

z/OS V2 R2 Sysprog Goody Bag

Bob Rogers

IBM Distinguished Engineer, Retired

Trident Services, Inc.; MSP Communications; IBM Corporation



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides education, professional networking and industry influence.

Copyright (c) 2015 by SHARE Inc.  Except where otherwise noted, this work is licensed under <http://creativecommons.org/licenses/by-nc-sa/3.0/>



Session objectives

The objective of this presentation is to overview some of the small and medium size enhancement to the base elements of z/OS 2.2 that would be of interest to Systems Programmers. The intention is not so much to be tutorial of the details but rather to provide context and history of the enhancements with respect to the problems they are intended to solve.


Abstract

This edition of the z/OS Goody Bag, in addition to many software-only enhancements, covers the numerous items supporting new or improved capabilities of the IBM z13 processor. Support for the z13 hardware includes SMT, SIMD, higher n-way and I/O fabric priority. Other software enhancements include parallel batch scheduling, CA level locking for RLS, dynamic JES2 Checkpoint tuning and expansion, private storage tracking with PFA. LOGREC deallocation and much more.

There are a lot of interesting enhancements so we are taking the Rolling Stone approach of “all the news that fits”.

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Trademarks



z/Architecture*
zEnterprise*
z/OS*
SHARE
Educate • Network • Influence

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

AIX*	DFSMSdfp	DS8000*	IBM*	Language Environment*	REXX	System z9*
BladeCenter*	DFSMSdss	Easy Tier	IBM eServer	MQSeries*	RMF	System z10
CICS*	DFSMSHsm	FICON*	IBM logo*	MVS	SYSREXX	Tivoli*
DataPower*	DFSMSrmm	FlashCopy*	IMS	Parallel Sysplex*	System Storage	WebSphere*
DB2*	DFSORT	HiperSockets	InfinBand*	PR/SM	System x*	z10 BC
DFSMS	DS6000*	HyperSwap*	Infoprint*	RACF*	System z*	z10 EC

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

Windows Server and the Windows logo are trademarks of the Microsoft group of countries.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

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

* Other product and service names might be trademarks of IBM or other companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g. zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Complete your session evaluations online at www.SHARE.org/Orlando-Eval



List of Topics

- *Support for SMT on z13*
- *Hiding New Hardware Facilities*
- *Greater Scalability*
- *zAAP-on-zIIP Mandatory on z13*


- *APF List Changes Recorded in SMF*
- *Enhanced System Symbols*
- *IPL Device and Volume Message*
- *Operator Console Enhancement*
- *2 Million Concurrently Open Data Sets*
- *SMFLIMxx parmlib member*
- *REGIONX JCL Keyword*
- *LOGREC Reallocation w/o IPL*
- *SLIP ACTION=CMD*
- *TSO Logon Password Preprompt*

- *PFA Private Storage Exhaustion Check*
- *REXX APIs for Health Checker Persistent Data*
- *Subsystem Deletion Support*
- *JES2 Eventlog for Step Return Codes*
- *JES2 Deadline Scheduling*
- *JES2 Job Execution Control*
- *JES2 Toleration of JES3 JECL*

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Simultaneous Multithreading

- ■ “Simultaneous multithreading (SMT) permits multiple independent threads of execution to better utilize the resources provided by modern processor architectures.”*
- ■ With z13, SMT allows up to two instructions per core to run simultaneously to get better overall throughput
- ■ SMT is designed to make better use of processors
- ■ On z/OS, SMT is available for zIIP processing:
 - ■ Two concurrent threads are available per core and can be turned on or off
 - ■ Capacity (throughput) usually increases
 - ■ Performance may in some cases be superior using single threading



Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Greater Scalability

- Up to 128 CPUs
- Up to 4 Terabytes of memory
- Support for 4 Subchannel Sets

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

zAAP-on-zIIP Mandatory on z13

- The IBM z13 system continues to support zIIPs but does not continue to support zAAPs.
- When IPLed on a z13, z/OS 2.2 will accept a specification of ZAAPZIIP=NO but will force zAAP-on-zIIP (ZAAPZIIP=YES)
- A warning message is issued.

