



International Technical Support Organization

z/OS Planned Outage ~~Avoidance~~ Control SHARE Atlanta 2012 Session 10650

www.ibm.com/redbooks

Frank Kyne
(kyne@us.ibm.com)



Redbooks[®]

IBM ITSO - International Technical Support Organization

© 2012 IBM Corporation. All rights reserved.

Intro

Who am I?

Objective of this session.

- This is not about eliminating planned IPLs. It is about you controlling them, rather than them controlling you.

PLEASE ask questions as I go along.

If you have suggestions or additional information, please come and chat after the session or email me or fill in the comment area on the evaluation forms.

Acknowledgements

The inspiration for this presentation is an IBM Redbook, *Planned Outage Avoidance Checklist, SG24-7328*, written by:

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

With invaluable help from Peter Relson, Marna Walle, and Robert Hering we updated this presentation with enhancements that have been introduced since that Redbook.

Acknowledgements

This is an updated version of a presentation that Marna Walle did at the Fall 2011 SHARE in Orlando.

We removed information about features that were available prior to z/OS 1.8 (and therefore are covered in the referenced Redbook) and focus on the enhancements since then.

Related sessions

It is a little late now, but did you attend the session on Parmlib management yesterday afternoon?

Excellent session (session number 10648) with some very smart (and hard-learned) lessons - download the presentations (there are 2 of them) from SHARE Web site and speak to the presenters if you have questions or ideas.

- One item that is especially relevant is Mary Anne's use of Symbolic Alias Facility (combined with Peter's IEASYMUP) to let you easily switch back and forth between different releases of products, potentially without an IPL.

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 new product for me?"

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

Addressing this challenge

1) Use Parallel Sysplex data sharing and dynamic workload routing to let you IPL when you need to, for as long as is required:

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

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

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

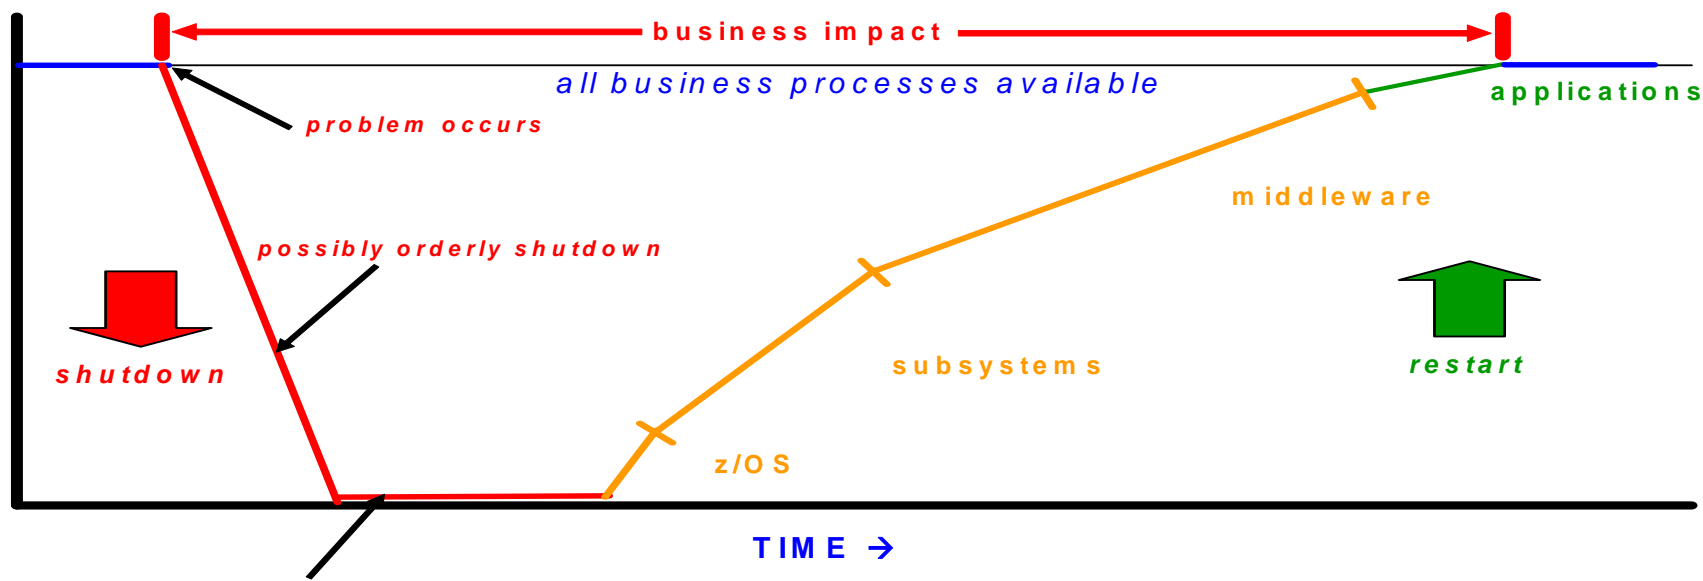
– Exploiting data sharing and dynamic workload routing lets you decouple application availability from the availability of a given system. So you can give your users the availability *they* need, and still have the system outages *you* need.

