flexible
- REUSEASID parm added to DIAGxx
- Ability to restart system rexx address space - AXRPSTRT
- SETOMVS AUTOCVT command lets you dynamically modify the AUTOCVT setting in BPXPRMxx
- New option on START command, to specify that named STC should use a reuseable ASID. Initially for LLA, DLF, and VLF.

Post-Redbook Outage Avoidance items

z/OS 1.10 Enhancements

- Dynamic JES2 exit support
- Ability to change sysplex root data set without sysplex IPL
- Ability to move from GRSRNL=EXCLUDE to full RNLs without a sysplex IPL
- z/OS UNIX RESOLVER address space, TCP/IP address spaces, DFSMSrmm address space, and the TN3270 address spaces now support ASID reuse.
- New SETRRS ARCHIVELOGGING lets you turn RRS archiving on and off without restarting RRS
- Basic HyperSwap lets you swap from primary to secondary DASD without an IPL
- Ability to dynamically add a CP (DYNCPADD in LOADxx) on z10

Post-Redbook Outage Avoidance items

z/OS 1.11 Enhancements

- Ability to point at specific parmlib AXR members when you use the AXRPSTRT proc to restart System REXX
- SETALLOC command changes values in ALLOCxx member without an IPL
- System Status Detection Partitioning Protocol may improve the chances of spin loop recovery completing successfully (thereby avoiding an IPL)
- ALTROOT statement lets you specify alternate sysplex root file system to dynamically switch to in case current sysplex root becomes unavailable
- Ability to specify maximum time that the system is set to be non-dispatchable during a dump - MAXSNDSP
- Enhancements to make dynamic LPA exit (CSVDYLPA) more usable

Post-Redbook Outage Avoidance items

z/OS 1.12 Enhancements

- CRITICALPAGING function for HyperSwap environments
 - Note, however, that an IPL is required to ENABLE this feature
- Ability to specify NOBUFFS action (SMF) at the log stream level
- VSAM CA Reclaim for KSDSs
- Support for non-disruptive CF Dump
- If a broken PDSE is encountered in LNKLIST during IPL, a message is now produced identifying the bad data set and IPL continues without that data set in the Link list.

Post-Redbook Outage Avoidance items

z/OS 1.12 Enhancements

- Enhancements to LLA and PROGxx processing
- CSVLLIX1 and CSVLLIX2 (LLA exits) added to dynamic exits facility
- Extended addressability support for catalogs
- New DEFERTND option to delay making address spaces non-dispatchable during an operator-initiated dump
- Ability to specify a hot-standby Sysplex Distributor
- Able to change number of Common Inet ports without OMVS restart
- HIS detects change in CPU speed without an IPL

Post-Redbook Outage Avoidance items

z/OS 1.13 Enhancements

- Ability to stop a JES2 job at the end of the current step
- Dynamically discontinue use of a JES2 spool volume or increase spool volume size
- Ability to dynamically add spool volumes to JES3
- Ability to change spool-related JES3 parms without an IPL
- DADSM and CVAF support for concurrent service
- Dynamic support for DADSM IGGPRE00 and IGGPOST0 exits

Post z/OS 1.7 Outage Avoidance items

z/OS 1.13 Enhancements

- New FORCE option of CMDS command
- New UNALLOC parameter for the SPIN keyword on the DD statement, to allow you to specify that output data set should be spun off without stopping and starting address space
- DEVMAN added to CATALOG, LLA, VLF, RESOLVER, TCP/IP, DFSMSrmm, and TN3270 to mark address spaces as reusable

Planned Outage Avoidance

Other Enhancements?

If you know of any other enhancements in this area that I have missed, please come and talk to me, or send me an email (kyne@us.ibm.com)

Also, I am trying to compile a list of items that still require an IPL, so if you would like, please send me that list as well.

Shameless advertising

Come to Poughkeepsie to take part in a project with other subject matter experts from all over the world to write a Redbook.....


- IBM covers all travel expenses, hotel, meal allowance, car, etc...
- Your mission is to learn as much as you can about the latest and greatest IBM technology and document your experiences
- Gain fame and fortune (well, at least, you will get your name on the front cover of a Redbook)

If you think you might be interested, keep an eye on <http://www.redbooks.ibm.com/residents.nsf/ResIndex/>

or sign up for automatic notification at

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

Questions?



Thanks!!



**And please hand in your evaluations (Session 9730)
so we can improve this material**



International Technical Support Organization

ibm.com/redbooks

Additional slides



© 2010 IBM Corporation. All rights reserved.

