


z/OS Problem Determination Update

SHARE Session 8696: March 1, 2011

Karla Arndt
System z - Core Technologies
Rochester, Minnesota
kka@us.ibm.com
507-253-3606

Anuja Deedwaniya
System z - Application Enablement
Poughkeepsie, New York
anujad@us.ibm.com
845-435-4160



SHARE
Technology • Connections • Results

Permission is granted to SHARE Inc. to publish this presentation paper in the SHARE Inc. proceedings; IBM retains the right to distribute copies of this presentation to whomever it chooses.

Trademarks

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

DB2*	System REXX
DB2 Connect	System z
DB2 Universal Database	Tivoli*
e-business logo	VM/ESA*
GDPS*	WebSphere*
Geographically Dispersed Parallel Sysplex	z/OS*
HyperSwap	z/VM*
IBM*	zSeries*
IBM eServer	
IBM logo*	
Parallel Sysplex*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a registered trademark of the Intel Corporation in the United States, other countries or both.

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

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.

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

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

* 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.

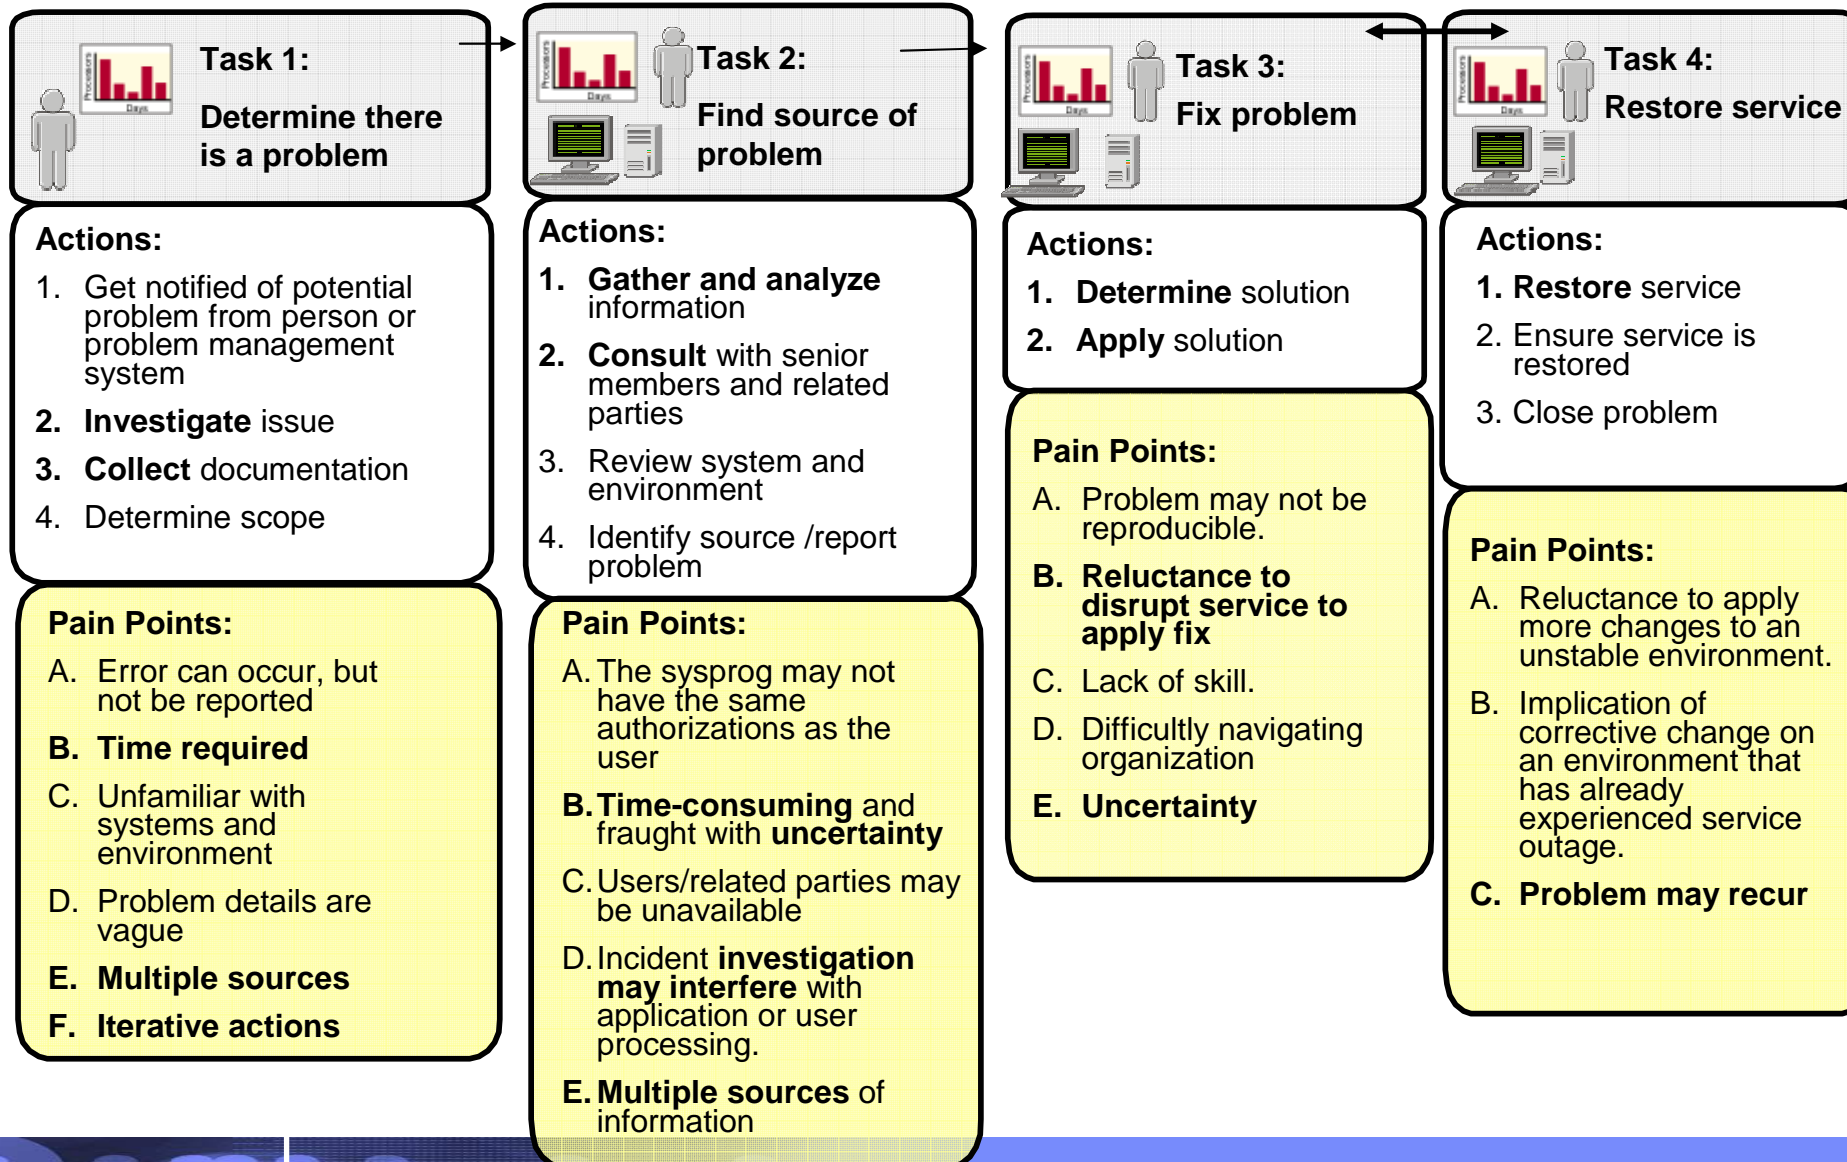
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products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

