


z/OS V1R13

Improvements in R13



This presentation covers the improvements specific to Release 13 of the z/OS® operating system. The developers provided internal education on the features that they authored; the resulting material has been condensed to PDF files and organized here for your reference.

Improving usability and skills (1 of 3)

- JCL Improvements with JES2
 - Stop journaled jobs on step boundaries
 - Forces job out of execution when the current step ends and re-queued for execution and held if requested
 - \$Ejxxxx,STEP[,HOLD]
 - Job-level return codes
 - JOBRC=HIGHEST(MAXRC), LAST(LASTRC), STEPRC(STEP,STEPNAME)
 - \$T JOBCLASS(x) is enhanced with JOBRC operand
 - Support for instream data sets in PROCs
 - //ddname DD *
 - SPIN= DD JCL (and dynamic allocation) support for spin interval specification similar to that on JESLOG
 - SPIN=(UNALLOC, interval|time|size)
- Remaining SDSF Sysplex functions no longer to require WebSphere® MQ:
 - WLM enclaves (ENC)
 - z/OS UNIX processes (PS)
 - Health checks (CK)
 - Resource monitor (RM) (JES2 only)

There have been a lot of improvements in the area of JCL with JES2. You can stop a journaled job on a step boundary. What that means is it will stop a job, force it out of execution, when the current step or the defined step ends.

It will be re-queued for execution, and it will be held, if requested. For example, the request is put in with \$EJxxxx (job number, step, optionally Hold).

For job level return codes, there is a new parameter called JOBRC, which controls how the job completion code is stepped. By default, it is set to the highest return code of any step; if the execution fails with an ABEND, the completion code is set to the last ABEND code. However, this parameter can be used to request that the job completion code be set to the return code of the last step.

There is also now support for in-stream data sets in procedures - you can now specify //ddname DD* and provide in-stream data.

SPIN=DD JCL and dynamic allocation support for spin interval values now is very similar to JESLOG, so you will specify SPIN=UNALLOC, then either the interval, time, or size when that needs to be spun off.

Remaining SDSF sysplex functions no longer require WebSphere MQ. In a sysplex if you wanted to have a sysplex view of (for example) WLM enclaves under SDSF, you needed to have WebSphere MQ implemented or you could not see the WLM enclaves from every system on a single panel. That restriction has been removed. Now SDSF collects that information using an XCF interface. So WLM enclaves, z/OS UNIX process, health checks resource monitoring for JES2 – that can all be done now, without the use of WebSphere MQ.

Improving usability and skills (2 of 3)

- New Catalog PARMLIB Member IGGCATxx
 - CATALOG=(xx,yy, ...) in IEASYSxx
 - Default is IGGCAT00
 - PARMLIB concatenation & multiple members supported
 - Catalog defaults taken if no PARMLIB member found
 - Support for specifying:
 - VVDSSPACE - VVDS space defaults
 - NOTIFYEXTENT – desired percent threshold of the extents allocated for a catalog before system issues IEC361I to warn that the catalog is becoming full
 - TASKMAX – Upper limit on CAS service tasks
 - Whether to enable extension records for user catalog aliases
- Warning message for usercatalog delete
 - Confirmation WTOR IDC1999I will be issued if a DELETE UCAT RECOVERY command is issued
 - DELRECOVWNG(YES) in IGGCATxx
 - You can enable/disable this enhancement by:
MODIFY CATALOG,{ENABLE|DISABLE},DELRECOVWNG
- Automatically fix SMS CDS data set attributes
 - Health check for NOREUSE in R12
 - The CDS is automatically altered to **REUSE** in R13 when it is defined as being NOREUSE

There is a new catalogue PARMLIB member being introduced in this release. Until now, the only way to customize the catalogue environment or cat, is the SYS1.NUCLEUS member SYSCATxx, or PARMLIB LOADxx member, there was only one line available there and it has long been filled, preventing new parameters from being introduced. So a new parameter called IGGCAT00 has been introduced. It supports PARMLIB concatenation; catalogue defaults will be taken if you do not specify a member, or the member is not found. The support for specifying VVDSSPACE, NOTIFYEXTENT, TASKMAX, and so on – now these parameters belong in IGGCAT00. So IGGCATxx will allow you define catalog system parameters and make permanent changes between IPLs. If you issue DELETE UCAT with the RECOVERY command, the system will issue a WTOR ICD1999I and ask you if you are sure that it was what you want to do. The way you implement that function is by DELRECOVWNG YES in IGGCATxx and you can enable or disable that dynamically by issuing the MODIFY CATALOG,{ENABLE|DISABLE},DELRECOVWNG command.

Starting in R13, the SMS CDS attribute is automatically altered to REUSE. Until now, there has been a health check for NOREUSE. Now control data sets that are defined with NOREUSE will be automatically set to REUSE beginning in R13.

Improving usability and skills (3 of 3)

- Automatic cross-sysplex UCB updates for DFSMSdss COPY or RESTORE operation and DFSMSHsm Fast Replication Backup and Recovery processing
 - Specify a new REFUCB keyword in DEVSUPxx:
 - ENABLE|DISABLE(REFUCB)
 - Designed to issue VARY automatically on sharing systems when these operations change volume serial, VTOC pointer
- Better O/C/EOV Messages
 - Additional information so you don't have to look up the message
 - New DEVSUPxx parameter to activate:
 - OCE_ABEND_DESCRIP = YES | NO
 - Example:
 - IEC145I 413-40,IFG0194F,RDASL1,RDSL1,SYSUT1,0920,,DATASET
 - ERROR DESCRIPTION:
 - THE DEVICE DOES NOT SUPPORT THE RECORDING MODE REQUESTED BY THE USER OR DETERMINED BY THE SYSTEM.
 - END ERROR DESCRIPTION: IEC145I

When you have a function you are performing that updates the UCB for an online volume and that may have been a target of DFSMSdss COPY or RESTORE operation, or an ICKDSS REFORMAT NEW RETARG operation, these operations may cause the volume serial and the RETARG location on the target volume to change.