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 working to reduce the amount of time it takes to stop and restart z/OS and its major subsystems - See Redbook SG24-7816 for more info.



For Unplanned - gather diagnostic information

For Planned - implement planned changes

Addressing this challenge

3) Let you change more things without performing an IPL.

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, but that are equally important.

Addressing this challenge

Less user impact

More changes per IPL

Higher risk of mistakes

Higher risk of a change being forgotten

Higher risk of problem rediscovery if service level not kept current


Fewer changes/less risk

Reduced chance of regression

High user impact

Bad user perception

Higher risk of IPL-induced problems



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.
- 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 5 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

Answer:

- At one time, nearly all of these would have required an IPL.
- Now, only one of them requires an IPL to change it (RSVNONR - reserving slots for non-reusable address spaces).
 - However IBM have been making changes to reduce the number of situations where address spaces go non-reusable, so the same RSVNONR value should let you run for longer between IPLs.

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

Answer: IT DEPENDS (of course)

Some products support dynamic updates to LPA-resident modules, some don't. The only way to know for sure is to look in the PTF text - if it supports dynamic update, it should say so in the text.

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

Answer:

- Prior to z/OS 1.10, a sysplex IPL was required to make this change.
- As of 1.10, you can make this change without stopping and restarting **ANYTHING** - no sysplex IPL, no system IPL, no USS shutdown, no need to shutdown any address spaces.
 - For more information, refer to section 18.1 in z/OS 1.10 Implementation Redbook (SG24-7605)

Planned outage avoidance

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

Enhancement: In z/OS 1.13 (and rolled back to z/OS 1.10 with OA29367), on the z/OS end of the link, you can specify an AID of *. So, as long as you install the adapters on the CF end in advance, you can now subsequently add InfiniBand adapters to the z/OS CPC, do a dynamic reconfig on just that CPC, and bring the new links online and start using them without PORing the standalone CF or restarting the CF LPAR.

Planned outage avoidance

Let's look in a little more detail at **SOME** of the Planned Outage avoidance enhancements in z/OS 1.8 to z/OS 1.13....

Planned outage avoidance - GRS

One system in a sysplex performs the role of Contention Notification System (CNS) when in GRS Star mode.

- That system gathers contention information from all systems in the sysplex, and then transmits it to all systems using ENF 51 signals.

The first system IPLed in the sysplex will be the CNS

- If that system goes down, the other systems race to take over that role.

Prior to z/OS 1.8, the only way to remove the CNS role from a system was to IPL that system.

z/OS 1.8 added the SETGRS CNS=sysname command to let you move the role to any system at any time, with no IPL.

- Might consider adding that command to your automation at IPL time to ensure that that role always resides on chosen system
- Ensure CNS has sufficient capacity to handle workload

Planned outage avoidance - GRS

z/OS 1.8 also added the ability to control (via GRSCNFxx member) the system-wide maximum number of concurrent ENQs by Authorized (ENQMAXA) and Unauthorized (ENQMAXU) address spaces.

It also provided the ability to dynamically change these values using the SETGRS ENQMAXA and SETGRS ENQMAXU commands (thereby eliminating a potential IPL to change those values).

