z/OS 1.13 Sysprog Goody Bag

Bob Rogers
IBM Corporation

August 11, 2011
Session Number 9723

Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

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.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.
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.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
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.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective 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.

z/OS 1.13 Preview*

A smarter operating system

Improving Usability and Skills
z/OSMF Software Deployment and Storage Management applications, User-level mount command for z/OS UNIX System Services, Automatic UCB updates, SDSF Sysplex functions to work without MQ, Catalog parmlib member, Better O/C/EOV Messages, Health Checks, …

Integrating new Applications and Supporting Industry and Open Standards
Java/COBOL interoperability, Improved Support for unnamed sections, ISPF Edit Macros, Subsystem and Unauthorized XTIOT support, dbx hookless debug, DFSORT improvements, Job level return codes, …

Scalability & Performance
Fully-shared zFS in a sysplex, RMODE 64 extensions, IFASMFDL improvements, 500K+ aliases per user catalog, Larger VVDSs, FREE=EOV, FTP support for large format data sets and EAS,…

Enhancing Security
RRSF over TCP/IP, LDAP improvements, SAF security for z/OSMF, NAS address checking and encryption negotiation, New restricted QNAMEs, PKI support for DB2 backstore, ICSF support for new HMACs, FTP & TN3270 password phrase support, …

Improving Availability
Warn before TIOT exhaustion, CMDS enhancements, Parallel FTP for dump xfer, PFA ENQ tracking, RTD improvements, zFS Refresh, DADSM Dynamic Exits, JES3 dynamic spool addition, Better channel recovery, More ASID reuse, …

Self Managing Capabilities
WLM and RMF to provide response time distribution for all goals, DFSMShsm Journal Backup and space management improvements, …

Extending the Network
IDS IPv6 support, NAT Traversal for IKEV2, NMI extensions, More VLANs per OSA port, more 64-bit TCP/IP, EE improvements, …

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

z/OS Release 13

- Message Flood Automation Enhancements
- CMDS FORCE
- LLA use of suffix on restart
- Program Execution above 2GB
- QSORT for Metal C
- IFASMFDL stop reading before end of logstream
- Allow Archive or Delete of an entire SMF logstream
- 64-bit Subspaces
- Instream data in procs and includes
- Job Return Code
- Suspend a job at next step
- Spin Any (JES 2)
- FREEVOL=EOV
- IEBCOPY Performance Enhancement for PDSs
- z/OS Problem Documentation Upload Utility

Message Flood Automation Enhancements [1]

- Allow a MSGFLDxx PARMLIB member to be read much earlier during system initialization, obviating the need for the built-in table that requires a code change to alter.

- Allow up to 1024 SPECIFIC messages to be specified. This is an increase from 50 (of which 35 are used by GDPS and HyperSwap).

- Enhanced Address Space handling
  - Allow up to 128 address spaces to be tracked and acted on simultaneously.
  - Track address spaces by both jobname and ASID.
  - Allow the ASID to be symbolically substituted into the CMD action command text, allowing it to be directed to a specific ASID.

Message Flood Automation Enhancements [2]

- Allow up to 64 REGULAR JOB statements and up to 64 ACTION JOB statements to be specified in a MSGFLDxx PARMLIB member.

- Enhance the output of the DISPLAY MSGFLD,STATUS command to indicate when a message flood is underway, the type of flood, the jobname and ASID of the program causing the flood, and how long the flood has been underway.

- Additional assistance in setting Message Flood Automation message rate thresholds.
  - DISPLAY MSGFLD,MSGRATE enhanced to provide threshold recommendations at higher message rate percentages and to allow the operators to restrict the output to the uppermost portion of the message rate graph.

CMDS FORCE

- Cancelling executing commands (via the **CMDS ABEND** command) can cause system/sysplex outages.
- New support was provided in z/OS V1R12 that allows a command to dynamically indicate when it needs to be non-abendable
- While not a good idea to cancel a running command, if the only way to prevent an IPL is to terminate the command, CMDS FORCE may do it.
- Now need a way to terminate a command even if the command indicates that it should not be cancelled via CMDS ABEND.
- A **CMDS FORCE** command is provided in z/OS V1R13
  - To control the use of the FORCE operand, the new facilities class profile is checked: **MVS.CMDS.FORCE**

LLA use of suffix on restart

- On releases prior to z/OS R13, when LLA terminates and is automatically restarted, the previous CSVLLAxx suffix is used.