Before the volume can be accessed by any system in the sysplex, or any remote system, the UCB must be refreshed with the new values. Automatically refreshing the UCB will eliminate the need to manually vary the volume offline and then online.

The DFSMSdss now builds an event notification ENF64, to notify listeners that a volume has been a target of a copy or restore operation, and therefore, the volumes here may change.

There are also now better open and close volume messages as in the example shown here. The old message, shown in red, was the IEC145I 413-40 message. If you implement this DEVSUPxx function with the OCE_ABEND_DESCRIP=YES parameter, then the blue part of the example message will show up, giving you a more verbose description of the error.

z/OS R13 simplification enhancements

- Health Checker Framework enhancements
 - Greater ability to schedule health checks
 - New SYNCVAL keyword in HZSPRMxx and MODIFY
 - Ability for checks to raise message severity as conditions change
- New Health Checks:
 - Two new checks for Allocation – intended to warn about potential Allocation deadlock conditions
 - Detects tape library initialization errors with suggestions on how to resolve.
- New Migration Checks for:
 - Warn when zFS configuration option is not set to sysplex=filesys
 - Verify new symlinks added to enable read-only root in R13
 - OA35636 and OA35605 for R11 and R12
 - Warn the z/OS Console Mode has not been set
 - Available for R10 and later with OA32930



The z/OS Health Checker framework is enhanced in release 13. It now allows you to specify when a health check should be run the first time and when it should be repeated. Using a new SYNCVAL keyword in the HZSPRMxx PARMLIB member, or on the Health Checker MODIFY command. This will allow you to run health checks more predictably at more predictable times. To increase the severity of their exceptions as system conditions change, you have the capability to change the message's severity at that time.

There are new health checks that are introduced as part of release 13. Two new checks for allocation are intended to warn about potential deadlock conditions pending; and to detect tape library initialization errors, giving suggestions on how to resolve them.

There are also migration specific health checks; one is intended to warn when the zFS configuration option is not set to sysplex=filesys. This health check is also available in releases 11 and 12 with APAR OA35465. Another health check verifies any new symlinks that are added to enable read only root in R13, rolled back to 11 and 12 with the APARS shown here.

There is a new health check to warn you that the z/OS console mode has not been set. It has also been available since release 10 with the APAR shown on this slide.

Scalability and performance (1 of 6)

- True cross-system sharing of zFS across a sysplex
 - Direct I/O from all sharing systems
 - No more function-shipping
 - Significant zFS file system performance improvements expected; most measurements showed a 50-150% improvement*
- IEBCOPY improvements
 - Much better performance expected for some operations; elapsed time reductions from 19-70%* for:
 - PDS-to-PDS COPY
 - PDS-to-sequential unload
 - PDS compress
 - Removed requirement for APF authorization

* Note: Performance improvements are based on internal IBM laboratory tests. Your results will vary. I/O performance improvements measured for fully shared zFS ranged from very small to 900%, with the majority of workload conditions tested falling between 50% and 150%. The actual amount of improvement will depend on the environment (monoplex or Parallel Sysplex®) and the type of file processing being done. IEBCOPY improvement will depend on the amount of data being copied, the record format, the record length, and the block size.

In R13, you can share the zFS file system across a sysplex without any doing function shipping. This means you can do direct I/O for read/write from all sharing systems within the sysplex. Test results show significant zFS file system performance improvements – of course, your results may vary.

There have been several IEBCOPY improvements – including much better performance for some operations; the lab measured elapsed time reductions from 19-70% for PDS to PDS copy, PDS to sequential unload, and PDS compress.

Also the requirement for IEBCOPY to be APF authorized has been removed.

Scalability and performance (2 of 6)

- High-Performance FICON® (zHPF) improvements planned for 4Q2011*
 - zHPF to support certain I/O transfers for QSAM, BPAM, and BSAM
 - Better I/O performance expected with no application changes
 - Extends current zHPF support for VSAM, Extended Format sequential, zFS, and PDSE data sets to support:
 - Basic nonextended format Physical Sequential data sets
 - Basic and large format sequential data sets
 - Requirements:
 - z/OS V1.13, z/OS V1.12, or z/OS V1.11 with PTFs documented in the hardware PSP buckets
 - zEnterprise™ System server with channels that support zHPF and a minimum Machine Change Level (MCL)
 - HMC V2.11.1
 - Support Element V2.11.1
 - IBM System Storage® DS8700 or DS8800 series with new DS8000® licensed machine code
 - Enable in IGDSMSxx member of PARMLIB: SAM_USE_HPF(YES|NO)
 - Default NO on z/OS R11-R12, YES on z/OS R13
 - See Statements of Direction

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

7

Improvements in R13

© 2012 IBM Corporation

High-Performance FICON (zHPF) improvements are planned for 4Q2011. In addition to support that is already there, QSAM, BPAM, and BSAM access methods will be supported for better I/O performance. Expect current zHPF support to be extended for VSAM, Extended Format sequential, zFS, and PDSE data sets. This support includes Basic nonextended format Physical Sequential data sets, and Basic and large format sequential data sets. In order to support these data sets, there are some prerequisites – you must be at z/OS V1.13, or at z/OS V1.12 or z/OS V1.11 with appropriate PTFs documented in the hardware PSP buckets. You must be running on a zEnterprise System server with channels that support zHPF and a minimum Machine Change Level (MCL) installed. The HMC firmware must be at level V2.11.1 – this also applies to Support Element V2.11.1. You will need IBM System Storage DS8700 or DS8800 series with new DS8000 licensed machine code. You enable this by adding to the IGDSMSxx member as a PARMLIB: and specify SAM_USE_HPF(YES or NO).