Planned outage avoidance - GRS

GRS is always present in a z/OS system. However, for those customers using another serialization product, you would run GRS in `GRSRNL=EXCLUDE` mode.

To migrate from that mode to full GRS mode used to require a sysplex IPL - z/OS 1.10 allows you to keep apps up and running on one system during the migration (with some restrictions).

- Use of `SET GRSRNL=xx` results in:

ISG880D WARNING: GRSRNL=EXCLUDE IS IN USE. REPLYING FORCE WILL RESULT IN THE USE OF SPECIFIED RNLS. REPLY C TO CANCEL

- See z/OS 1.10 Implementation Redbook for restrictions

Note that going back to `GRSRNL=EXCLUDE` mode would require a sysplex IPL.

Planned outage avoidance - GRS

GRS provides commands to display information about ENQs that it manages - `D GRS,C`, `D GRS,ANALYZE`.

However, GRS also manages latches and these can also experience contention.


If you experience latch contention, but are not aware of that, the end result might be IPLing a system in an attempt to clear up the problem.

z/OS 1.11 added the ability to get latch information from GRS, by using the `D GRS,ANALYZE,LATCH` command. With that information, you can address the offending address space, instead of IPLing the system to make the problem go away.

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 *optionally* uses two PDSE address spaces, SMSPDSE and SMSPDSE1. The default continues 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 required in order to turn SMSPDSE1 on initially.



PDSE_SMSPDSE1 HealthCheck



Planned outage avoidance - PDSE

Prior to z/OS 1.12, if the linklist contained a corrupted PDSE, the system would wait state when it attempted to read that data set during IPL.

- The wait state would be entered without displaying the name of the offending PDSE.
- This not only failed the IPL, you ALSO had to analyze a SAD to identify the broken PDSE, extending the length of the outage.

Starting with z/OS 1.12, the system will issue a message indicating that it had a problem ("IGW704A One or more corrupted PDSE were found in LINKLIST please check the SYSLOG for more details"), and the IPL then continues without that data set in the link list.

- This gives you the opportunity to restore the data set or remove it from link list and either add it to link list dynamically or re-IPL after the problem has been addressed.

Planned outage avoidance - PDSE

The `V SMS,PDSE|PDSE1,ANALYSIS` command can help identify damaged PDSEs.

The `D SMS,PDSE|PDSE1,CONNECTIONS,DSN(dsname)` command (delivered with z/OS 1.13) helps you identify who is using the named data set.

The `V SMS,PDSE|PDSE1,REFRESH,DSN(dsname)` command (also delivered in z/OS 1.13) discards directory pages for the named data set from the PDSE|PDSE1 caches.

Planned outage avoidance - PDSE

z/OS 1.13 also contains a new batch utility - IEBPDSE - that helps you detect damage in the named PDSE(s):

```
//STEPCHK2 EXEC PGM=IEBPDSE,PARM='DUMP'  
//SYSLIB DD DSN=IBMUSER.SIMPLE.V2.PDSE,DISP=OLD  
// DD DSN=IBMUSER.SIMPLE.V3.PDSE,DISP=OLD
```

If you wish, you could use this utility as part of your system shutdown procedure to ensure that all PDSEs in the link list are error-free.

- If errors are detected, program ends with non-zero return code. This could be built into your automation to let the operator decide if the IPL could be postponed, or the procedures changed so that the damaged PDSE is left out of link list until it can be repaired.
- Recommend including the PARM="DUMP" parameter to get dump that may help diagnose the problem.

Planned outage avoidance - zFS

Per z/OS 1.7 announcement:

- The Hierarchical File System (HFS) functionality has been stabilized
- All requirements for UNIX file services are expected to be addressed in the context of zFS only

Prior to z/OS 1.10, it was possible to migrate HFS files to zFS dynamically (using BPXWH2Z), however a sysplex IPL was required to move the sysplex root to zFS.