- But, if a customer stops and then restarts LLA and omits a suffix, none is used, defaulting to "only LNKLST".

- With z/OS R13, if LLA is stopped and restarted and a suffix is omited, the previous CSVLLAxx suffix is used.

- One can explicitly specify LLA=NONE can be to achieve the previous behavior of "only LNKLST".

Program Execution above 2GB

- When z/Architecture support was first introduced in z/OS, greater compatibility was maintained by not doing the work to support program execution above 2GB.

- Since then, pressure has been slowly building to provide this support

- The first task is enabling the z/OS Supervisor to deal with the full 128-bit z/Architecture PSW

- This is precisely the support that is enabled in z/OS R13

- This support can only be exploited by code that does not invoke any z/OS or other services and does not need to be loaded into memory by z/OS

- Despite these severe limitations, there are numerous cases of high frequency code which is created "on the fly" by products like DB2and DFSort, which could benefit from this capability

- Diagnostic support:
  - IPCS can already handle displaying data above 2G
  - SLIP can already handle ranges above 2G
  - GTF is extended to capture 128-bit PSW information
  - System trace is extended to capture 128-bit PSW information

QSORT for Metal C

- qsort() is an array sorting function that is part of many C standard libraries used to sort arrays.

- Metal C users have requested that qsort() be added to the Metal C Runtime Library so that it could be used within their applications.

- qsort() has been added to the Metal C Runtime Library.
  - Support for both ILP32- and LP64-compiled programs

IFASMFDL stop reading before end of log

- Prior to z/OS R13, IFASMFDL always reads until the end of the logstream regardless of the specified end date and time.
  - SMF data set recording did not have this issue since each data set contained only a fixed amount of data.
- In z/OS R13, a new options are provided to request that IFASMFDL stop reading the logstream before the end: SMARTENDPOINT and SMARTEPOVER(xxxx).

  - SMARTENDPOINT processing uses the following rules to find the logical end point in the logstream:

    – SMARTENDPOINT has been available for the DUMP option at R12 and below with APAR OA31737. At R13 SMARTENDPOINT is supported for ARCHIVE and DELETE.
    – Take the specified end time plus the SMARTEPOVER value. That is the logical smart end point time.
    – For each SID found in a logstream keep a table entry and mark that SID complete when all SIDs have hit the smart end point time.
    – Once all SIDs are accounted for, stop reading.

  - The SMARTEPOVER(xxxx) option specified a value between 0000 and 0200 (2 hrs) that controls the SMARTENDPOINT processing.
    – The default is 0200. Further guidance available in the MVS System Management Facilities (SMF)

Archive or Delete of an entire SMF logstream

- Prior to z/OS R13, it is not possible to ARCHIVE or DELETE an entire logstream. This limits the configuration options customers have when setting up SMF logstream recording.

- A method is provided in z/OS R13 to allow the ARCHIVE and DELETE operations to operate on the entire logstream.

- There are no new keywords or options for this support.

    - If the specification happens to include the entire logstream the operation is now allowed

    - Without this support if an ARCHIVE or DELETE of an entire logstream is attempted IFASMFDL will fail and an IFA832I message will be issued to the job log.