Remember that the default value is NO on z/OS R11 and R12, but has been changed to YES on z/OS R13.

Scalability and performance (3 of 6)

Taking z/OS storage volumes to the extreme - EAV

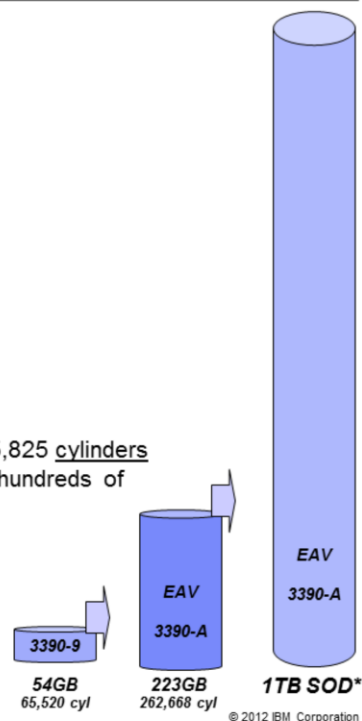
- z/OS R10 introduced Extended Address Volumes with support for VSAM (incl. zFS)
- z/OS R11 added Extended Format Sequential and support for data sets spanning the 64K cylinder line
- z/OS R12 added:
 - PDS and PDSE (including load modules and program objects libraries)
 - Sequential, BDAM, GDG, LPALIB, LPA list, link list data sets, SYSn.IPLPARM, SVCLIB
 - Catalogs, VVDSs, JES2 and JES3 spool and checkpoint, JES3 JCT
 - DFSMSrmm, DFSMSHsm data sets
 - Standalone Dump data set and AMASPZAP support
 - VSAM AIX® support in Language Environment
- z/OS R13:
 - SDSF support for output data sets
 - FTP support for SMS-managed and non-SMS-managed PS; basic and large format; PDS and PDSE; and GDG data sets
- No support for above the line for:
 - Imbed and Keyrange attributes, incompatible CA sizes for VSAM
 - NUCLEUS, SVCLIB, LOGREC, VTOC, VTOCIX,
 - RACF® databases, Page data sets, HFS data sets
 - PARMLIB concatenation data sets
 - XRC Control, Master, or Cluster non-VSAM data sets

Extended Address Volumes are the 223 gigabyte volumes that were used back in release 10. The key point on this chart is to show what has been added in release 13. In release 13, extended address volumes now support allocations of SDSF output to data sets. You can also have FTP support for SMS-managed and non-SMS-managed physical sequential data sets, basic and large format data sets, PDS and PDSE data sets, and GDG data sets.

What is still not supported “above the line” – that is above the 65k cylinder boundary – is the Imbed and Keyrange attributes for a VSAM dataset. Also, NUCLEUS, SVCLIB, LOGREC, VTOC, VTOCIX, RACF databases, Page data sets, PARMLIB concatenation data sets, XRC Control, Master, or Cluster non-VSAM data sets are still not supported.

Scalability and performance (4 of 6)

- One TB Extended Address Volumes*
 - z/OS R13 planned to support EAVs up to 1 TB volume
 - Requirements:
 - IBM System Storage DS8700 or DS8800
 - New DS8000 licensed machine code
 - Availability planned for 4Q11 (SoD)
 - Also planned for z/OS R12 with PTFs
- VVDS maximum size increase
 - For VVDSs in and out of EAS
 - Maximum VVDS space increased from 5,460 tracks to 5,825 cylinders
 - Increases practical maximum number of data sets from hundreds of thousands per volume to millions per volume



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

9

Improvements in R13

R13 has a Statement of Direction that says the 223 gig volume that has been supported has been taken to the next level, so one terabyte Extended Address Volumes are planned to be supported.

To use one terabyte Extended Address volumes, you need release 13, or release 12 with the appropriate PTFs. You also need IBM System Storage DS8700 or DS8800, with the new LIC levels for the machine. Availability for this is planned by mid December 2011. Also in plan is support on z/OS R12 with PTFs.

In order to put more data sets on these larger volumes, the VVDS maximum size has been increased for VVDSs both in and out of Extended Address Space. The Maximum VVDS space increased from 5,460 tracks to 5,825 cylinders. This increases the practical maximum number of data sets from hundreds of thousands per volume to millions per volume.

Scalability and performance (5 of 6)

- New JCL parameter, FREEVOL=EOV
 - Specifies that a tape for part of a multivolume data set be available at end of volume rather than end of step
 - Allow other jobs to use the tape immediately
 - Allow overlapped processing of multivolume tape data sets
- Support for more aliases per user catalog
 - z/OS R12 increased the maximum catalog size dramatically
 - Existing limit on number of aliases is about 3,500 (depending on alias lengths)
 - New limit in z/OS R13 is over 500,000 (depending on alias lengths)
 - New catalog connector extension record (Type V)
 - Catalog PARMLIB member (IGGCATxx) keyword
 - EXTENDEDALIAS(YES|NO)
 - **Do not specify YES until all systems that will process the catalog are at R13!**
 - New command:
 - MODIFY CATALOG,[ENABLE|DISABLE](EXTENDEDALIAS)
 - **Do not issue until all systems that will process the catalog are at R13!**

Customers in today's fast paced world are under pressure to finish their batches as quickly as possible because the batch window is shrinking very rapidly. These performance improvements are designed to make that happen. A new JCL parameter, FREEVOL=EOV (End of Volume) specifies that a tape for part of a multivolume data set be available at end of volume rather than end of step. This allows other jobs to use the tape immediately, allowing overlapped processing of multivolume tape data sets.