Hiding New Hardware Facilities

- When migrating to a new machine with new facilities, it is often desirable to initially avoid using the new facilities – better to do some “compatibility mode” testing first.
- The MACHMIG statement in LOADxx supports doing this.
- The z13 Vector Extension Facility can be hidden from software with the MACHMIG option VEF ().
- There is no need for for a MACHMIG option to hide z13 SMT capability since it is only actually visible to the operating system. Just not turning it on will suffice to prevent it from being used.

APF List Changes recorded in SMF

- Dynamic changes to the list of Authorized Program Facility (APF) libraries has been supported for many year.
- But, the changes to APF lists have not been recorded in a durable and easy to audit manner.
- z/OS 2.2 introduces a new SMF type 90 subtype 37 record to document dynamic APF changes.
- Here's what in the record:
 - Function
 - Origin (CSVAPF, SET PROG, SETPROG with PROGxx suffix)
 - Data set name
 - Volume ID
 - Timestamp of the operation
 - Jobname of the requester
 - CHKEY (e.g., stepname) of the requester
 - Console ID of the requester
 - Security product user token (Utoken) of the requester

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Enhanced System Symbols

- System Symbols were introduced with two limitations:
 - the symbol names could not be long than 8 characters
 - the symbol values could not be longer than the symbol names.
- z/OS 2.2 raises the limit on keyword length to 16 character.
- The character “_” may now be used within or at the end of a symbol name.
- When a symbol name ends with an “_”, its value can be up to 44 characters in length.
- The size of the system symbol table is increased by 75%
- There are some restrictions, so read the rules and enjoy.

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

IPL Device and Volume Message

- It's really hard to tell that the wrong device has been IPLed until something really bad happens.
- z/OS 2.2 introduces a console message identifies the subchannel set, device number and volume that were the target of the IPL.
- *IOS128I IPL DEVICE: sdddd VOLUME: vvvvvv*
- Even if nothing really bad was going to happen, this message makes it easier to associate an archived syslog with a specific set of maintenance on a volume.

Operator Console Enhancements

- Operators often need to logon to multiple systems in a sysplex.
- This leads to two problems:
 - Operators must use multiple userids to log on to multiple consoles. This is inconvenient and potentially confusing.
 - Unattended operator consoles present a security exposure.
- z/OS 2.2 address both of these problems
 - It allow the same userid to be concurrently logged on to multiple operator consoles within a sysplex.
 - It provide a timeout value such that if an operator does not provide any input from the console within that time value, the userid is automatically logged off.
 - the VARY CN() command can be used to set or change a console's timeout value.
 - It is not supported for: System Console, EMCS consoles, Subsystem consoles, Printer consoles or Consoles in USE(MS|SD)

2 Million Concurrently Open Data Sets

- In z/OS 2.2, control structures and processing are changed to make it possible and reasonable to maintain up to 2,000,000 concurrently open datasets.
 - Scheduler Workarea (SWA) blocks are indirectly addressed by SWA Virtual Address (SVA) tokens. The bits of an SVA are now interpreted as indexes rather than offset to accommodate more SWA blocks.
 - The infrastructure for managing SWA and maintaining SYSDSN are move about 2GB to relieve pressure of virtual below 2GB.
 - Serializing with the SYSZTIOT resource is restructured to shorten exclusive hold time.

SMFLIMxx to replace IEFUSI

- IEFUSI is an old and complicated installation exit point
- In z/OS 2.2, most of the capability of IEFUSI can be achieved through a new parmlib member.
- Filters: (wildcarding i
 - JOBNAME(1-8 values) JOBACCT(1-8 values)
 - STEPNAME(1-8 values) STEPACCT(1-8 values)
 - PGMNAME(1-8 values) USER(1-8 values)
 - JOBCLASS(1-8 values) SUBSYS(1-8 values)
- Limits:
 - REGIONBELOW(1K to 16M) REGIONABOVE(1M to 2G)
 - SYSRESVBELOW(1K to 16M) SYSRESVABOVE(1M to 2G)
 - MEMLIMIT(up to 15P)
- Multiline WTO written when a rule is applied
- Still need IEFUSI to set limits for Dataspaces and Hiperspaces