Problem Determination Update - Agenda

- Requirements for Problem Determination Improvements
- z/OS Problem Determination Simplification
 - Problem identification: Runtime Diagnostics
 - Soft Failure detection: Predictive Failure Analysis
 - Problem Data Management: z/OSMF Incident Log
- Documentation
- Related SHARE sessions

z/OS Problem Determination Scenario



Conclusions: Problem Determination in a complex environment

Installation Pain Points

Risk to the business <ul style="list-style-type: none"> • The impact of the symptoms • Risk of recurrence • Impact in getting system stabilized • Mean time to recovery too long
Complexity of performing the task
Troubleshooting a live system and recovering from an apparent failure
Data collection very time-consuming
Significant skill level needed to analyze problems, interact with IBM and ISVs to obtain additional diagnostic info





Requirement Areas

Detect “sick, but not dead” event BEFORE it causes problems
Diagnose the cause in real time to allow operations to mitigate event inquiries
Manage / capture data to determine cause of problem <ul style="list-style-type: none"> • Allow problem to be fixed to prevent recurrence

PD Simplification

Simplifying tasks

- z/OS Service Aids**
- Dumps
 - Logs
 - Traces
 - Tools



Soft Failure Detection: Predictive Failure Analysis (PFA)

- | | |
|------------------------|----------------------|
| Common Storage Usage | SMF Arrival Rate |
| Frames and Slots Usage | JES2 Spool Usage |
| Logrec Arrival Rate | Enqueue Request Rate |
| Message Arrival Rate | |

System Symptom Analysis: Runtime Diagnostics

- Analyze a sick system upon request via operator command
- Component errors
- Address space behavior (ENQ, CPU usage, Loops, etc.)

Problem Data Management: Incident Log

- Improve FFDC for system-detected problems
- Diagnostic data “snapshots” for transient data
- FTP Incident with all diagnostic data
- User interface to display summary, detail, drive actions

Soft Failures: What is a soft failure?

“Sick, but not dead” or Soft failures





- 80% of business impact, but only about 20% of the problems
- Long duration
- Infrequent
- Unique
- Any area of software or hardware
- Cause creeping failures
- Hard to determine how to isolate, diagnose
- Hard to determine how to recover
- Hard for software to detect internally
- Probabilistic, not deterministic

PD Simplification

Simplifying tasks

- z/OS Service Aids**
- Dumps
 - Logs
 - Traces
 - Tools



Soft Failure Detection: Predictive Failure Analysis (PFA)

- | | |
|------------------------|----------------------|
| Common Storage Usage | SMF Arrival Rate |
| Frames and Slots Usage | JES2 Spool Usage |
| Logrec Arrival Rate | Enqueue Request Rate |
| Message Arrival Rate | |

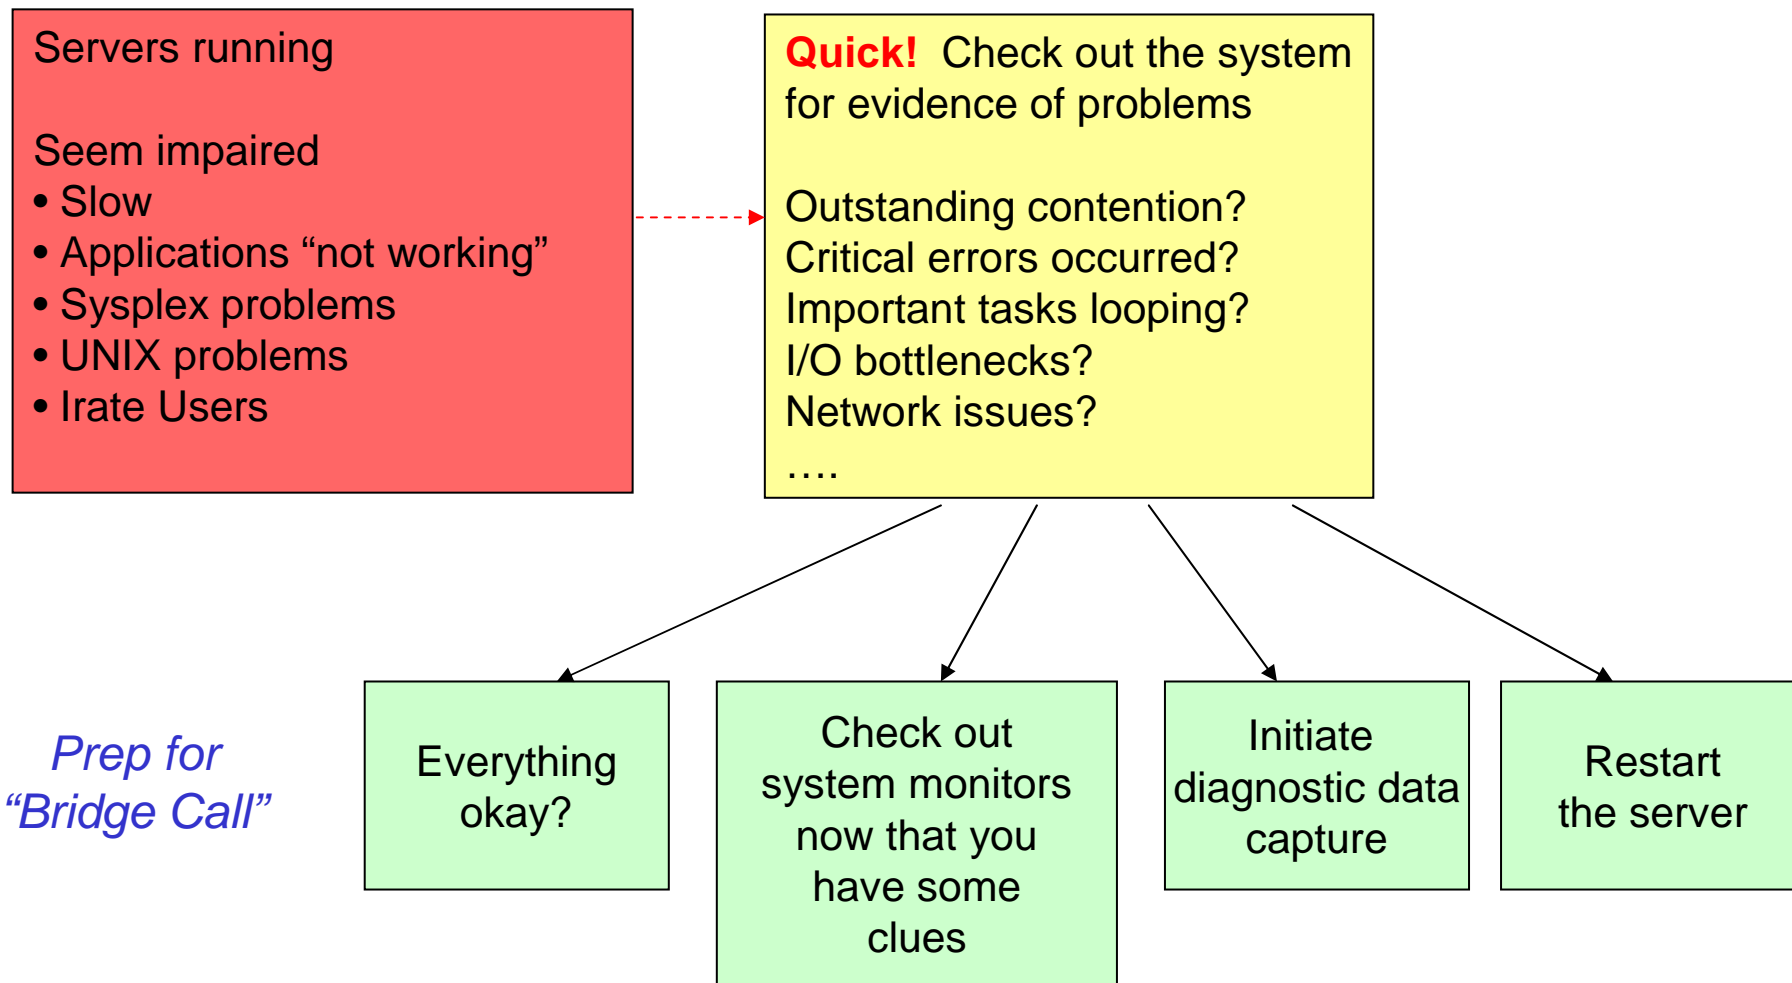
System Symptom Analysis: Runtime Diagnostics

- Analyze a sick system upon request (operator command)
- Component errors
- Address space behavior (ENQ, CPU usage, Loops, etc.)