Planned outage avoidance RACF

RACF Started Class Names Table

- How many people still use the old Started Class Names Table to associate Started Tasks with RACF userids?
 - This requires an IPL every time you want to add a new STC
- RACF 1.9 (1989!) introduced the RACF STARTED class that allows you to assign STCs to RACF userids using the panels or RACF commands - no module updates (and no IPLs!) required

Planned outage avoidance RACF

Security Server Dynamic CDT support

- Starting with z/OS V1R6 the Dynamic Class Descriptor Table (CDT) provides the means to add, change, and delete installation-defined classes in the CDT without IPLing the system.
- New RACF class name, CDT, is used to hold the definitions of the installation-defined classes.
- You can use RDEFINE and RALTER commands to define classes.
- Then use command SETROPTS RACLIST(CDT) REFRESH.

Planned outage avoidance RACF

Security Server Removal of the Router Table

- Back in z/OS V1R6, the IBM-supplied portion of the RACF router table (ICHRFROX) was removed:
 - This eliminates the need to provide an entry in the installation-defined router table (ICHRFR01) for every installation-defined class.
 - Most installation-defined classes will not require any change to ICHRFRO1. The most likely candidate for this is DSNR (DB2).
- No longer need to IPL when an installation-defined class is being added.

Planned outage avoidance RACF

Security Server Dynamic Template enhancements

- Templates now have a level indicator which can be used to compare two sets of templates and determine which is the latest. Prevents you from installing a downlevel set of templates onto a RACF database.
- During initialization at IPL, RACF will determine whether the database has the right level of templates and if not, RACF will ignore the templates in the database and use those from IRRTEMP2 automatically.
- When a PTF that changes the templates is applied to a live system, you can run IRRMIN00 and have RACF recognize the new templates without an IPL.
- Prevents complete reinitialization of a RACF database, if that database is "live" on the system where IRRMIN00 is run.

Planned outage avoidance RACF

Security Server changes that still require a sysplex-wide IPL:

- Updates to the RACF Data Set Names Table (DSNT) require a sysplex-wide IPL
- Updates to the RACF Range Table require a sysplex-wide IPL

Therefore, monitor database usage and plan on adding a database if necessary as far in advance as possible (to take advantage of any planned sysplex IPLs)

Planned outage avoidance Commands

Sometimes commands can run for a long time, locking out other commands - **VARY ON** or **VARY OFF**, for example.

To address this situation, the **CMDS** command was introduced.

CMDS command lets you:

- **DISPLAY** commands that are running and waiting to execute.
- **SHOW** detailed information about specific commands.
- **REMOVE** commands awaiting execution.
- **ABORT** running commands.

See messages **IEE062I** and **IEE063I** for sample output.

Planned outage avoidance Clock change

Time Change Considerations

- Most IBM z/OS products now support the summer/winter time change. However, applications still need to be reviewed and there are still some 'gotchas':
 - RMF will have missing data unless kept in separate data set
 - Tivoli Omegamon products recycle required
 - Remove use of EDCLLOCL usermod in C/C++ and Language Environment.
- TWS clarification of support in 8.2 for time change and sysplex timer
 - see DOC APAR PK06007 (8.1 requires recycle)

Planned outage avoidance zFS

Migration to zFS

- As of z/OS 1.7, HFS was functionally stabilized - all future enhancements will be in zFS - therefore you should start planning on migrating all HFS files to zFS
- IBM provides a tool under TSO called BPXWH2Z to help you migrate file systems:
 - However, the file system should not be R/W when you convert it.
 - ▶ Use FSINUSE tool to see if anyone is using it.
 - Userid that runs the tool must be UID(0).
 - Currently must be initiated from ISPF panels.
 - Converting the version root HFSs can be done using rolling IPLs.
 - Converting the system or sysplex root no longer requires a sysplex-wide IPL (as of z/OS 1.10).

Planned outage avoidance Unicode

Dynamic UNICODE updates

- Parmlib member CUNINIxx contains statement used to control the conversion environment:
 - Pointed to by UNI=xx in IEASYSxx.
- Starting with z/OS 1.7, you can update the conversion environment definitions dynamically using SETUNI command.