REGIONX JCL Keyword

- z/OS 2.2 provides a new JCL keyword for JOB and EXEC statements: REGIONX
- REGIONX allows separate specifications for below-the-line and above-the-line region values.
- Examples:
 - REGIONX=(6M,1G)
 - REGIONX=(0M,0M)
- Above-the-bar limit is still specified with MEMLIMIT KEYWORD

LOGREC Reallocation w/o IPL

- In the past, LOGREC has been inflexible. There have been a number of important capabilities that were not supported:
 - Performing maintenance on the DASD backing the LOGREC data set.
 - Switching from LOGSTREAM to DATASET after a logger outage without having the LOGREC data set pre-allocated.
 - Resizing the LOGREC data set.
 - Saving an old LOGREC data set and switch to a new data set.
- In z/OS 2.2 the SETLOGRC command is enhanced:
 - When switching to DATASET, allocate the data set even if IPLed with LOGSTREAM or IGNORE.
 - When switching from DATASET, deallocate the data set.
 - Allowing specification of DATASET or LOGSTREAM resource.
- The system parameter LOGREC= in IEASYSxx is enhanced to allow for specification of log stream name.

SLIP ACTION=CMD

- For facilitate problem determination and to better ensure first failure data capture, it should be possible to have operator commands issued in response to a SLIP trap event.
- In z/OS 2.2, SLIP supports a new action, ACTION=CMD, to issue up to 8 operator commands.

- CMD is allowed with other actions, but not WAIT or IGNORE

- The commands are provided in quotes. A command can be broken into parts with nested parentheses. E.g.

```
SLIP SET,C=071,SDATA=(ALLNUC,SQA,CSA,LPA,LSQA,RGN,  
ALLPSA,SUM,TRT),REASON=20,ID=SS71,CMD=('command 1',('command 2  
part 1', 'command 2 part 2'),('command 3 part 1','command 3 part  
2','command 3 part 3')),A=(SVCD,CMD),END
```

- A new DESC= keyword provides way to show a description in the WAIT state messages or dump header.

```
SLIP SET,...,ACTION=SVCD,DESC='dump description',END
```

TSO Logon Password Preprompt

- TSO Logon can present information about a USERID even in the absence of providing a valid password.
- The existence of valid TSO/E user IDs can be determined by the message response to incorrect LOGONs. Once identified, further user information such as account number, logon procedure and initial command can be discovered.
- This situation is remedied by support being provided in z/OS 2.2 (and 1.13 and 2.1) via APAR OA44855
- A new option, PASSWORDPREPROMPT, is provided. When active it will require all logon information to be successfully entered before any indication of the validity of the user ID is provided.

PFA Private Storage Exhaustion Check

- Detects future exhaustion of private storage under 2G in 6 storage locations within individual address spaces
 - Private user region
 - Private authorized area
 - Private user + private authorized (i.e., private below the line)
 - Extended private user region
 - Extended private authorized area
 - Extended private user + extended private authorized (i.e., private above the line)
- Does not detect exhaustion due to
 - Fragmentation
 - Fast increases of usage that are on a machine-time scale or even faster than one collection interval

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

REXX APIs for Health Checker Persistent Data

- The Health Checker framework allows persistent data which is saved across IPLs to be maintained for individual health checks.
- Unfortunately, the current Read/Write interfaces are only available as Assembler services.
- There are growing number of checks are being written in REXX because they are so easy to implement.
- z/OS 2.2 provides REXX callable functions to read and write persistent data.
- Now the REXX checks can take advantage of a consistent and easy interface for maintaining persistent data.

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Subsystem Deletion Support

- Subsystems have been hard to “undefined”.
 - Subsystem names are reserved when the subsystem is defined, even if it fails to successfully initialize.
 - If a subsystem fails, even if it can be disabled or deactivated, the subsystem name is still reserved and cannot be reused.
- In z/OS V2R2, the subsystem is not defined if the INITRTN cannot be found or is not APF-authorized.
 - This applies to any INITRTN defined via IEFSSNxx, the SETSSI ADD command or IEFSSI macro.
- Also a SETSSI DELETE command to remove a subsystem.
 - This is a logical deletion. Subsystem routines still running are not terminated and no storage associated with the subsystem is freed.
 - SETSSI DELETE,S=SUB1,FORCE
 - SETSSI DELETE,S='abcd',FORCE
 - SETSSI DELETE,S='C1C2C3F1'X,FORCE