Problem Data Management: Incident Log

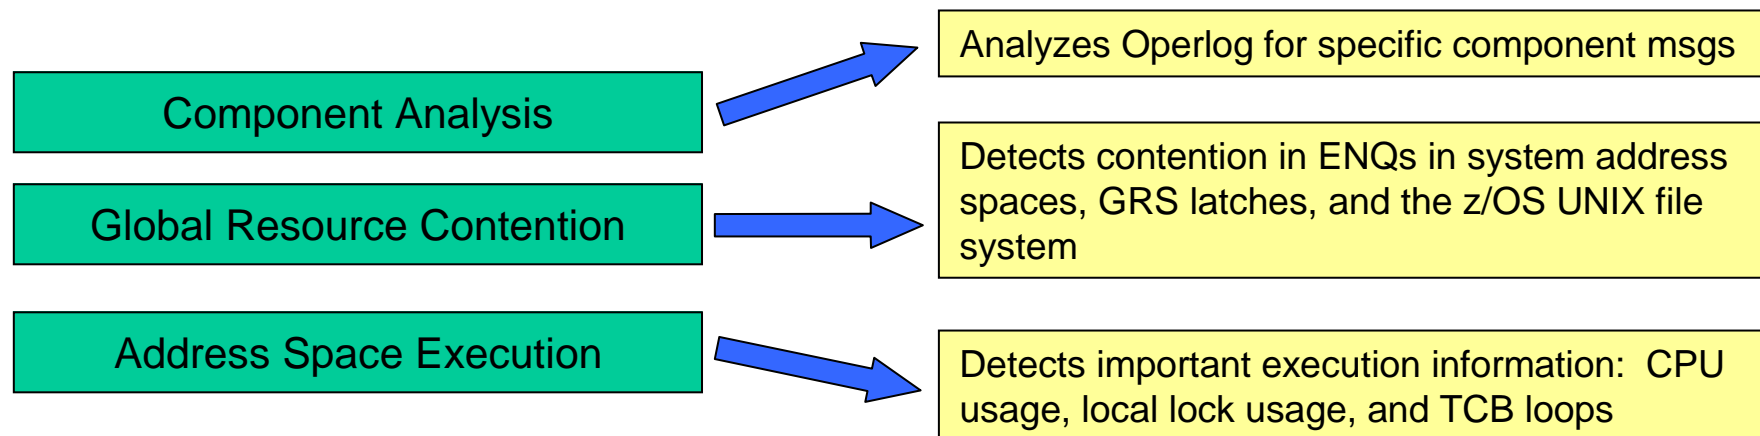
- Improve FFDC for system-detected problems
- Diagnostic data “snapshots” for transient data
- FTP Incident with all diagnostic data
- User interface to display summary, detail, drive actions

Problem Scenario ... is this typical?



Runtime Diagnostics

- Analyzes a “sick, but not dead” system in a timely manner
- Performs analysis similar to a very experienced system programmer
 - But faster – goal of 60 seconds or less
 - More comprehensive
 - Looks for specific evidence of “soft failures”
 - Provides suggested next steps
- **Runtime Diagnostics**
 - Is not automation or a monitor
 - Takes no corrective action
 - Has no background processing and minimal dependencies on system services



Runtime Diagnostics Benefits

- Reduces the skill level needed by a system programmer for investigating soft failures
 - Provides timely, comprehensive analysis at a critical time period
 - *Also great productivity aid for experienced system programmers!*
- Allows you to *quickly discover next actions* to take such as
 - which jobs to cancel
 - what to investigate further
 - Such as classes of resources or a single address space using a monitor like RMF or Tivoli Omegamon
- Use Runtime Diagnostics ...
 - when the help desk or operations reports a problem on the system
 - to get ready for the “bridge call”
 - when PFA detects abnormal behavior

Runtime Diagnostics Invocation

- z/OS 1.12 – Started task – “Run” the analysis via a START command
 - START HZR,SUB=MSTR
 - Invokes HZR PROC
 - Will only run on R12 system, but other systems in the Sysplex do not need to be R12
 - Can override HZROUT to specify a data set, for example:
 - //HZROUT DD DISP=SHR,DSN=MY.DATA
 - START HZR,SUB=MSTR,DSN=MY.DATA,DISP=SHR
- z/OS 1.13 – Address space – started with the START command above
 - Address space needs to be available for PFA integration
 - Recommend to start address space at IPL
 - “Run” the analysis via a MODIFY command
 - f hzr,analyze
 - Migration Action: If you used Runtime Diagnostics in z/OS 1.12, ensure you update the hzrproc to point to PGM=HZRINIT instead of PGM=HZRMAIN.

Runtime Diagnostics Invocation (continued)

- The output of Runtime Diagnostics is a multi-line WTO
 - Can also be directed to a sequential dataset using HZROUT DD
- **SYSNAME** option targets system other than HOME
 - Operlog and ENQ analysis are done for specified system
 - **Operlog is suggested to allow message analysis**
 - Example: `OPTIONS=(SYSNAME=SYS2)`
 - z/OS 1.12 – SYSNAME option on START command
 - z/OS 1.13 – SYSNAME option on MODIFY command
- **DEBUG** option for use under IBM Service guidance
 - Takes a dump to help debug analysis
 - Options specific to type of analysis and when found or not found
 - Example: `OPTIONS=(DEBUG=(LOOP,NOENQ))`
 - z/OS 1.12 – DEBUG option on START command
 - z/OS 1.13 – DEBUG option on MODIFY command

Runtime Diagnostics Output

- Success

```
f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT 974
SUMMARY: SUCCESS
REQ: 001 TARGET SYSTEM: SY1      HOME: SY1      2010/12/21 - 11:30:57
INTERVAL:  60 MINUTES
EVENTS:
  FOUND: 05 - PRIORITIES: HIGH:05  MED:00  LOW:00
  TYPES: CF:04
  TYPES: HIGHCPU:01
```

- Qualified Success – Example of Operlog not connecting

```
f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT 751
SUMMARY: QUALIFIED SUCCESS - SOME PROCESSING FAILED
REQ: 001 TARGET SYSTEM: SY1      HOME: SY1      2010/12/21 - 11:25:55
INTERVAL:  60 MINUTES
EVENTS:
  FOUND: 02 - PRIORITIES: HIGH:02  MED:00  LOW:00
  TYPES: HIGHCPU:01
  TYPES: LOOP:01
PROCESSING FAILURES:
  OPERLOG....IXGCONN REQ=CONNECT ERROR.....RC=00000008 RS=0000080B
```

Runtime Diagnostics: Critical Message Analysis

- Component-specific, critical messages in OPERLOG
 - “Needles in a haystack”
 - Looks one hour back, if available
 - For some messages, additional analysis done
 - Groups related messages into a single event
 - Weeds out shortage and relieved critical messages
 - In some cases, will only show last message if a critical message for the same resource name is repeated, say every 10 minutes
 - Message summary found listed in Runtime Diagnostics output

```
EVENT 02: HIGH - CF          - SYSTEM: SY1      2011/02/15 - 14:47:03
IXC585E STRUCTURE LIST01 IN COUPLING FACILITY TESTCFN,
PHYSICAL STRUCTURE VERSION C7565A8D E48F6410,
IS AT OR ABOVE STRUCTURE FULL MONITORING THRESHOLD OF 80%.
ENTRIES:  IN-USE:           491 TOTAL:           583,   84% FULL
ELEMENTS: IN-USE:           508 TOTAL:          1167,   43% FULL
        ERROR: INDICATED STRUCTURE IS APPROACHING FULL MONITORING THRESHOLD.
        ACTION: D XCF,STR,STRNAME=strname TO GET STRUCTURE INFORMATION.
        ACTION: INCREASE STRUCTURE SIZE OR TAKE ACTION AGAINST APPLICATION.
```