Under support for more aliases per user catalog z/OS R12 increased the maximum catalog size dramatically. The existing limit on number of aliases is about 3,500 (depending on alias lengths). The new limit in z/OS R13 is over 500,000 (again, depending on alias lengths). To take advantage of these higher limits, you need to define a new catalog connector extension record called Type V, and you need to update your PARMLIB IGGCATxx member with the keyword EXTENDEDALIAS(YES or NO). **Do not specify YES until all of the images in the sysplex that are using that catalog are at Release 13!** That is, you should not implement this function in a mixed environment, but wait until all the images are at release 13.

You can do this dynamically as well with the MODIFY CATALOG command as shown here, but again, **do not do this until all systems that will process the catalog are at R13.**

Scalability and performance (6 of 6)

- IFASMF DL Improvements
 - Avoid reading to end of log stream
 - IFASMF DL starts reading a log stream at a point (approximately) representing a specified time
 - New SMARTENDPOINT keyword specifies that IFASMF DL should stop reading a log stream when a point representing double the maximum MAXDORM value (2 hours) is reached
 - SMARTEPOVER keyword tells IFASMF DL to use a specified value rather than 2 hours
 - SMARTENDPOINT available on z/OS R10 and up with OA31737
 - SMARTEPOVER available on z/OS R10 and up with OA34374
- Allow entire log stream to be archived or deleted
 - Treat log streams as though they were SMF data sets
 - Will reset log stream starting point to next new block
 - Available for z/OS R11 and up with the PTF for APAR OA34589
- RMODE 64 The next step...
 - Allow execution of enabled code above 2G
 - Support for code above 2G that calls no system services and is not loaded by normal system "load" methods
 - Handle and resume after I/O and external interrupts

The SMF DUMP program for processing SMF log streams has been enhanced to reduce the time it takes to extract the data from the log stream. The new function is designed to allow you to specify a new SMARTENDPOINT keyword when running the utility to limit the amount of data read from the log stream that is being processed. IFASMF DL is also enhanced to allow you to specify that an entire SMF log stream is archived or deleted. Archive and delete processing is designed to process all the SMF data that that exists in a log stream, rather than a subset of the data. This is intended to allow you to easily migrate SMF data management processes based on archiving the entire content of the SYS1.MAN SMF data sets to use SMF log streams.

The SMARTENDPOINT keyword is available on versions 10 or later. The APARs you will need are listed here.

The Archive and Delete processing has been rolled back to release 11.

RMODE 64 now allows you to have an executable above the bar, or above the 2 gigabyte line, with a couple of restrictions. Code that is above the bar cannot call system services, and cannot be loaded by normal system load methods. You are also responsible to resume after any I/O or system interrupt has been processed.

Improving availability (1 of 3)

- JES2 SPOOL Migration*
 - Dynamically remove a SPOOL volume using \$T M SPOOL
 - Also, can enlarge an existing spool data set using \$TSPOOL,SPACE
 - For example, in combination with Dynamic Volume Expansion
 - Planned availability 4Q11 with the PTF for APAR OA36158
- JES3 Dynamic SPOOL Addition
 - Add a SPOOL volume without a JES3 restart using the *MODIFY CONFIG command
- Improved Channel Recovery
 - Remove paths to all devices affected by a path error
 - Avoids repeated recovery for path errors as I/Os are driven to more devices along the path
- zFS Internal Restart
 - Automatic recovery from severe PFS layer problems
 - Remounts all mounted zFS file systems
 - Accessing open files may result in I/O errors or EAGAIN until refresh completes
 - No configuration changes incorporated during restart
 - Can also be operator initiated with new commands
 - MODIFY ZFS,REFRESH to refresh zFS
 - MODIFY ZFS,NSVALIDATE to validate control blocks and refresh if needed

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

12

Improvements in R13

© 2012 IBM Corporation

JES2 will help you now in release 13 and allow you to discontinue the use of a spool volume dynamically in a relatively short period. There is a new \$T M SPOOL command that will allow you to specify that either a new spool data set of equal or greater size on a different volume will be used to replace an existing spool data set, OR that another existing spool data set with sufficient contiguous space be used to replace an existing spool data set.

Also a new TSPOOL, SPACE command will allow you to increase the size of an existing spool data set. This function is intended to help you improve availability when removing spool volumes from a JES2 system, or a JES MAS, and expected to be much faster than draining the spool volume.

There is also an improvement in channel recovery in z/OS release 13. IOS will track path related errors and automatically remove failing paths from all affected devices for the affected control units. This capability should reduce the time it takes the system to recover from path related errors and help prevent system performance problems that can occur when significant time is spent in repetitive channel error recovery. Less time will be spent loading up the error recovery modules and running the error recovery.

The zFS internal restart enhancements that are part of release 13 will automatically recover from any physical file system layer problems. After recovering, it will perform the remounts of all mounted zFS file systems. It will start accessing the files that the applications may be running, which may result in I/O errors or EAGAIN until the refresh completes. That is the only thing that you need to be aware of when the zFS internal restart occurs. No configuration changes are incorporated when the internal restart is happening. A z/OS command, MODIFY ZFS,REFRESH allows an operator to initiate an internal restart. You can also do a validation of the control blocks by issuing the MODIFY ZFS,VALIDATE command.