Planned outage avoidance - zFS

Starting with z/OS 1.10, this can now be done dynamically:

- Ensure current root is mounted read only
- Create new sysplex root zFS file
- Mount new sysplex root so we can copy old sysplex root into it
- `/samples/copytree -a / /tmp/temp_mp` to copy contents
- Unmount new root
- `F OMVS,NEWROOT=new_sysplex_root_name,COND=YES`
 - This will switch to new sysplex root and mount it R/O

Planned outage avoidance - zFS

D OMVS,F

BPXO045I 13.55.29 DISPLAY OMVS 989

OMVS 0010 ACTIVE OMVS=(00,FS)

TYPENAME	DEVICE	-----STATUS-----	MODE	MOUNTED	LATCHES
----------	--------	------------------	------	---------	---------

HFS	1	ACTIVE	READ	03/08/2012	L=14
NAME=OMVS.SYSPLEX.ROOT				13.52.19	Q=14

PATH=/
OWNER=#@\$3 AUTOMOVE=Y CLIENT=N

F OMVS,NEWROOT=OMVS.TRAINER.SYSPLEX.ROOT,COND=YES

IEF196I IEF285I OMVS.SYSPLEX.ROOT KEPT

IEF196I IEF285I VOL SER NOS= #@\$#M1.

BPXF246I THE SYSPLEX ROOT FILE SYSTEM MIGRATION PROCESSING 009
COMPLETED SUCCESSFULLY.

IEF196I IEF285I OMVS.SYSPLEX.ROOT KEPT

IEF196I IEF285I VOL SER NOS= #@\$#M1.

D OMVS,F

BPXO045I 14.06.55 DISPLAY OMVS 013

OMVS 0010 ACTIVE OMVS=(00,FS)

TYPENAME	DEVICE	-----STATUS-----	MODE	MOUNTED	LATCHES
----------	--------	------------------	------	---------	---------

ZFS	1	ACTIVE	READ	03/08/2012	L=53
NAME=OMVS.TRAINER.SYSPLEX.ROOT				14.06.22	Q=53

PATH=/
OWNER=#@\$3 AUTOMOVE=Y CLIENT=N

Planned outage avoidance - zFS

This process can also be used to:

- Move from zFS back to HFS sysplex root
- Move from one zFS sysplex root to another
- Move from on HFS sysplex root to another

All systems in the sysplex must be running z/OS 1.10 or later to be able to use this process.

For more information, refer to section 3.7 in z/OS Distributed File Service zSeries File System Implementation z/OS V1R11, SG24-6580 Redbook

Planned outage avoidance - zFS

While the sysplex root does not (should not) contain much, it is used by the entire sysplex and constitutes a single point of failure.

z/OS 1.11 delivered the ability to define an alternate sysplex root data set:

- **ALTROOT** parameter in BPXPRMxx
- All systems in the sysplex must be z/OS 1.11 or later
- Works for both HFS and zFS
- Can switch to alternate automatically.
- If you update current sysplex root, **YOU** are responsible to copy those changes over to the alternate.

For detailed implementation information and restrictions, refer to section 16.8 in z/OS 1.11 Implementation Redbook, **SG24-7729**.

Planned outage avoidance - Moving TCDB

A Tape Configuration Database is a special type of ICF catalog, used to contain information about tapes in a tape library.

The TCDB name is fixed:

- You can use a HLQ other than SYS1 if you specify the HLQ in the SYSCAT parm in LOADxx member
- TCDB name is not specified anywhere else (SMS or anything)

The TCDB will typically be shared by all members of the sysplex.

So, how do you move it without taking the sysplex down?

Planned outage avoidance - Moving TCDB

- 1) Stop RMM and OAM on all systems that share the TCDB
- 2) Quiesce all tape libraries that use the TCDB
 - `V SMS,LIBRARY(libname), ALL, OFFLINE`
- 3) Verify integrity and structure of the Volcat
- 4) Export volcat to a temporary file
- 5) Disconnect from Master catalog
- 6) Import to new volume
- 7) Import Connect to Master catalog (depending on number of master catalogs the volcat is cataloged in)
- 8) Bring libraries back online
- 9) Restart RMM and OAM