Planned outage avoidance Linkage Indexes (LXs)

Non-reuseable linkage indexes

- Prior to z/OS 1.6, system LXs were not reusable. Every time a user of a system LX goes away, the LX becomes dormant. Eventually you run out of LXs and need to re-IPL.
- In z/OS 1.6, the number of LXs was increased from 2048 to 32K and LXs can be reusable (requires exploitation by the products)
 - NSYSLX max value was 512 prior to 1.6, 512 (12-bit), 8192 (24-bit) after 1.6
- In z/OS 1.7, RRS added support for the use of reusable LXs by the products using its services.
- For more information, see <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FLASH10273>

Planned outage avoidance Linkage Indexes (LXs)

Non-reuseable linkage indexes

- z/OS 1.3 introduced monitoring - msg issued at 85% usage:

```
-IEA059E ASID SHORTAGE HAS BEEN DETECTED
-IEA060I ASID SHORTAGE HAS BEEN RELIEVED
-IEA061E REPLACEMENT ASID SHORTAGE HAS BEEN DETECTED
-IEA062I REPLACEMENT ASID SHORTAGE HAS BEEN RELIEVED
-IEA063E SYSTEM LX SHORTAGE HAS BEEN DETECTED
-IEA065E NON-SYSTEM LX SHORTAGE HAS BEEN DETECTED
-IEA066I NON-SYSTEM LX SHORTAGE HAS BEEN RELIEVED
-IEA070E SYSTEM BIG LX SHORTAGE HAS BEEN DETECTED
-IEA071I SYSTEM BIG LX SHORTAGE HAS BEEN RELIEVED
-IEA072E NON-SYSTEM BIG LX SHORTAGE HAS BEEN DETECTED
-IEA073I NON-SYSTEM BIG LX SHORTAGE HAS BEEN RELIEVED
```

- Recommend adding code to your automation package to generate alerts if any of the above are issued.

Planned outage avoidance GRS

GRS availability features:

- You can change GRS RNLs dynamically using `SET GRNRNL=xx` as long as all members of the GRS complex (ring or star) are members of the sysplex (the RNLs are stored in sysplex CDS).
- Use `D GRS,ANALYZE` command to quickly identify blocked or holding users.
- `SYNCHRES` to make `RESERVES` synchronous - can be turned on and off dynamically (`SETGRS` command), system by system.
- Ability to change Contention Notification System by command (`SETGRS CNS=system_name`) - previously you had to shutdown the current CNS to get the function to move to another system.

Planned outage avoidance XES

Enhancements in hang detection in XES

- Previously, a structure connector that failed to reply to a rebuild request was difficult to detect and could result in unnecessary IPLs.
- XES added hung member detection, which identifies the guilty member with message IXL041E.
- Again, use automation to highlight these messages:
 - z/OS 1.12 adds function to SFM related to this condition.

Planned outage avoidance VTAM

VTAM Generic Resources

- When an application registers as a generic resource, information about the GR name is stored in various places in VTAM (and in the ISTGENERIC structure).
- If you want to reuse that name for something else, all that information must be deleted.
- Prior to 1.7, this meant defining a new GR structure, stopping ALL VTAMs, and bringing them up with the new structure - similar impact to a sysplex IPL.
- In 1.7, there is a new VTAM command to delete GR information dynamically:
 - F NET,GR,GRNAME=neta.grappl,OPTION=DELETE
- For more information, refer to section 6.1.3.7 "Removing a Generic Resource", in *SNA Network Implementation Guide*.

Planned outage avoidance Health Checker

z/OS HealthChecker:

- VERY popular with customers - may warn you of problems before they become critical, thereby avoiding an IPL.
- Everyone should run every check and either address any "issues" that are raised, or adjust the HealthChecker parms so it understands this situation is normal in your shop. Then run it continually.
 - Note that the message numbers reflect the component that owns the check (IXCHxxx) rather than HealthChecker (HZS)

Planned outage avoidance Hardware

Hardware related:

- HyperSwap - if you have the ability to use HyperSwap, make sure it is enabled.
- Make sure you specify large enough max subchannels in HCD to allow for numerous dynamic changes:
 - Only devices that do not support dynamic reconfiguration are those that use old definitions that pre-date dynamic reconfig support - 3274, for example.
- Monitor for and apply New Function service for new devices in a timely manner:
 - There is little point in having the ability to dynamically add new devices if you need an IPL specifically to pick up service related to the new device.