Improving availability (2 of 3)

- CMDS command enhancements
 - CMDS ABEND,CMD=xxxxxxxx,ID=nnnn introduced many releases ago
 - Enhanced in R12 to enforce “non-abendable” commands
 - CMDS FORCE command added for z/OS R13; intended to be used when only alternative is IPL
- Parallel FTP tool now part of z/OS
 - IBM z/OS Problem Documentation Upload Utility
 - New program name, **AMAPDUPL**
 - Alias **MTFTPS** for compatibility
- Message flood automation processing improvements:
 - Increase message ID limit from 50 to 1024
 - Allow up to 128 address spaces to be tracked per system
 - Allow the default message set to be identified in a PARMLIB member
 - Intended to increase the scope of message flood automation, improve its usability, and help improve system availability

The CMDS command enhancements in version 12 allow killing a command that was not ABEND-able.

In release 13 there is a new command called FORCE, so you can force that command if all else fails. Be aware you should not use this all the time, because it is really not intended for the operator to use, but is provided for when there is really no other alternative other than to IPL the machine. Also the IBM z/OS problem documentation upload utility has now become a part of z/OS so you do not need to download it separately; the name is AMAPDUPL and the alias name MTFTPS is kept for compatibility reasons.

Message flood automation or MFA has many enhancements listed here, including increasing the message ID limit from 50 to 1024, and allowing up to 128 address spaces to be tracked per image. The default message set can now be identified in a PARMLIB member. These things are intended to increase the scope of the message flood automation, improve its usability, and help improve system availability.

Improving availability (3 of 3)

- IPL devices in subchannel sets other than 0
 - IPL, IODF, SADMP volumes supported for IPL from Subchannel Set 1 or 2
 - Allow use of PPRC secondary devices for IPL after primary fails
 - Requires:
 - zEnterprise System (z196 or z114)
 - HMC V2.11.1
 - Support Element V2.11.1
 - Minimum Machine Change Level (MCL)
 - Also available for z/OS R11 and R12 with the PTF for APAR OA35140
- DADSM dynamic exits support
 - IGGPRE00
 - IGGPOST0
- DADSM and CVAF support for concurrent service
 - Dynamically update without IPL to help improve system and application availability
- ASID Reuse
 - DEVMAN address space now reusable
 - CATALOG, LLA, VLF, z/OS UNIX RESOLVER, TCP/IP, DFSMSrmm, and TN3270 already reusable

Now from an IPL point of view, z/OS 1.13 will allow devices used early during IPL processing – or ‘nip’ time - to be accessed using subchannel set 1 or subchannel set 2. This is to allow the use of PPRC secondary devices defined using the same device number and a new device type in alternate sub channel set to be used for IPL, IODF volume, and standalone dump volumes when needed. This support requires a zEnterprise System at z196 or z114; HMC firmware code needs to be v2.11.1, as should be the support element. The machine change level also has to be current.

This support is also made available down to release 11, with the PTF for the APAR listed on this chart.

The DADSM dynamic exit now supports the IGGPRE and IGGPOST exits.

DADSM and CVAF components have now been designed to commit, dynamically updating their programs without IPL. This should improve system and application availability.

To make ASIDs more reusable up to this point, CATALOG, LLA, VLF, and the other address spaces listed here have released their ASIDs when they terminate so that the ASID is reusable. DEVMAN has now been added to this list as and now it is also reusable.

z/OS Predictive Failure Analysis and runtime diagnostics

- **z/OS PFA checks intelligently analyze the rate and trends of specific z/OS system resources and report on z/OS system anomalies, potentially helping you to identify and avoid failures. PFA checks analyze:**
 - Common storage usage checking (R10), LOGREC arrival rate detection (R10)
 - Frame and slot usage checking (R11), Message arrival rate detection (with z/OS R11)
 - SMF Message Arrival rate detection (R12), ENQ and SPOOL utilization tracking (R13)
 - Additional customization (R12)
 - You can specify atypical jobs and address spaces to be excluded from learning algorithms
- **Runtime Diagnostics helps you to pin point the source of a 'soft failure' by looking at z/OS system resources in real time and reporting on areas of contention. RTD checks for:**
 - Select critical messages in the Operlog (z/OS R12)
 - Address spaces with high CPU usage or which might be in a loop (z/OS R12)
 - Address spaces suspended in local lock contention (z/OS R12)
 - System address spaces that are ENQ "waiters" (z/OS R12)
 - Evaluates GRS latch and z/OS UNIX System Services file system latch contention (z/OS R13)
- **Autonomics (R13) - RTD is automatically invoked from PFA when PFA detects very low SMF arrival rates, RTD problem notification sent to PFA. (R13)**
- **F HZR,ANALYZE**

z/OS PFA checks intelligently analyze after modeling the rate and trend of specific z/OS system resources that may cause problems if the resource gets exhausted. The implementation of PFA is partly done in Health Checker, so the PFA health checks run in Health Checker.

You can see here what was done in release 10 and 11. What has been added to release 13 is the ENQ and SPOOL utilization tracking.

The things that an operator typically does when there is a system problem, runtime diagnostics do that programmatically. If PFA detects that there are issues, it can then call the runtime diagnostics to run those tests. If you think there are problems you can issue MODIFY HZR, ANALYZE command on a system, or on all the system in the sysplex.

Integrating new applications and supporting industry and open standards

Simplified XCF interfaces for passing messages in a Parallel Sysplex

- New services allow a server to be established to process messages and for messages to be sent across the sysplex without first joining an XCF group
- Intended to make it easier to exploit XCF services for applications that do not require the member management and monitoring provided by the XCF group services interfaces

Support to allow tasks having subtasks in a WLM enclave to leave it, and for subtasks of a task joining an enclave to be joined

- Available now on z/OS V1.11 and z/OS V1.12 with the PTFs for APARs OA33344 and OA33406

z/OS XML System Services will support a binary XML format

- Extensible Dynamic Binary XML (XDBX)
- XDBX supports a subset of XML
- Expected to provide performance improvements for validating parsing operations compared to conventional XML text documents
- **Planned to be enabled on z/OS V1.13 with the PTF for APAR OA36712 in 4Q11**

A lot of work has been done to simplify the XCF interfaces for passing messages between peers in a parallel sysplex. So that is the precise API that SDSF now is taking advantage of to lift the restriction of using WebSphere MQ that was mentioned a few slides back.