Runtime Diagnostics: ENQ Contention Checking

- Looks for a system address space that is an ENQ “waiter” for over 5 seconds
- Lists both waiter and blocker
- Equivalent to D GRS,AN,WAITER

```
f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT 581
SUMMARY: SUCCESS
REQ: 004 TARGET SYSTEM: SY1      HOME: SY1      2010/12/21 - 13:51:32
INTERVAL: 60 MINUTES
EVENTS:
FOUND: 04 - PRIORITIES: HIGH:04  MED:00  LOW:00
TYPES: HIGHCPU:01
TYPES: LOOP:01  ENQ:01  LOCK:01
-----
EVENT 01: HIGH - ENQ      - SYSTEM: SY1      2010/12/21 - 13:51:32
ENQ WAITER  - ASID:0038 - JOBNAME:IBMUSER2 - SYSTEM:SY1
ENQ BLOCKER - ASID:002F - JOBNAME:IBMUSER1 - SYSTEM:SY1
QNAME: TESTENQ
RNAME: TESTOFAVERYVERYVERYVERYLOOOOOOOOOOOOOOOOOOOOOOONGRNAMEI234567...
ERROR: ADDRESS SPACES MIGHT BE IN ENQ CONTENTION.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE BLOCKING JOBS AND
ACTION: ASIDS.
```


Runtime Diagnostics: CPU Analysis

- Takes two quick samples over 1 second interval
- Any task using > 95% of a single CPU is considered a potential problem
- The usage reported might be > 100% if an address space has multiple TCBs and several are using a high percentage of the capacity of a CPU

```
f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT 581
SUMMARY: SUCCESS
REQ: 004 TARGET SYSTEM: SY1      HOME: SY1      2010/12/21 - 13:51:32
INTERVAL: 60 MINUTES
EVENTS:
FOUND: 04 - PRIORITIES: HIGH:04  MED:00  LOW:00
TYPES: HIGHCPU:01
TYPES: LOOP:01 ENQ:01 LOCK:01
-----
EVENT 02: HIGH - HIGHCPU          - SYSTEM: SY1      2010/12/21 - 13:51:33
ASID CPU RATE:99%      ASID:002E      JOBNAME:IBMUSERX
STEPNAME:STEP1      PROCSTEP:      JOBID:JOB00045  USERID:IBMUSER
JOBSTART:2010/12/21 - 11:22:51
ERROR: ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MIGHT BE LOOPING.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
```

Runtime Diagnostics: Local Lock Suspension

- Lists any address space where its local lock suspension time is over 50%

```
f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT 581
SUMMARY: SUCCESS
REQ: 004 TARGET SYSTEM: SY1      HOME: SY1      2010/12/21 - 13:51:32
INTERVAL: 60 MINUTES
EVENTS:
  FOUND: 04 - PRIORITIES: HIGH:04  MED:00  LOW:00
  TYPES: HIGHCPU:01
  TYPES: LOOP:01 ENQ:01 LOCK:01
-----
EVENT 04: HIGH - LOCK              - SYSTEM: SY1      2010/12/21 - 13:51:33
HIGH LOCAL LOCK SUSPENSION RATE - ASID:000A JOBNAME:WLM
STEPNAME:WLM      PROCSTEP:IEFPROC  JOBID:+++++++  USERID:+++++++
JOBSTART:2010/12/21 - 11:15:08
  ERROR: ADDRESS SPACE HAS HIGH LOCAL LOCK SUSPENSION RATE.
  ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
```

Runtime Diagnostics: Loop Detection

- Investigates all tasks in all address spaces looking for TCB loops
 - Takes a snapshot of the system trace
 - Looks for consistent, repetitive activity that typically indicates a loop
- When both HIGHCPU and LOOP events occur for the same job, there is a high probability that the task in the job is in a loop.
- Normal, corrective action is to cancel the job.

```
f hzr,analyze
HZR0200I RUNTIME DIAGNOSTICS RESULT 581
SUMMARY: SUCCESS
REQ: 004 TARGET SYSTEM: SY1      HOME: SY1      2010/12/21 - 13:51:32
INTERVAL: 60 MINUTES
EVENTS:
  FOUND: 04 - PRIORITIES: HIGH:04 MED:00 LOW:00
  TYPES: HIGHCPU:01
  TYPES: LOOP:01 ENQ:01 LOCK:01
-----
EVENT 02: HIGH - HIGHCPU          - SYSTEM: SY1      2010/12/21 - 13:51:33
ASID CPU RATE:99%      ASID:002E  JOBNAME:IBMUSERX
STEPNAME:STEP1      PROCSTEP:          JOBID:JOB00045 USERID:IBMUSER
JOBSTART:2010/12/21 - 11:22:51
  ERROR: ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MIGHT BE LOOPING.
  ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
EVENT 03: HIGH - LOOP            - SYSTEM: SY1      2010/12/21 - 13:51:14
ASID:002E  JOBNAME:IBMUSERX  TCB:004FF1C0
STEPNAME:STEP1      PROCSTEP:          JOBID:JOB00045 USERID:IBMUSER
JOBSTART:2010/12/21 11:22:51
  ERROR: ADDRESS SPACE MIGHT BE IN A LOOP.
  ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
```

Runtime Diagnostics: z/OS UNIX Latch Contention

- New in z/OS 1.13
- If z/OS UNIX latch contention or waiting threads exist for > 5 minutes in z/OS UNIX, a Runtime Diagnostics OMVS event is created.
- Normal action is to issue D OMVS,W,A to get the ASID and job names of the waiters

```
F HZR,ANALYZE
HZR0200I RUNTIME DIAGNOSTICS RESULT 692
SUMMARY: SUCCESS
REQ: 009 TARGET SYSTEM: SY1      HOME: SY1      2010/12/21 - 14:24:29
INTERVAL: 60 MINUTES
EVENTS:
  FOUND: 02 - PRIORITIES: HIGH:02  MED:00  LOW:00
  TYPES: OMVS:01
  TYPES: LOCK:01
-----
EVENT 01: HIGH - OMVS          - SYSTEM: SY1      2010/12/21 - 14:24:29
ASID:000E - JOBNAME:OMVS
MOUNT LATCH WAITERS: 1
FILE SYSTEM LATCH WAITERS: 0
XSYS AND OTHER THREADS WAITING FOR z/OS UNIX: 1
  ERROR: z/OS UNIX MIGHT HAVE FILE SYSTEM LATCH CONTENTION.
  ACTION: D OMVS,W,A TO INVESTIGATE z/OS UNIX FILE SYSTEM LATCH
  ACTION: CONTENTION, ACTIVITY AND WAITING THREADS. USE YOUR SOFTWARE
  ACTION: MONITORS TO INVESTIGATE BLOCKING JOBS AND ASIDS.
```

Runtime Diagnostics: GRS Latch Contention



- New in z/OS 1.13
- Obtains latch contention information from GRS
- Omits z/OS UNIX file system latch contention
- Returns the longest waiter for each latch set

```
F HZR,ANALYZE
HZR0200I RUNTIME DIAGNOSTICS RESULT 692
SUMMARY: SUCCESS
REQ: 002 TARGET SYSTEM: SY1      HOME: SY1      2010/12/21 - 14:32:01
INTERVAL: 60 MINUTES
EVENTS:
  FOUND: 02 - PRIORITIES: HIGH:02  MED:00  LOW:00
  TYPES: LATCH:02
-----
EVENT 01: HIGH - LATCH          - SYSTEM: SY1      2010/12/21 - 14:32:01
LATCH SET NAME: SYSTEST.LATCH_TESTSET
LATCH NUMBER:3                CASID:0039  CJOBNAME:TSTLATCH
TOP WAITER - ASID:0039 - JOBNAME:TSTLATCH - TCB/WEB:004E2A70
TOP BLOCKER- ASID:0039 - JOBNAME:TSTLATCH - TCB/WEB:004FF028
ERROR: ADDRESS SPACES MIGHT BE IN LATCH CONTENTION.
ACTION: D GRS,AN,LATCH,DEP,CASID=0039,LAT=(SYSTEST.L*,3),DET
ACTION: TO ANALYZE THE LATCH DEPENDENCIES. USE YOUR SOFTWARE
ACTION: MONITORS TO INVESTIGATE BLOCKING JOBS AND ASIDS.
```