See <http://www.ibm.com/support/docview.wss?uid=isg3S1001551> for details

Planned outage avoidance - DEVSUP member

What exactly does the DEVSUP member control?

```
COMPACT={ [YES|NO] }
TAPEAUTHDSN = { YES|NO }
TAPEAUTHF1 = { YES|NO }
TAPEAUTHRC4 = { ALLOW|FAIL }
TAPEAUTHRC8 = { FAIL|WARN }
VOLNSNS={ [YES|NO] }
MEDIAn=xxxx
ERROR=xxxx
PRIVATE=xxxx
ALVERSION={ 3 | 4 | FORCE3 | FORCE4 }
TAPEBLKSZLIM={ nnnnn | nnnnnK | nnnnnM | nG }
COPYSDB={ YES | SMALL | LARGE | INPUT | NO }
ENFORCE_DC_MEDIA={ ALLMEDIATY | MEDIA5PLUS }
MTL_NO_DC_WORM_OK
{ENABLE | DISABLE} feature
NON_VSAM_XTIOT={ YES|NO }
TAPE_MULTI_VOLUME_ANOMALY={ALLOW | FAIL }
OCE_ABEND_DESCRIP={ YES|NO }
DDRSIZELIM={ xxxx | xxxxM | 1000M }
EXPIRATION_MESSAGE={NEVER | ALWAYS }
```

Planned outage avoidance - DEVSUP member

Starting with z/OS 1.8, you can activate changes to the DEVSUP member by using the SET DEVSUP=xx command.

This should eliminate the need to IPL to make those changes effective.

- Note that you can also use the DEVSERV QL command to change the setting of *some* of the keywords - BUT REMEMBER TO UPDATE THE DEVSUPxx MEMBER, OR THE CHANGES WILL BE LOST AT THE NEXT IPL.

Planned outage avoidance - RACF exits

The RACF password exits (ICHPWX01 and ICHPWX11) must be installed in LPA.

- This means that a change to the exit requires an IPL with CLPA to activate it.

In z/OS 1.9, the sample ICHPWX11 exit was changed to call a System Rexx exec.

- Now, all you have to do is update the Rexx exec and save it.

The sample ICHPWX01 exit that is shipped with z/OS does NOT call System Rexx. However... a sample IS provided at <http://www.ibm.com/systems/z/os/zos/features/racf/downloads/rexpwexit.html>

Just like the ICHPWX11 sample, this lets you change the function in the exit without having to do an IPL.

Planned outage avoidance - System Rexx

And on the subject of System Rexx....

z/OS 1.9 added the ability to restart the System Rexx address space without doing an IPL.

- First you have to stop the current System Rexx address space.
 - It is non-cancellable, so you have to FORCE it.
 - Stopping the AXR address space will abend any running System Rexx execs:
 - ▶ Issue `F AXR,SYSREXX STATUS,DETAIL` to get a list of running System Rexx execs
- Then you restart System Rexx address space using `S AXRPSTRT`

Planned outage avoidance - System Rexx

z/OS 1.11 enhanced System Rexx by letting you concatenate other libraries to `SYS1.SAXREXEC`.

- Means that you no longer need to place your execs in an IBM-delivered data set.
- Concatenation (and other parameters) are defined in AXRxx member in Parmlib.
 - IEASYSxx contains AXR parameter to specify which AXR member(s) you want to use.
- The `S AXRPSTRT` proc supports specification of one or more AXR members:
 - `S AXRPSTRT,AXR=(00,01)`