XCF provides a simplified system of interfaces for passing messages within a sysplex. These services allow a server to be established to process messages and for messages to be sent across the sysplex without first doing a join to an XCF group. This should make it easier to exploit XCF services or API's for applications that do not require the member management and monitoring that is provided by the XCF group service interface.

WLM has several enhancements in release 13. Tasks having subtasks in a WLM enclave can now leave, and subtasks of a task joining an enclave can be joined. This has been made available to release 12 and 11 with the APARs shown here.

z/OS XML System Services now support binary XML format. XDBX supports a subset of XML. It is expected to provide performance improvements for validating parsing operations compared to conventional XML text documents. This is planned for delivery with APAR OA36712 in late 2011.

Enhancements in security (R13)

- **IKEv2**
 - Initial support with z/OS R12 Communications Server. z/OS R13 adds Network Address Translation (NAT) traversal support for IKEv2 over IPv4.
- **System SSL, ECC**
 - z/OS R12 Communications Server added support for Elliptic Curve Cryptography (ECC), ECDSA (Elliptic Curve Digital Signature Algorithm).
 - z/OS R13 to extend System SSL ECC support for:
 - Creating ECC-style certificates in key database files or ICSF PKCS#11 tokens
 - Creating ECC-style certificates through the Certificate Management Services (CMS) API
 - Enabling ECC for TLS V1.0 and TLS V1.1 handshakes (RFC4492)
 - ECC certificate support with Crypto Express3 Coprocessor (on zEnterprise server)
- **ICSF support for additional HMAC algorithms**
 - Support for FIPS-198, support planned for SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512
- **TN3270 and FTP support for password phrases**
- **SAF Based security for z/OSMF**
 - New general resource class ZMFAPLA
 - Similar to EJBROLES
 - Intended to use instead of repository based security

Support for the Internet key exchange version 2 was provided by Communications Server in release 12. In release 13, Communications Server has added the support for NAT traversal support for IKEv2 over IPv4.

In R12, Communications Server added support for Elliptic Curve Cryptography and Elliptic Curve Digital Signature Algorithm. In R13 support is extended to include creating ECC-style certificates in key database files, creating ECC-style certificates through the certificate management services APIs, enabling ECC for TLS V1 and TLS V.1. handshakes, and ECC certificate support with Crypto Express3 Coprocessor on z196 and z114.

ICSF now supports FIPS 198 and support is planned for SHA 1, 224, 256, 384 and 512. In addition that, TN3270 and FTP now support password phrases in addition to passwords.

When z/OSMF came out up until release 12 they provided their own authentication methods by providing a repository based security. Now in release 13 of z/OSMF, SAF based security is allowed, meaning there is a new general resources class that you can define in RACF - ZMFAPLA is very similar to EJBROLES definitions.

z/OS Security Server – RACF

Helping to address security and compliance** guidelines

▪ RACF

- RACF Remote Sharing Facility (RRSF) support for TCP/IP (in addition to SNA APPC)
- Support for generating Elliptic Curve Cryptography (ECC) secure keys (using Crypto Express3 Cryptographic Coprocessors (CEX3C) available with zEnterprise servers)

▪ Tivoli® Directory Server for z/OS (LDAP)

- Support for SHA-2 and salted SHA-2 hashing of user password attributes. Addresses:
 - Need for stronger hashing and cryptographic algorithms
 - Enhanced interoperability with distributed IBM TDS, openLDAP, and other LDAP servers.
 - The National Institute of Standards and Technology (NIST) policy for the use of hash functions.
- Support for LDAP administrators to delegate LDAP administrative authority
 - Can improve LDAP administration flexibility, help improve auditability, and help improve security
- Support for DB2® 9 for z/OS (5635-DB2) backend for scalability of large LDAP deployments
- Improved interoperability between z/OS applications and Microsoft Active Directory environments for Kerberos
- Support for RFC 2696 and RFC 2891 for improved LDAP sorted search performance

** It is the customer's responsibility to identify, interpret, and comply with laws or regulatory requirements that affect its business. IBM does not warrant that its products or services will ensure that the customer is in compliance with the law.

18

© 2012 IBM Corporation

The RACF Remote Sharing Facility up until now has supported SNA advanced peer-to-peer communication, but in Release13, it will now support sharing over TCP/IP as well.

Tivoli Directory Server - use to be called LDAP - has a lot of enhancements for Release 13. SHA2 and salted SHA2 hashing using password attributes should address the need for stronger hashing and cryptographic algorithms and provide enhanced interoperability with Tivoli Directory Server running on distributed platforms, open LDAP, and other LDAP servers.