64-bit Subspaces

- Until now, 64-bit virtual has not been mapped into subspaces (which are used by CICS for storage isolation)
  - It has been necessary to issue the BSG instruction to switch to Base Space for the access virtual above 2GB
  - Performance is degraded with BSG
- CICS transactions run in subspace mode
  - Faster access to 64-bit private or shared storage is desirable
  - System services that exploit 64-bit private or shared storage cannot be called in subspace mode
- Solution in z/OS R13
  - Allow direct access to 64-bit private or shared storage in subspace mode
  - Storage isolation only for private storage ranges below 2GB (unchanged)
  - No subspace storage isolation for high virtual storage, i.e. 64-bit virtual storage ranges cannot be assigned to a subspace
  - This support is provided in z/OS V1R12 with PTF UA58696
- IARV64 system service now available in subspace mode.
  GETSTOR, DETACH , PAGEOUT / PAGEIN , DISCARDDATA, CHANGEGUARD, GETSHARED, PAGEFIX / PAGEUNFIX , PROTECT / UNPROTECT , SHAREMEMOBJ , CHANGEACCESS ,LIST , GETCOMMON
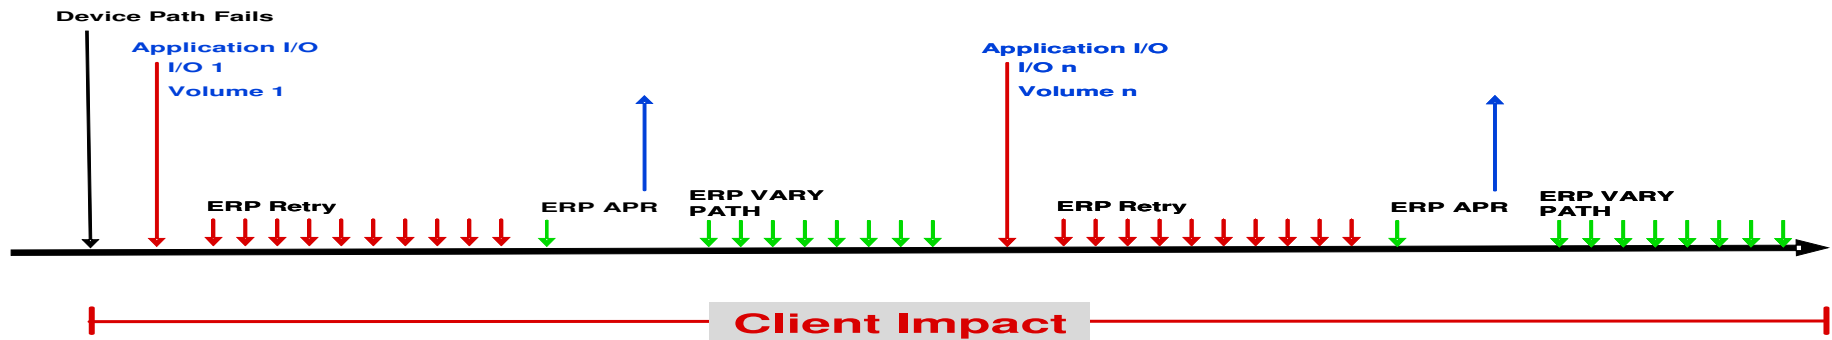 * Unauthorized callers can only issue request types in *blue*

Improved Channel Path Recovery

- Currently, when various types of path-related errors occur, IOS removes the path only from the one device on which the error occurred.

- It would be beneficial if z/OS were more pro-active when these types of path-related errors occur. This could reduce the time for z/OS to recover from path-related errors by performing recovery for the entire scope of devices impacted by the failing resource all at one time.

- In z/OS R13, IOS will track path-related errors at the Control Unit level and will, at a threshold point (# of failures in specific time interval), respond by removing the failing path from all devices in the Control Unit.

- In the IECIOSxx PARMLIB Member at IPL or for SET IOS=

  ```
  RECOVERY  PATH_SCOPE={CU|DEVICE}  default: DEVICE
            ,PATH_INTERVAL=nn        default: 10
            ,PATH_THRESHOLD=nnn      default: 10
  ```