Planned outage avoidance Hardware

Hardware related:

- IF you have a need to move storage between LPARs, specify RSU (in terms of MB or GB) in IEASYSxx, and RESERVED STORAGE value for target LPAR:
 - Reconfigurable element size depends on CPC generation.
- Define spare (placeholder) LPARs.
- Define all LPARs with RESERVED CPs (and zIIPs/zAAPs if approp):
 - Recommend INITIAL+RESERVED=max number of CPs possible on that CPC model.
- When ordering the CPC, go through the plan-ahead process with your IBMer - to allow future non-disruptive growth:
 - When ordering upgrades, ensure that you can get from the current configuration to the target one without a sysplex IPL.
- Remember that CFLevel upgrades do NOT require a POR.

Planned outage avoidance

What does still need an IPL ?

- System Symbols - can use IEASYMUP, but be careful!
- Parmlib members - BLSCECT, BLSCUSER, CNIDTRxx, CONFIGxx, DEVSUP, EPHWPOO, IEAAPPOO, IEAFIX00, IEAPAKxx, IEASYMxx, IEASYSxx, IOEPRMxx, LOADxx, MSTJCLxx, and NUCLSTxx, still require an IPL to modify.

Planned outage avoidance JES2

JES2 SPOOL partitioning and affinities

- To minimize the impact of loss of a SPOOL volume, there are two things you can do:
 - Limit jobs or job classes to a subset of the spool volumes using the FENCE statement in JES2PARM.
 - Create affinities between a SPOOL volume and a system or set of systems. It is NOT possible to do this in JES2PARM - you must use the \$TSPOOL command to define the associations:
 - ▶ `$TSPOOL(jessp1),SYSAFF=(SYSA,SYSB)`

Planned outage avoidance

One general point before we start:

- *The current (running) sysres set, and the associated USS Version HFS/zFS, should be considered READ ONLY and should ONLY ever be updated in exceptional circumstances*

Some best practices:

- IBM-provided libraries (LINKLIB, SVCLIB, etc) should only ever be updated with SMP
 - PLEASE don't try to manually copy updated modules - it ALWAYS causes problems
- Every sysres set that you can IPL should have an SMP target zone associated with it
- The dynamic change capabilities are designed for *additions* to your environment (new HW, new SW) - they are NOT designed for mass apply of service

Planned outage avoidance

What does still need an IPL ?

- New PLPA and COMMON page data sets
- RACF Data set name table, RACF range table changes (SYSPLEX IPL)
- Adding JES3 Spool volumes - FIXED in z/OS 1.13
- Adding MCS and SMCS consoles
- Backouts for GRS STAR, BRLM and PDSE Sharing (SYSPLEX)
- Changing MAXUSER, RSVSTRT, RSVNONR, and NSYSLX values
- Service:
 - Mass PTF apply
 - Updates to NUCLEUS or LPA (unless specifically supported by vendor)
- Changes to the USS SWA() and SYSPLEX() parameters

Planned outage avoidance

Dynamic changes

- Good, because they help you avoid an IPL
- BUT, they require foolproof system management and processes to avoid regressing them in the future

Should you use SET aaa=xx, or SETaaa parm.....?

- Using SET aaa=xx has the benefit of ensuring that you update the Parm member AND it syntax checks the change
 - Also, not every Parmlib command that supports a SETABC command also supports SET ABC=
- But... what if many people changed that member - do you activate changes that you are not ready for yet?
- There is no "right" answer - it varies from installation to installation

Planned outage avoidance

What are the considerations for fewer IPLs?

- When you DO have one, many more changes are squeezed into one outage
 - This increases risk - lot of changes, little time = mistakes
- If you changed something dynamically and forgot to reflect the change back into Parmlib, there is less chance you will remember the change the longer it is until the next IPL
- If you have a way of keeping the applications available (data sharing and workload balancing), there should be less pressure to reduce the number or duration of planned outages

Planned outage avoidance Catalog

DFSMSdfp - catalog space monitoring

- Prior to z/OS 1.5, there was no warning when a catalog is about to run out of space.
- Now, z/OS issues a message when a specified percentage of maximum extents is reached:
 - IEC361I CATALOG catalogname (comptype) HAS REACHED xxx% OF THE MAXIMUM
 - Issued once per catalog per extent and reissued for each subsequent catalog extent
- Invoked by issuing F CATALOG,NOTIFYEXTENT(xxx) where xxx default is 80. 0 or 100 will disable the function
- F CATALOG,ALLOCATED will display the percentage of allocated extents for each catalog in the list in message IEC348I (rightmost column under %)
- z/OS 1.13 increases limit on catalog size

This is NOT Poughkeepsie.....



Neither is this.....





Welcome
Reception



ME



YOU