Tivoli Directory Server now supports LDAP administrators delegating LDAP administrative authority. You can now use DB2 version 9 to improve the scalability of large LDAP deployments. R13 supports improved interoperability between z/OS applications and Microsoft Active Directory environments for Kerberos. And finally, support for RFC 2696 and RFC 2891 improves LDAP sorted search performance.

z/OS and the IBM zEnterprise Unified Resource Manager

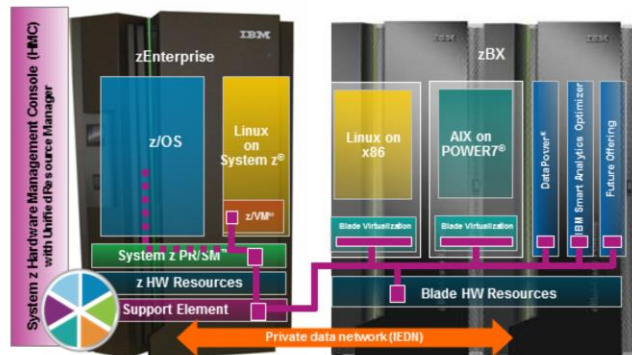
- z/OS on zEnterprise (z196 or z114) connects with the zEnterprise BladeCenter® Extension seamlessly
- **Unified Resource Manager defines the ensemble, provisions the new management and data networks, and can manage virtual environment.**
- **IBM zEnterprise Unified Resource Manager:**
 - can manage 'virtual servers' (z/VM and blade)
 - can monitor 'virtual server' workloads and z/OS workloads.

z/OS needs

- New HCD (or HCM) definitions for the new networks
 - New OSA CHPIDs: OSM for management network and OSX for data network
 - IBM SOD* HiperSockets™
- z/OS Communications Server configuration to enable z/OS to participate in a zEnterprise ensemble

Additional z/OS function ...

- New z/OS agent (Guest Platform Management Provider) can send high level z/OS WLM data to zEnterprise Unified Resource Manager
- z/OSMF Resource Monitoring can monitor z/OS and Linux workloads



19

Improvements in R13

© 2012 IBM Corporation

This is what zEnterprise hybrid computing could look like - a z196 or a z114 attached to a system z Blade Center Extension. For z/OS to participate in a hybrid computing environment, new HCD or HCM based definitions are needed. There are two OSA CHPIDs, one is OSM and one is OSX. OSX is the intra-ensemble data network, which is where the customer data flows between the hybrid ZBX and the zEnterprise system, and the OSM, which is the management network, which nobody touches.

There is an additional z/OS function that is an agent that runs everywhere at the endpoint of the purple line depicting a box ending up in each compute resource that is what's called GPMP, or Guest Platform Management provider. It runs inside the operating system and collects data on behalf of the IBM Unified Resource Manager.

Monitoring of z/OS as well as Linux workloads is also implementing using RMF™. The Unified Resource Manager itself resides in the hardware management console as firmware, and also in the support element inside the zCEC.

The IBM zEnterprise Unified Resource Manager - which is sometimes called zManager - can manage all virtual servers, including z/VM and virtual servers running on the blades, but it can only monitor virtual servers that are running workloads under z/OS control.

Self-managing capabilities

RMF monitoring for zEnterprise ensembles

- RMF provides CIM-based performance data gatherers for:
 - Linux on System z, AIX Systems and Linux on IBM System x®
 - Designed to provide a consistent monitoring solution for zEnterprise ensembles
- Response time distributions calculated by WLM and reported by RMF for velocity and discretionary goals**

- As for response time goals, reported in 14 “buckets”
- Unlike response time goal reporting, mid-points can be recalculated and changed from time to time

RMF support for additional contention reporting

- For system suspend lock, GRS enqueue, and GRS latch contention
- New Postprocessor Serialization Report available in XML output format
- New SMF Type 72 subtype 5 records
- Help make it easier to respond to serialization-related performance issues.

RMF monitoring for zEnterprise ensembles was introduced in the picture on the previous slide. RMF provides common information model based performance data gatherers that could be running on Linux on System z, on the zCEC, AIX operating system running on Power 7 blades in the zBX, and Linux running on IBM System x blades in the zBX. By the end of 2011, Windows server should run in the zBX on the System x blade servers. The intent here is to give you a consistent monitoring solution for the entire zEnterprise ensemble.

Response time distribution is calculated by WLM and RMF and reported for velocity and discretionary goals. RMF support for additional contention reporting will be provided for system suspend locks, GRS enqueues, and GRS latch contention. There is a new SMF type 72 subtype 5 record which will contain that. Hopefully it will help make it easier to respond to any performance related problems that may have anything to do with serialization.

z/OS R13 - The foundation for modern batch

- Expand existing COBOL applications!
 - **The z/OS Batch Runtime environment**, provides Java-to-COBOL interoperability, for transactional updates to DB2, and for sharing database connections between Java and COBOL (R13)*
Ideal for processing for computationally intensive programs and extensions
 - Use JZOS Batch Toolkit for z/OS for efficient use of z/OS System interfaces for Java batch (IBM Java SDKs for z/OS)
 - Leverage specialty engines!
- Simplified programming!
 - JES2 JCL enhancements provide in-stream data in catalogue procedures, more options on setting job return codes, and the ability to stop and hold a job at the end of a step (not just at the end of the job) give much more granularity and control (z/OS R13)
 - An new REST API allows you to submit z/OS batch jobs and retrieve z/OS batch job information from distributed systems as well as z/OS systems; and is intended to make z/OS batch processing much more accessible to distributed systems and web-based processes (z/OS and z/OSMF R13)
- Shorter batch windows
 - Allow overlapping processing for multi-volume data sets (FREEVOL=EOV, R13)
 - Avoid recalling migrated datasets, just to delete them (IEFBR14, R11)
 - “Pipe” data between two batch jobs to enable these jobs to perform reads and writes concurrently (BatchPipes®, 5655-D45)
- Real time batch
 - WebSphere Compute Grid delivers a resilient, highly available, secure, and scalable runtime with container-managed services for batch applications
 - Capable of supporting 24x7 batch and OLTP processing, and parallel computing

* Prerequisites:

IBM 31-bit SDK for z/OS, Java Technology Edition, Version 6.0.1 (5655-R31)