JES2 Eventlog for Step Return Codes

- Just having access to the Job return code is often not enough
- In z/OS 2.2, JES2 will automatically create an additional spool data set named EVENTLOG for batch jobs, started tasks and TSO users.
- It will contain information about each step in a job and the completion code for that step.
- EVENTLOG contains key SMF Type 30 subtype records
 - Subtype 1. Job start or start of other work unit.
 - Subtype 4. Step Total.
 - Subtype 5. Job termination or termination of other work unit.
- EVENTLOG contains machine readable records
 - Not intended to be human readable
 - EVENTLOG is non-printable, non-spinnable
 - Not SYSOUT, cannot be accessed by SAPI/PSO/FSS
 - Viewed via spool data set browse

JES2 Deadline Scheduling

- JES2 has historically lacked basic job scheduling controls. Function which has been available in JES3.
- JES2 introduces new keywords on SCHEDULE JCL statement:
 - HOLDUNTl to keep a job in a held state until a specified time.
 - STARTBY to specify a desired time for a job to start.
 - WITH to specify that a job should run on the same system where a reference job is currently executing.
- The target time for a job can be displayed via JES2 command
 - \$dj19,holduntl
 - \$HASP890 JOB(SCHTEST) HOLDUNTl=(2015.029,13:55:00)
 - \$dj19,startby
 - \$HASP890 JOB(SCHTEST) STARTBY=(2015.029,13:55:00)
- STARTBY function is controlled on a job class level by a new attribute, PROMO_RATE. This job promotion rate controls how much a job can be moved up the execution queue in one STARTBY aging cycle (1 minute). The default value is 0, which disables STARTBY for a job class.

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

JES2 Job Execution Control

- z/OS 2.2 provides a new feature in JES2 similar in function to JES3 dependent job control. It introduces new concepts: the jobgroup and the jobset.
 - A JOBGROUP statement names the jobs in a jobgroup.
 - A JOBSET statement names jobs within a jobgroup that have the same dependency relationships.
 - A SCHEDULE statement associates jobs with a jobgroup.
 - A CONCURRENT statement facilitates the simultaneous execution of jobs within a jobgroup.
 - A form of dependent job control is provided with BEFORE and AFTER statements with WHEN clauses to check return codes.
 - A special job is created to do logging for a jobgroup.
 - New commands are provided to display, hold, release, cancel, purge or modify a jobgroup.

JES2 Toleration of JES3 JECL

- Application (JCL) developers want consistent processing of JCL and JECL regardless of which JES is being used.
- JES2 takes a step toward unifying JES2 and JES3 JECL.
- For each JES3 JECL statement type, a choice is given to IGNORE, WARN or FAIL as a reaction to the JECL statement.
- It can be requested to process JES3 MAIN statements.
- It's also possible to control the processing of JES2 type JECL statements with a choice to PROCESS, IGNORE, WARN or FAIL as a reaction to the JECL statement.
- Here are the meanings of the reactions:
 - PROCESS - statement is processed
 - IGNORE - statement is ignored (treated as a comment)
 - WARN - statement is processed. However, a warning message is issued to record the occurrence of a JECL statement.
 - FAIL - an error message is issued and job is failed with a JCL error

Tamper Resistant SMF Data

- A problem today with SMF data:
 - SMF records contain critical information about an enterprise and are archived for long durations. They are generally shared among various departments for a number of activities. But, there is no built in protection of the SMF data.
- Solution provided in z/OS 2.2
 - Use digital signatures to detect a change, addition/removal of an SMF record from a group of records.
- The SMF data is signed on the way to System Logger
 - As each record is written to the logstream it is hashed. A running hash is maintained for each type/subtype. Periodically, the hash is digitally signed and a signature record is recorded to the logstream. On the global interval a signature is created for all data hashed during the interval and recorded to the logstream.
 - IFASMF DL understands signature records and will optionally move them with the records of an associated SMF type/subtype
 - IFASMF DP can verify a set of SMF records has not been tampered with when signature records are available.

Complete your session evaluations online at www.SHARE.org/Orlando-Eval