- DISPLAY IOS,RECOVERY now displays the following new information:

  ```
  IOS103I hh.mm.ss RECOVERY OPTIONS

        PATH RECOVERY SCOPE IS BY CU

        PATH RECOVERY INTERVAL IS nn MINUTES

        PATH RECOVERY THRESHOLD is nnn ERRORS

                    - or -

        PATH RECOVERY SCOPE IS BY DEVICE
  ```
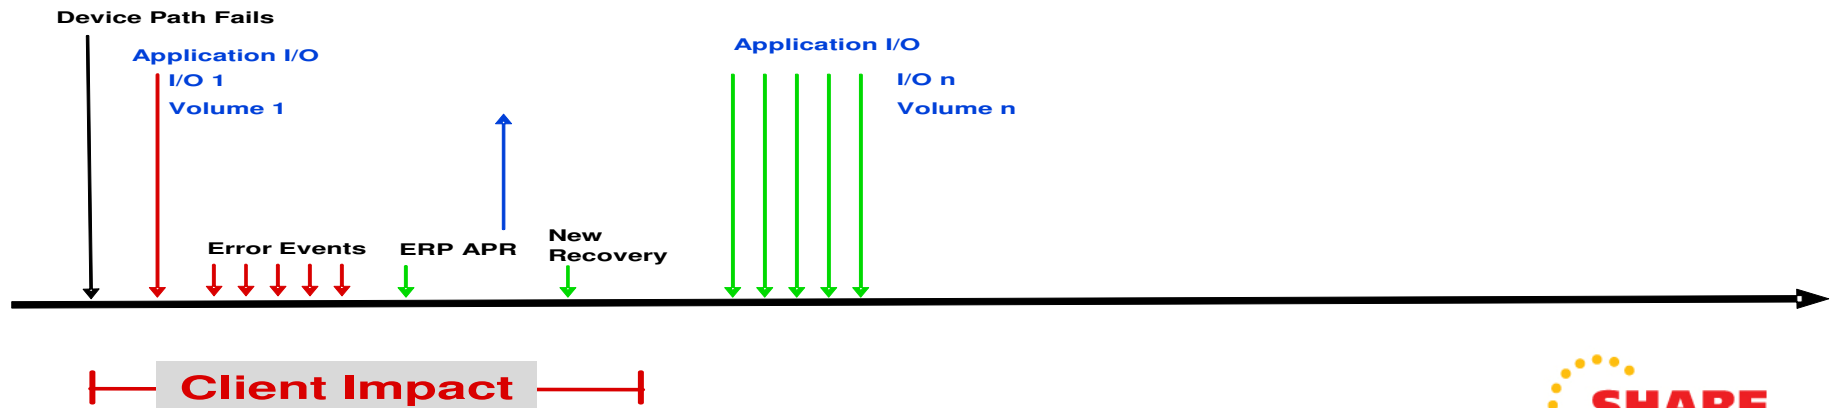
Improved Channel Path Recovery

I/O Recovery for Failing Path - Before

Device Path Fails

Application I/O
I/O 1
Volume 1

Application I/O
I/O n
Volume n

ERP Retry

ERP APR

ERP VARY
PATH

ERP Retry

ERP APR

ERP VARY
PATH

Client Impact

I/O Recovery for Failing Path - After

Device Path Fails

Application I/O
I/O 1
Volume 1

Application I/O
I/O n
Volume n

Error Events

ERP APR

New
Recovery

Client Impact

Path-related error indication

- When a path-related error occurs, it can be difficult to determine where the problem resides. Is it in the switch card attached to the channel, the link between the channel and switch or what?

- In z/OS R13, IOS issues a new message with the IOS050I and IOS051I messages issued for path-related errors to identify the component that detected the error.

IOS050I CHANNEL DETECTED ERROR ON ddddd,yy,op,stat,
      PCHID=pppp

IOS054I ddddd,pp ERRORS DETECTED BY comp, comp,…

Where *comp* is one or more of the following:

  CHANNEL, CHAN SWITCH PORT, CU SWITCH PORT, CONTROL UNIT

Instream data in procs(JES2)

- Instream data is supported in procedures and includes
  - For example:
    ```
    //PROC1   PROC
    //ASTEP     EXEC PGM=xyz
    //DD1         DD *
        This is instream data
    //                    PEND
    ```
  - Prior to this support you could not have instream data in a JCL proc
  - Support also allows instream data in include statements wherever a DD is allowed
  - This is particularly useful when parameters are provided in a file (for example, BCDBATCH)
  - Requires z/OS 1.13 and JES2 1.13 on both the converting system and the initiating system
  - Not supported for JES3 or MSTR subsystems

Job Return Code (JES2)

- Historically, there has not been a formal definition of the job return code. The defacto standard for JES2 has been to return the maximum return code of the job.
- z/OS 1.13 introduces a new JOB Card parameter - JOBRC

   //ABCDEFG JOB …,JOBRC=options

   Where options are:

   STEP,proc.stepname – return code of a specific step
   MAXRC – highest return code of all steps run
   LASTRC – return code of the last step run

- A new message in JOBLOG which give JOBRC rather than max RC
- SDSF will display the JOBRC on the job display line

Suspend a job at next step (JES2)

- ## If a job is restartable
  - A new operator command is defined to suspend the job at the next step

    $E J123,STEP
  - The job is then requeued for execution
- ## This allows for a somewhat faster non-disruptive shutdown
- ## The job must complete the step it is running.
  - This remains an issue for long running steps

Spin Any (JES 2)

- Provides the ability to specify automatic spin options in JCL

  - Eliminates the requirement to take down a long running process to release spool space

  - Provides a time, size, or command based option to spin output

- **Syntax:**

  ```
  SPIN= { NO }
       { UNALLOC }
       { (UNALLOC,'hh:mm') }
       { (UNALLOC,'+hh:mm') }
       { (UNALLOC,nnn [K|M]) }
       { (UNALLOC,NOCMND) }
       { (UNALLOC,CMNDONLY) }
  ```

FREEVOL=EOV

- Previously, the volumes of multivolume tape data set remained allocated until after all volumes were processed.

- z/OS R13 supports a new FREEVOL=EOV keyword on the DD statement to enable earlier freeing.

  - This allows you to specify that each volume of a multivolume tape data set that is being read be made available for other processing once the processing for that volume is finished

  - It is intended to allow overlapped processing for multivolume data sets, which can speed batch processing.

IEBCOPY Performance for PDSs

- Performance is improved when copying one partitioned data set (PDS) to another using IEBCOPY

- IEBCOPY is planned to exploit 31-bit storage for track buffers

- The requirement for APF authorization is planned to be removed

Problem Documentation Upload Utility

- The Problem Documentation Upload Utility (PDUU) is a utility that facilitates sending large amounts of documentation in an efficient manner
  - Previously, PDUU was an independent tool available for downloading from a web site
  - z/OS R13 has included this in the z/OS operation system as a z/OS Service Aid
    - The PDUU utility program name is AMAPDUPL with MTFTPS as an alias entry point for compatibility
    - AMAPDUPL resides in SYS1.MIGLIB, so a STEPLIB is not necessary
- Data sets are compressed, optionally encrypted, and sent in parallel using multiple, simultaneous FTP sessions
  - There can be up to 20 FTP sessions running simultaneously
  - There are two work files for each FTP session which are dynamically allocated and can range in size from 1MB to 9,999MB
- PDUU always compresses the input data before it is written to the work data sets
  - No need to use an additional tool such as AMATERSE or TRSMAIN to compress the input data set before using the PDUU
- Encryption using the 192-bit triple Data Encryption Standard can be requested by using CIPHER_KEY keyword
  - This uses a z/Architecture instruction that is built-in on all processors starting with the z990/z890

SSRBs above the 2GB Bar

- Virtual storage below 2GB is becoming ever more constrained

- Increased use of preemptible SRBs puts even more pressure on virtual below 2GB

- In z/OS R13, the SSRB control block is restructured and most of the fields are moved to a new control block, the SSRX, in common storage above 2GB.

- Although no interfaces are changed, there may be programs which reference the SSRB and need to change due to movement of the field or fields.

Improved Predictive Failure Analysis

- Predictive Failure Analysis (PFA) is enhanced with additional checks and integration with Runtime Diagnostics
    - PFA_ENQUEUE_REQUEST_RATE Check
        - Detects a damaged address space or damaged system by comparing the number of enqueue requests per CPU second to the rate expected.
    - PFA_JES_SPOOL_USAGE Check
        - Detects abnormalities in persistent jobs by modelling the change in their JES2 spool usage (i.e., the change in the number of track groups used from one collection to the next)
    - Detect if certain metrics are too low by using predictive technology along with the results of Runtime Diagnostics