DB2 V9.1 for z/OS (5635-DB2) or later with PTFs

IBM Enterprise COBOL for z/OS V4.1 (5655-S71) or later

Modern batch is a new component introduced in Release 13. IBM research was performed and about 90% of the respondents that participated in that research consider batch to be mission critical, the majority choosing to run it on System z. Central to batch processing is the COBOL programming language, and COBOL is simple, efficient, robust, and scalable. COBOL assets are almost everywhere and capable of supporting billions of transactions a day. Top analysts do agree that COBOL can be modernized to help revolutionize the batch process. So z/OS 1.13 introduces many new batch programming capabilities. The z/OS batch runtime environment is intended to form the foundation for resilient, highly available, secure, and scalable services for batch applications. The z/OS batch runtime is intended to provide the framework for Java to COBOL interoperability for transactional updates to DB2 and for sharing database connections between Java and COBOL. The intent really is to enable you to reuse valuable COBOL assets by developing new or enhancing existing batch applications with Java. With this capability, for example, you can share a database connection between COBOL and Java, allowing you to use Java subroutines directly instead of using Java stored procedures. The Java portions of your batch applications can be eligible for IBM System z, ZAAP processor or on the IBM System z ZIIP processor or using the ZAAP on ZIIP capability that was introduced in Release 11.

Initially Java, DB2, and COBOL support is provided with IBM 31 bit SDK for z/OS, Java Technology Edition Version 6.1, DB2 version 9.1 for z/OS or later with PTFs and the IBM Enterprise COBOL for z/OS version 4.2 .

z/OS and IPv6



- **IPv4 address pool is exhausted February 3, 2011**
 - <http://www.ipv6news.info/2011/02/04/ipv4-address-pool-is-exhausted/>
 - Now the IPv4 Internet only has the stock of IPv4 addresses held by the regional registrars and Internet Service Providers (ISPs) to keep it going.
- **z/OS is IPv6 certified!** (http://jtc.fhu.disa.mil/adv_ip/register/certs/ibmzosv110_dec08.pdf)
- **z/OS Communications Server is adding function for IPv6 networks:**
 - **For z/OS R11**
 - Support RFC4941 and RFC5095; and the AES-based AES-XCBC-MAC-96 and AES-XCBC-PRF-128 algorithms - intended to meet new government IPv6 standards
 - **For z/OS R12**
 - Health checks for IPv4 and IPv6 routing
 - Support for DFSMSrmm, IKEv2, ability to Send DNS Queries over IPv6, support for security-related RFC3484 and RFC5014
 - **For z/OS R13**
 - Support for IPv6 intrusion detection security equivalent to that provided for IPv4, integrated with the Configuration Assistant (in z/OSMF)
 - Support for IPv6 checksum and segmentation offload enhancements and for LPAR-to-LPAR checksum offload for both IPv4 and IPv6 packets available with OSA-Express4S QDIO (announced July 12 2011)
 - TCP/IP segmentation and checksum processing on OSA card and not on CP

The world has run out of IPv4 addresses, but z/OS is IPv6 certified. You can read about that certification at the second link on this slide.

This slide lists what has been enhanced for releases 11, 12, and 13. In release 13, IPv6 detection security equivalent to what is provided in IPv4 is integrated into z/OSMF and part of z/OS Communications Server Configuration Assistant. There is now support for IPv6 checksum and segmentation offload enhancements; at the same time LPAR to LPAR checksum offload for both IPv4 and IPv6 packets and it is available with OSA Express4S, which is a new feature on the z196 and z114 in QDIO format. TCP/IP segments and checksum processing runs on OSA hardware - it does not use the general purpose engine within the zCEC.

Networking

- More flexibility for specifying reserved TCP/IP port ranges
- New CSSMTP server design for better memory and JES resource management when retrying mail send operations
- Improved z/OS system resolver processing when name servers are unresponsive
- More VLANs per OSA port
 - Define up to 32 VLANs per OSA port per IP version
- Autonomic recovery for APPN routing tree corruption
- New design to monitor for CSM-constrained conditions and taking specified recovery actions
- Enterprise Extender connectivity tests initiated using the DISPLAYNET,EEDIAG,TEST=YES command when firewalls block ICMP messages expected to complete more quickly

You can now specify and reserve TCP/IP port ranges in your TCP configuration files. The new CSSMTP server has implemented better memory and JES resource management when doing mail send operations. The z/OS system resolver processing has been improved when it is trying to talk to domain name servers that happen to be unresponsive. There are now more VLANs per OSA port, it used to be 8 and now you can define it up to 32 VLANs per OSA port. The reason for this is that you need more VLANs to support the z hybrid environment where a zBX may be attached to the system zCEC. New also in release 13 is autonomic recovery for tree corruption when that happens in APPN routing.

There is a new design to monitor for Communications Server CSM constraint conditions.

The Enterprise Extender connectivity test can be started by issuing the command DISPLAYNET, EEDIAG, TEST=YES, when firewalls may be blocking ICMP messages and it should complete more efficiently and more quickly.

Summary

- New features for V1R13 in:
 - Availability
 - Simplification
 - Hardware support
 - Economics and platform efficiency
 - Application middleware workload enablement
 - Installation and migration
 - Security

This presentation has covered the new enhancements in release 13 of the z/OS operating system. You can also explore the technical details of the new features under the categories of availability, simplification, hardware support, economics and platform efficiency, application middleware workload enablement, installation and migration, and security.

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, AIX, BatchPipes, BladeCenter, DataPower, DB2, DS8000, FICON, HiperSockets, Parallel Sysplex, POWER7, PR/SM, RACF, RMF, System i, System Storage, System Storage DS, System x, System z, Tivoli, WebSphere, z/OS, z/VM, and zEnterprise are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. Microsoft, Windows, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2012. All rights reserved.