PD Simplification

Simplifying tasks

- z/OS Service Aids**
- Dumps
 - Logs
 - Traces
 - Tools



Soft Failure Detection: Predictive Failure Analysis (PFA)

- | | |
|------------------------|----------------------|
| Common Storage Usage | SMF Arrival Rate |
| Frames and Slots Usage | JES2 Spool Usage |
| Logrec Arrival Rate | Enqueue Request Rate |
| Message Arrival Rate | |

System Symptom Analysis: Runtime Diagnostics


- Analyze a sick system upon request (operator command)
- Component errors
- Address space behavior (ENQ, CPU usage, Loops, etc.)

Problem Data Management: Incident Log

- Improve FFDC for system-detected problems
- Diagnostic data “snapshots” for transient data
- FTP Incident with all diagnostic data
- User interface to display summary, detail, drive actions

How PFA detects soft failures

- Causes of “sick, but not dead”
 - **Damaged systems**
 - Recurring or recursive errors caused by software defects anywhere in the software stack
 - Serialization
 - Priority inversion
 - Classic deadlocks
 - Owner gone
 - **Resource exhaustion**
 - Physical resources
 - Software resources
 - Indeterminate or unexpected states
- Predictive failure analysis uses
 - *Historical data*
 - *Machine learning and mathematical modeling*to detect abnormal behavior and the potential causes of this abnormal behavior
- Objective
 - Convert “sick, but not dead” to a correctable incident



Predictive Failure Analysis Enhancements by Release

- **z/OS 1.10 (SPE)**

- Common storage usage check
 - CSA + SQA below the line
 - eCSA + eSQA above the line
- LOGREC arrival rate check
 - Key 0
 - Keys 1 to 7
 - Keys 8 to 15

- **z/OS 1.11**

- Frames and slots usage check
 - Persistent address spaces
- Message arrival rate (WTO/WTOR) check
 - Chatty, persistent address spaces
 - Non-chatty, persistent address spaces
 - Non-persistent address spaces
 - Total system

- **z/OS 1.12**

- SMF arrival rate check
 - Same categories as message arrival rate check
- Modeling improvements
 - More granular Common storage usage check
 - CSA, SQA, ECSA, ESQA, CSA+SQA, ECSA+ESQA
 - Supervised learning (excluded jobs list)
 - Dynamic modeling
- Performance and serviceability enhancements

- **z/OS 1.13**

- JES spool usage check
 - Persistent address spaces
 - JES2 only
- Enqueue request rate check
 - Chatty, persistent address spaces
 - Total system
- Integration with Runtime Diagnostics to detect rates that are too low
 - Message arrival rate check, SMF arrival rate check, and Enqueue request rate check.




PFA and Runtime Diagnostics Summary

- PFA detects a soft failure before it impacts your business
- **Runtime Diagnostics** helps you analyze a soft failure, diagnose the problem, and take corrective action in a timely manner
- One main source of information for both: *z/OS Problem Management G325-2564-XX*
 - PFA IEA presentations
 - http://publib.boulder.ibm.com/infocenter/ieduasst/stgv1r0/index.jsp?topic=/com.ibm.iea.zos/zos/1.11/Availability/V1R11_PFA/player.html
 - http://publib.boulder.ibm.com/infocenter/ieduasst/stgv1r0/index.jsp?topic=/com.ibm.iea.zos/zos/1.12/Availability/V1R12_Availability_PFA_Enhancements/player.html
 - *z/OS Hot Topics Newsletters* -- http://www.ibm.com/systems/z/os/zos/bkserv/hot_topics.html
 - #20 (GA22-7501-16) -- *Fix the Future with Predictive Failure Analysis* by Jim Caffrey, Karla Arndt, and Aspen Payton
 - #23 (GA22-7501-19) – *Predict to prevent: Let PFA change your destiny* by Jim Caffrey, Karla Arndt, and Aspen Payton
 - #23 (GA22-7501-19) – *Runtime to the Rescue! Using Runtime Diagnostics to find out your problems fast* by Bob Abrams, Don Durand, and Dave Zingaretti
 - *IBM Systems Magazine - Mainframe Edition*
 - PFA A Soft Touch by Karla Arndt, Jim Caffrey, and Aspen Payton
 - http://www.ibmssystemsmagmainframedigital.com/nxtbooks/ibmsystemsmag/mainframe_20101112/index.php#/48

PD Simplification

Simplifying tasks

- z/OS Service Aids**
- Dumps
 - Logs
 - Traces
 - Tools



Soft Failure Detection: Predictive Failure Analysis (PFA)

- | | |
|------------------------|----------------------|
| Common Storage Usage | SMF Arrival Rate |
| Frames and Slots Usage | JES2 Spool Usage |
| Logrec Arrival Rate | Enqueue Request Rate |
| Message Arrival Rate | |

System Symptom Analysis: Runtime Diagnostics

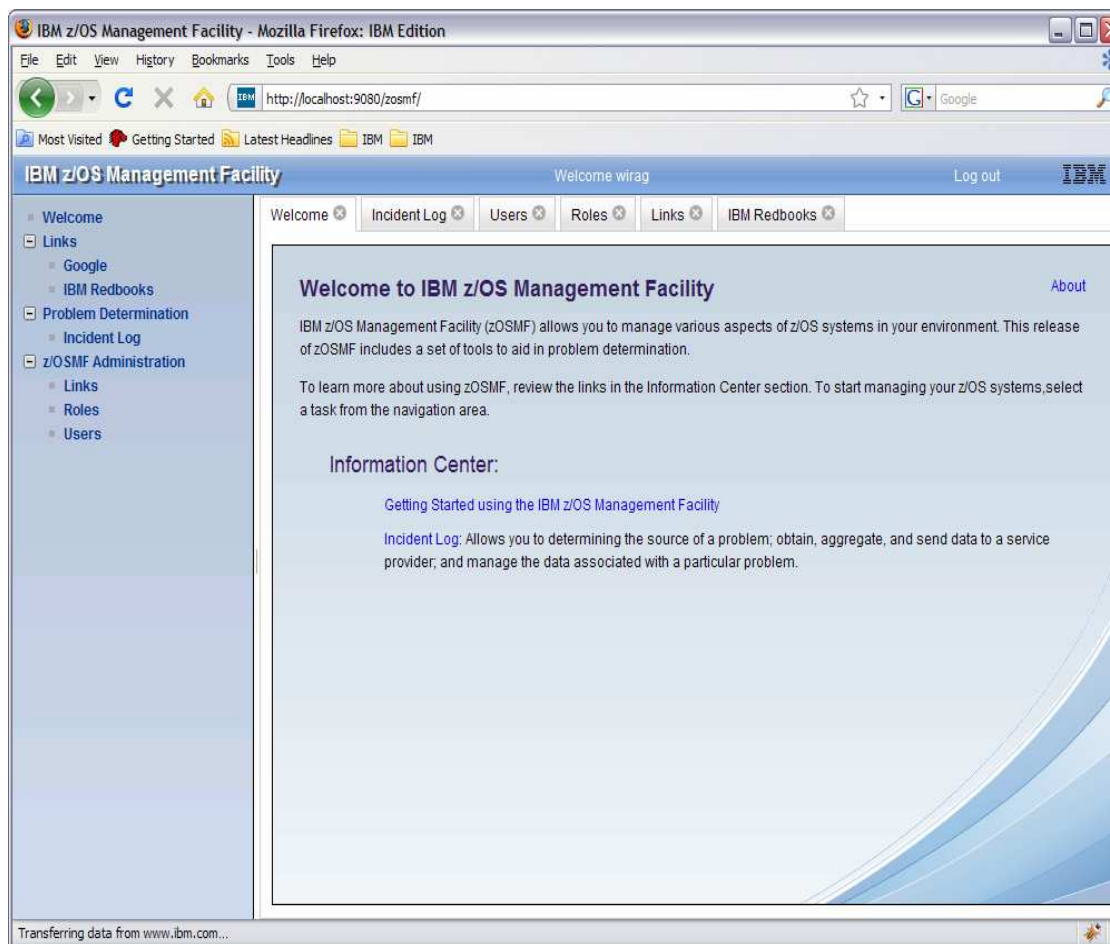
- Analyze a sick system upon request (operator command)
- Component errors
- Address space behavior (ENQ, CPU usage, Loops, etc.)