PFA for "too low" exceptions

- Integration with Runtime Diagnostics to Detect Too Low
  - Detects a *damaged or hung address space or system* based on rates being too low
  - When PFA detects an abnormally low condition,
    - Runtime Diagnostics is executed
    - If the results of Runtime Diagnostics indicate a problem,
      - *the exception is issued*
      - *the PFA prediction report includes the Runtime Diagnostics output*
  - Conditions detected:
    - AIRH190E – Tracked job enqueue request rate lower than expected
    - AIRH211E – Total system enqueue request rate lower than expected
    - AIRH206E – Tracked job message arrival rate lower than expected
    - AIRH207E – Other persistent job message arrival rate lower than expected
    - AIRH153E – Total system message arrival rate lower than expected
    - AIRH208E – Tracked job SMF arrival rate lower than expected
    - AIRH209E – Other persistent job SMF arrival rate lower than expected
    - AIRH175E – Total system message SMF arrival rate lower than expected

PFA for "too low" exceptions

- "Too low" exception message sent as WTO

- Runtime Diagnostics output included in PFA report

- Prediction report and result message available in SDSF (sdsf.ck)

- Prediction report relevant to comparison category causing exception

```
                    Message Arrival Rate Prediction Report
          (Heading information intentionally omitted.)

     Persistent address spaces with low rates:
                                            Predicted Message
                              Message          Arrival Rate
          Job                 Arrival
          Name      ASID       Rate      1 Hour      24 Hour       7 Day

          JOBS4     001F       1.17       23.88        22.82       15.82
          JOBS5     002D       2.01        8.34        11.11       12.11

     Runtime Diagnostics Output:

     Runtime Diagnostics detected a problem in job:    JOBS4
        EVENT 06: HIGH  - HIGHCPU - SYSTEM: SY1 2009/06/12   - 13:28:46
        ASID CPU RATE: 96% ASID: 0027 JOBNAME:    JOBS4
        STEPNAME: DAVIDZ PROCSTEP:   DAVIDZ JOBID: STC00042 USERID:
     ++++++++
        JOBSTART: 2009/06/12   - 13:28:35
     Error:
        ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
     Action:
        USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
     ----------------------------------------------------  -----------------
     -----
        EVENT 07: HIGH  -  LOOP - SYSTEM: SY1 2009/06/12   - 13:28:46
        ASID: 0027 JOBNAME:    JOBS4 TCB: 004E6850
        STEPNAME: DAVIDZ PROCSTEP: DAVIDZ JOBID: STC00042 USERID:
     ++++++++
        JOBSTART: 2009/06/12   - 13:28:35
     Error:
        ADDRESS SPACE A PPEARS TO BE IN A LOOP.
     Action:
        USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.

          (Additional output intentionally omitted.)
```

z/OSMF Improvements

- **z/OSMF Capacity Provisioning Manager application**
  - Designed for easy monitoring of CPM status

- **z/OSMF Configuration Assistant for Communications Server**
  - Multiple release configuration support (both R12 and R13 systems)
  - Sysplex-wide policy definitions
  - IP address discovery from stacks

- **Expanded SAF-based security for z/OSMF user authorization and roles**
  - In addition to current z/OSMF security
  - Intended to be used in place of the current z/OSMF repository-based authorization support

- **Consolidated workload monitoring**
  - With RMF and z/OSMF you can monitor z/OS, AIX®, and Linux workloads
  - Monitor across zHybrid ensembles and other network-accessible AIX and Linux systems from within z/OSMF

- **z/OSMF support for application linking**
  - Allow z/OSMF applications to link directly to others via URL
  - Both in-context linking and simple linking
 - Intended to make it simpler to navigate across apps…such as…

z/OSMF and ISPF

- Enables system programmer to perform tasks from one interface.
  - Also makes ISPF applications URL Web-accessible for linking and launching from other applications