Planned outage avoidance - ALLOCxx member

What exactly does the ALLOC member control?

```

SPACE PRIMARY(nnnnnnn) SECONDARY(nn) DIRECTORY(n)
    {TRK } MEASURE {CYL } {AVEBLK BLKLNTH(nnnn) {ROUND } }{ {NOROUND} } {CONTIG}
PRIM_ORG {MXIG } {ALX } {RLSE } {NORLSE}
UNIT NAME(unit-name) UNITAFF(unit-name) REDIRECTED_TAPE {TAPE}{DASD}
TIOT SIZE(nn)
SDSN_WAIT WAITALLOC {NO } {YES }
VOLUME_ENQ POLICY {WTOR } {CANCEL} {WAIT} {WTOR }
VOLUME_MNT POLICY {CANCEL} {WTOR }
SPEC_WAIT POLICY {WAITHOLD} {WAITNOH} {CANCEL } MAXNWAIT(nnn) POLICYNW {WTOR } {CANCEL} {WTOR }
ALLC_OFFLN POLICY {WAITHOLD} {WAITNOH} {CANCEL } MAXNWAIT(nnn) POLICYNW {WTOR } {CANCEL}
ALLOCxx
CATLG_ERR FAILJOB {YES} {NO } ERRORMSG {YES} {NO }
2DGT_EXPDT POLICY {ALLOW}{WARN} {FAIL}
VERIFY_VOL POLICY {YES} {NO}
SYSTEM IEFBR14_DELMIGDS(LEGACY|NORECALL)
    TAPELIB_PREF(EQUAL|BYDEVICES)
    REMIND_INTV(xxx)
    TEMPDSFORMAT(UNIQUE|INCLUDELABEL)
    VERIFY_UNCAT(FAIL|TRACK|MSGTRACK|LOGTRACK)
    MEMDSENQMGMT(ENABLE|DISABLE)

```

Planned outage avoidance - ALLOCxx member

Prior to z/OS 1.11, changes to any of the parameters in the ALLOCxx could only be activated via an IPL.

z/OS added two new commands:

- D ALLOC,OPTIONS (to display the settings)
- SETALLOC parm=value (to change individual parameters)

These let you implement changes to the parameters controlled by the ALLOCxx member **WITHOUT** an IPL:

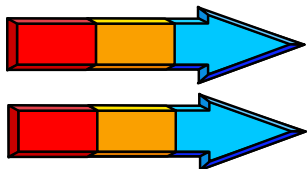
- Note that there is **NO** SET ALLOC=xx command. So you cannot update ALLOCxx and dynamically activate the new member.
- **REMEMBER** to update the ALLOCxx member after activating the change using the SETALLOC command. If you forget, next IPL will regress your change.

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.

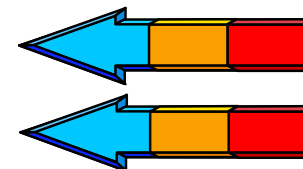
z/OS 1.9 added a new REUSASID keyword in DIAGxx member.

- Indicates to enable new support in z/OS to let some of these address spaces to be reused.
- New option on START command (REUSASID=YES), indicating that the task should use a "reuseable" address space.
 - Requires support in the product being started. Now supported by LLA, DEVMAN, CATALOG, DLF, RESOLVER, TCPIP, DFSMSrmm, TN3270, and VLF



IEA_ASIDS HealthCheck

IEA_LXS HealthCheck



Planned outage avoidance - RRS log stream

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

- For some systems, restarting RRS is as disruptive as an IPL.

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

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

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 SPIN=UNALLOC JCL keyword in z/OS 1.13 (similar to JESLOG, now for all)
 - ▶ Lets you spin off SYSOUT output every xxx lines, or every xx minutes
 - Alternative is to use SEGMENT=xxx on SYSOUT DD stmt.

Planned Outage Avoidance - VSAM record mgmt 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.
- See z/OS 1.12 Implementation Redbook (section 7.6.3) for more information about how to use this.

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 your LOGR CDS format level is lower than this, 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 - make sure you define these messages to your automation.
- 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....
 - ▶ And you can use IXCMIAPU to rename a log stream

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 - Hardware

CF LPAR changes:

- You can dynamically add engines to a CF LPAR, up to the RESERVED number from LPAR profile.
- You can configure on CHPIDs that have been added using z/OS (or z/VM) dynamic I/O reconfiguration.
- Most CFCC patches can be activated without restarting the CF LPAR.
- However, adding or removing storage from a CF LPAR DOES require a deactivate/activate

Did you know that you can add MORE engines than the LPAR RESERVED number? Delivered on z10 and z/OS 1.10, enabled by default.

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 routing 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/atmastr.nsf/WebIndex/PRS1876>
- SHARE Session 9703 "PROGxx and LLA Enhancements z/OS 1.12" 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* - one for each 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 SMSPDSE and SMSPDSE1 hiperspaces
- 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 LNKLST command enhanced to make it more flexible
- REUSASID 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 come across any, please send me that list as well.

Questions?



Thanks!!



And please hand in your evaluations (Session 10650)