Problem Data Management: Incident Log

- Improve FFDC for system-detected problems
- Diagnostic data “snapshots” for transient data
- FTP Incident with all diagnostic data
- User interface to display summary, detail, drive actions

IBM z/OS Management Facility

- IBM z/OS Management Facility (z/OSMF), a zero priced product, simplifies, optimizes and modernizes the z/OS **system programmer** experience
- Task oriented, Web browser based user interface
- Integrated user assistance
- **z/OSMF Benefits:**
 - Help improve system programmer productivity, and make the functions easier to understand and use
 - Makes the day to day operations and administration of the mainframe z/OS systems easier to manage for both new and experienced system programmers



Problem Determination in a complex environment

Installation Pain Points

<p>Risk to the business</p> <ul style="list-style-type: none"> • The impact of the symptoms • Risk of recurrence • Impact in getting system stabilized • Mean time to recovery too long
<p>Complexity of performing the task</p>
<p>Troubleshooting a live system and recovering from an apparent failure</p>
<p>Data collection very time-consuming</p>
<p>Significant skill level needed to analyze problems, interact with IBM and ISVs to obtain additional diagnostic info</p>



Requirement Areas

<p>Detect “sick, but not dead” event BEFORE it causes problems</p>
<p>Diagnose the cause in real time to allow operations to mitigate event inquiries</p>
<p>Manage / capture data to determine cause of problem</p> <ul style="list-style-type: none"> • Allow problem to be fixed to prevent recurrence

Focus on Problem Determination

- **Focus on Problem Determination capability - Incident Log:**
 - The incident log and underlying z/OS diagnostic data gathering greatly improves the tasks related to:
 - Identifying system-detected problems (related to SVC dumps taken by the system)
 - Collecting diagnostic materials related to a problem and sending materials to IBM or another company's support area
 - Tell the system to take the next dump for a previously-recognized problem

z/OSMF Problem Determination – Incident Log

- **Auto-capture basic diagnostic materials, triggered when the dump is written to a data set, managed via PARMLIB member**
 - **Initial focus is on Abend and user initiated SVC dumps**
 - Improved FFDC for system-detected problems
 - Diagnostic data “snapshots” for transient data: Snapshots of 30 min Operlog or Syslog, 1 hr Logrec detail, and 4-hour Logrec summary
 - Incident Log will support the creation of diagnostic log snapshots based on the SYSLOG dataset or OPERLOG logstream and LOGREC data sets or LOGREC sysplex log streams
 - Allow doc to be tersed and FTP'd to IBM (or ISV) without having to keep track of where logs are archived via easy to use interface
 - Simplify informing DAE to take the next dump for the incident's symptom string
- **Functions include:**
 - Display list of incidents (Filter/ sort/ configure/ delete)
 - Display properties – view list of diagnostic data, logs
 - Set properties: associate problem number and tracking id (R11), new fields and more customization capabilities (R12)
 - Send diagnostic data via FTP: Manage FTP jobs status and define FTP Profiles (firewall) (R11), support for encrypted and parallel FTP (R12)
 - Send additional user-defined diagnostic data
 - Enhance scope of diagnostic log snapshots created (R12)
 - Allow next dump

Incident Log – Summary Information

Many fields, set tracking IDs

Incident Type	Description	Date and Time (GMT)	Problem Number	Tracking ID	Release	Notes	Pr
ABEND S00C4	COMPID=5665-28502,ISSUER=IKJEFT05. SDUMP FROM THE TMP ESTAE ROUTINE	Apr 27, 2010 1:42:04 PM	12981,999,000		V1R12	data sent by Anuja	
ABEND S06C4	COMPON=ATR,COMPID=SCRRES,ISSUER=ATRCMRE	Jun 1, 2010 3:36:34 PM	12345,123,123		V1R12		BC
ABEND S00C4	COMPON=CNZ,COMPID=SC1CK,ISSUER=CNZMIREC	May 15, 2010 3:55:44 AM	12345		V1R12		BC
User Initiated	SLIP DUMP ID=0005	May 6, 2010 4:04:26 PM	12345		V1R12		
User Initiated		Jun 25, 2010 1:36:22 PM			V1R12		
User Initiated		Jun 24, 2010 8:52:04 PM			V1R12		
User Initiated		Jun 24, 2010 8:27:15 PM			V1R12		
User Initiated		Jun 24, 2010 8:20:41 PM			V1R12		
User Initiated		May 6, 2010 3:16:54 PM			V1R12		
User Initiated		May 6, 2010 3:11:36 PM			V1R12		
User Initiated		May 6, 2010 2:55:22 PM			V1R12		
ABEND U0013	ABEND=U0013,RC=0000,COMPON=SDSF-ESTAE,COMPID=5694-A01,ISSUER=ISFSTAE,SDSF ABEND ROUTINE	Apr 15, 2010 6:47:53 AM			V1R12	data sent by Anuja	SD

Total: 12, Filtered: 12, Selected: 0
 Refresh Last refresh: Jul 20, 2010 10:17:23 PM local time (Jul 21, 2010 2:17:23 AM GMT)

Incident Log – Summary Information

Screen scrolled to the right >>>

Easy identification! Identifies what product and component. The product could be DB2 or CICS or any IBM or vendor product, as long as it produces a SVC dump that the system can recognize.

Welcome Incident Log

Incident Log

Incident Type Filter	Description Filter	Date and Time (GMT) Dates from Apr 12, 2010 12:00:00 AM	Problem Number Filter	Tracking ID Filter	Release Filter	Notes Filter	Product Filter	Component Name Filter	Component ID Filter
<input type="checkbox"/> ABEND S00C4	COMPID=5665-28502,ISSUER=IKJEFT05. SDUMP FROM THE TMP ESTAE ROUTINE	Apr 27, 2010 1:42:04 PM	12981,999,000		V1R12	data sent by Anuja			
<input type="checkbox"/> ABEND S05C4	COMPON=ATR,COMPID=SCRRS,ISSUER=ATRCMREI	Jun 1, 2010 3:36:34 PM	12345,123,123		V1R12		BCP	RRS	5752SCRRS
<input type="checkbox"/> ABEND S00C4	COMPON=CNZ,COMPID=SC1CK,ISSUER=CNZMIREC,	May 15, 2010 3:55:44 AM	12345		V1R12		BCP	COMTASK	5752SC1CK
<input type="checkbox"/> User Initiated	SLIP DUMP ID=0005	May 6, 2010 4:04:26 PM	12345		V1R12				
<input type="checkbox"/> User Initiated	TEST1	Jun 25, 2010 1:36:22 PM			V1R12				
<input type="checkbox"/> User Initiated	DUMP1	Jun 24, 2010 8:52:04 PM			V1R12				
<input type="checkbox"/> User Initiated	DUMP1	Jun 24, 2010 8:27:15 PM			V1R12				
<input type="checkbox"/> User Initiated	DUMP1	Jun 24, 2010 8:20:41 PM			V1R12				
<input type="checkbox"/> User Initiated	SLIP DUMP ID=0003	May 6, 2010 3:16:54 PM			V1R12				
<input type="checkbox"/> User Initiated	ZMFUSR2 LOOP	May 6, 2010 3:11:36 PM			V1R12				
<input type="checkbox"/> User Initiated	SLIP DUMP ID=0002	May 6, 2010 2:55:22 PM			V1R12				
<input type="checkbox"/> ABEND U0013	ABEND=U0013,RC=0000,COMPON=SDSF-ESTAE,COMPID=5694-A01,ISSUER=ISFSTAE,SDSF ABEND ROUTINE	Apr 15, 2010 6:47:53 AM			V1R12	data sent by Anuja	SDSF	SDSF	566548801

Incident Log – Incident Details

The screenshot displays the IBM z/OS Management Facility web interface in Mozilla Firefox. The browser title is "IBM z/OS Management Facility - Mozilla Firefox: IBM Edition" and the address bar shows "https://mysystemz:32208/zosmf/". The page header includes "Welcome zmaadm" and "Log out".