- **Up to 4 panes**
  - Panes can be sized. Each pane can have multiple ISPF sessions, tabs can be moved between panes
  - Shows ISPF Menu bar, Command line, Function keys
  - Open and cascade multiple windows
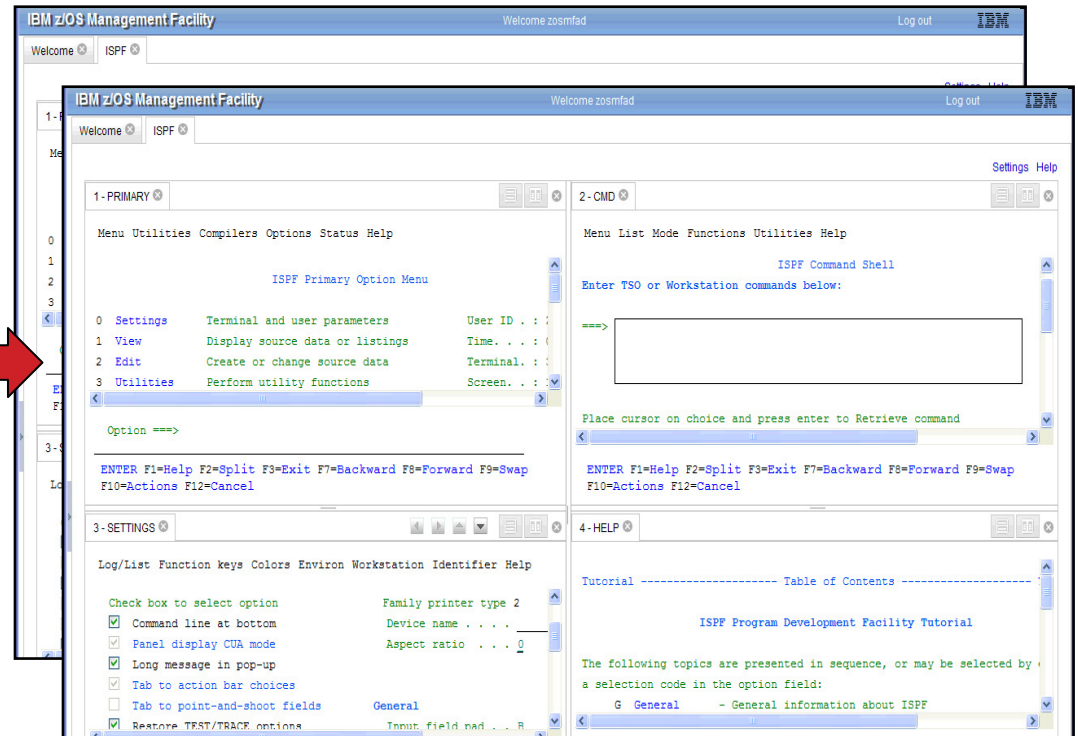
- **Also available in tabbed format**
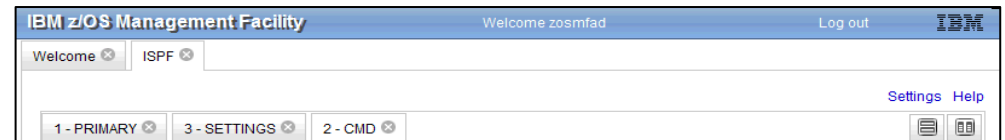
New programmatic interface for z/OS

- A RESTful web service (also called a RESTful web API, or REST API) is a simple web-based publishing service implemented using HTTP.
- z/OSMF R13 Introduces a new REST API (HTTP(s)) interface to z/OS for submitting and accessing batch job information.
- This interface is an open API which can be driven locally on a z/OS system or driven from a remote system (The remote system need only support the HTTP(s) protocol)
- REST API web services can be used by: web applications (javascript/AJAX, Flex(Flash), etc) and other web service clients, such as Java, PHP, Perl, etc
- The REST API web service will connect to both JES2 and JES3, as well as select secondary subsystems
- There is no GUI interface for the REST API.
  - This item provides the customer a programmatic interface for web applications via z/OSMF

JOBS REST API for z/OS

- Today's options:
  - Allocate and open internal reader
  - TSO/ISPF submit,
  - FTP "interface-level2"
  - Java z/OS submit interface

- Complex programming
- Security protocol limitations

- New RESTful HTTPs based API
  - Secure
  - Firewall friendly
  - Text like

- Any web-based, Java, PHP, Perl application, etc. supporting HTTP

- Improve access to z/OS batch
  - Increase client coverage
  - Add a non-controversial remote interface
  - Align with Web technologies
  - Make z/OS a little more approachable

- z/OSMF R13
- z/OS JES2
- z/OS JES3

Submit JCL, get status, retrieve output files, change jobclass, cancel job, purge job

- Invoke the power of z/OS batch from virtually any non-z/OS application