The left sidebar contains a navigation menu with the following items:

- Welcome
- Links
- Problem Determination
 - Incident Log

 A "Refresh" button is located below the menu.

The main content area is titled "Incident Log" and "View Diagnostic Details". It features two tabs: "General" and "Diagnostic Data". A green arrow points to the "Diagnostic Data" tab, which is highlighted by a blue callout box containing the text: "Tab shows lists of data (logrec and error log)".

The "Diagnostic Data" tab displays the following incident information:

- Incident type: ABEND
- Incident description: COMPON=WEBSPPHERE Z/OS, COMPID=5655N0200,ISSUER=BBORLEXT,ABEND IN (MODULE NAME NOT KNOWN)
- Date and time (GMT): Oct 9, 2009 7:16:49 PM
- System name: DCEPLX
- System name: S1
- Problem number: [input field]
- Tracking ID: DB: 5868,Scr: XR-125
- Component name: [input field]
- Component ID: 5655N0200
- z/OS release: V1R12
- Product: [input field]
- Abend code: S00D6
- Reason code: 00000027
- CSECT: ...
- Load module: ...
- Symptom string: ...
- Notes: Application problem

At the bottom of the page, there are "OK", "Apply", and "Cancel" buttons. A status bar at the very bottom indicates "Transferring data from localhost..." and "localhost:9443".

Incident Log – Diagnostic Data

IBM z/OS Management Facility - Mozilla Firefox

https://9.12.41.62:32208/zosmf/

IBM z/OS Management Facility Welcome debug13 Log out

- Welcome
- Configuration
 - Configuration Assistant
- Links
 - ShopzSeries
 - Support for z/OS
 - System z Redbooks
 - WSC Flashes
- Users

Refresh

Incident Log View Diagnostic Details

View Diagnostic Details

General Diagnostic Data

Data Type	Source	Sysplex	System
<input type="checkbox"/> SVC dump	ZMFDUMP.DYNZOS12.P03.D100114.T153115.SV000	SVPLEX6	P03
<input type="checkbox"/> Error log	CEA.L00.C56360D8.A3D56F9E	SVPLEX6	P03
<input type="checkbox"/> Operations log	CEA.Q00.C56360D8.A3D56F9E	SVPLEX6	P03

Total: 3, Selected: 0

Attachments

To send additional information that you think is relevant for this incident, such as a trace, use the **New...** action in the following table to specify the files to send. You can attach up to five additional files per send. The information about the attachments is discarded when you close the panel.

New...

Data Type	Source
There is no data to display.	

Total: 0, Selected: 0

Send View Status

OK Apply Cancel

Done 9.12.41.62:32208

Incident Log – Send Diagnostic Data

IBM z/OS Management Facility - Mozilla Firefox
 https://dceimgne.pdl.pok.ibm.com:32208/zosmf/

IBM z/OS Management Facility Welcome zmfusr1 Log out IBM

Incident Log > Send Diagnostic Data

Send Diagnostic Data

Use this wizard to prepare and send diagnostic data to a predefined FTP destination.

Review the selected diagnostic data and enter a problem number. If the data to send is incorrect, click **Cancel** to exit the wizard. To select the data you want to send, use the Incident Log or View Diagnostic Details panel.

Incident	Description	Date and Time (GMT)
ABEND S0EC3	COMPON=WEBSPPHERE Z/OS, COMPID=5655N0200,ISSUER=BBORLEXT,ABEND IN (MODULE NAME NOT KNOWN)	Aug 4, 2009 9:26:29 PM

Diagnostic Data to Send		
Data Type	Sysplex	System
SVC dump	CFCIMGNE	DCEIMGNE
Error log	CFCIMGNE	DCEIMGNE
Operations log	CFCIMGNE	DCEIMGNE
Error log summary	CFCIMGNE	DCEIMGNE

* Problem number:
 If the problem number is an IBM PMR number, check this box to verify the syntax.

< Back Next > Finish Cancel

Done dceimgne.pdl.pok.ibm.com:32208

Incident Log: Defining a FTP destination

The screenshot shows the IBM z/OS Management Facility web interface in a Mozilla Firefox browser. The page title is "New FTP Destination" under the "Incident Log > FTP Destinations > New" path. The interface includes a left-hand navigation menu with options like "Welcome", "Configuration", "Links", "Performance", "Problem Determination", and "z/OSMF Administration".

The main content area contains the following configuration fields:

- * System:** testcase.boulder.ibm.com
- * Path name:** /toibm/mvs
- Port number (must be between: 1-65535):** 21
- Use anonymous user ID and password
- Description:** (empty text field)
- Transfer method:**
 - FTP
 - z/OS Problem Documentation Upload Utility (Parallel FTP with optional encryption)
- Work data set prefix:** (empty text field)
- * Work data set size (MB):** 100
- Work data set data class:** (empty text field)
- Work data set storage class:** (empty text field)
- * Number of FTP sessions:** 3

At the bottom of the form are "OK" and "Cancel" buttons. A callout box on the right side of the form contains the text: "Option to use parallel FTP with encryption when sending data to IBM *".

Incident Log – Delete Incident

The screenshot shows the IBM z/OS Management Facility web interface in Mozilla Firefox. The browser address bar shows <https://dceimgne.pdl.pok.ibm.com:32209/zosmf/>. The page title is "IBM z/OS Management Facility" and the user is logged in as "Welcome pegusr".

The "Incident Log" section is active, displaying a table of incidents. A "Confirm Delete" dialog box is overlaid on the table, asking for confirmation to delete selected incidents. The dialog includes a warning icon and the text: "IZUP835W You are about to delete the selected incidents and all associated information. Do you want to continue?". Below this, it lists "Selected incidents" with a scrollable list containing: "ABEND S0913 Feb 23 2009 12:37:10 PM COMPON=COMP". There is also a checkbox for "Allow next dump" and "OK" and "Cancel" buttons.

On the left side of the interface, a menu is visible with the following items: Welcome, Links, Problem Determination, Incident Log. A blue callout box with a white border is positioned over the "Delete Incident..." menu item, which is circled in blue. The callout box contains the following text: "Set Tracking ID...", "Set Problem Number...", "Delete Incident...", "Send Diagnostic Data...", "View Diagnostic Details", "FTP Job Status", and "Allow Next Dump...".

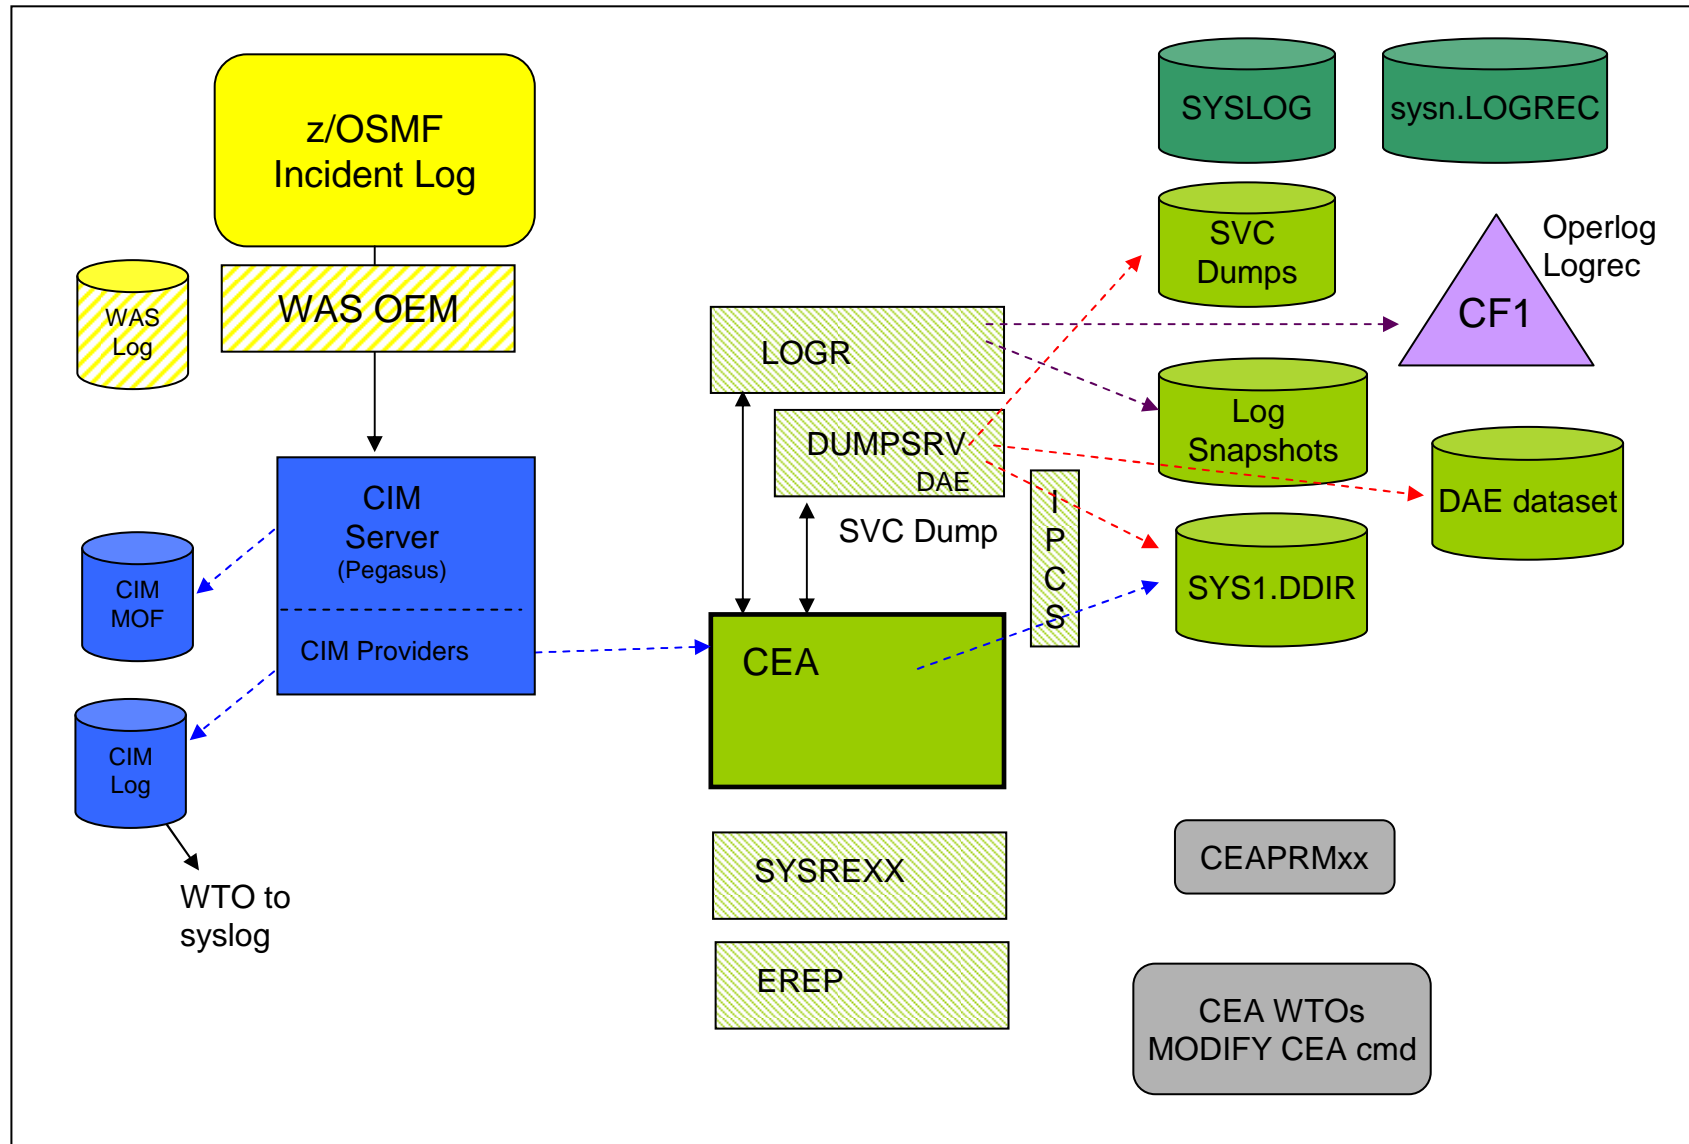
Incident Type	Description	Date and Time	Sysplex	System	Problem Number	Tracking ID	Release
<input type="checkbox"/>	ABEND=40D, RC=...	Past 90 days	Filter	Filter	41336,180,000		V1R1*
<input type="checkbox"/>	User Initiated				41336,180,000		V1R1*
<input type="checkbox"/>	ABEND S0913				12345		V1R1*
<input checked="" type="checkbox"/>	ABEND S0913				45678,057,649	3456789	V1R1*
<input type="checkbox"/>	User Initiated				41336,180,000	ar12345	V1R1*
<input type="checkbox"/>	User Initiated				41336,180,000		V1R1*
<input type="checkbox"/>	User Initiated				41336,180,000		V1R1*
<input type="checkbox"/>	ABEND S00C4				41336,180,000	ar12345	V1R1*
<input type="checkbox"/>	ABEND S00C4						V1R1*
<input type="checkbox"/>	User Initiated	Feb 18 2009 7:16:39 AM	CFCIMGNE	DCEIMGNE			V1R1*
<input type="checkbox"/>	User Initiated	Feb 17 2009 2:01:16 PM	CFCIMGNE	DCEIMGNE	41336,180,000		V1R1*

z/OSMF Problem Determination – Incident log *Benefits*

	Without z/OSMF Incident Log **	With z/OSMF Incident Log **
Recognizing a system-detected (dumped) problem occurred	Requires 5 to 7 manual steps, plus skill on effective use of IPCS to extract data from each of the dumps. Up to 5-6 minutes	Display in 1 click. Greatly reduced skill required As little as 5 seconds
Allow new dump to be taken for the same symptom	Requires 7 to 12 manual steps, plus skill on effective use of IPCS to locate the dump data set, obtain the symptom string, get into the IPCS DAE display, locate the matching symptom string (could be non-trivial) and indicate TakeNext on the IPCS display Up to 15 minutes	Make the update happen in 3 mouse clicks As little as 10 seconds
Collecting and sending diagnostic data	Requires 7 to 15 manual steps, plus skill to locate the right log files, build and run jobs, rename the output datasets, and use an FTP job to send the different data sets to the target destination. Up to 20 minutes Up to 30 minutes for sysplex components	Send the material in 8 clicks: <ul style="list-style-type: none"> •Select the incident materials •Specify the FTP destination information •Send the material •Check whether the information was FTP'd successfully As little as 30 seconds

“So easy, even a marketing professional can use it!” – Gita Grube Berg, IBM System z Marketing

** Based on IBM laboratory results, your results may vary



Additional Incident Log V1R12 enhancements

- The **Incident Log task** under the Problem Determination category is enhanced with the following new functions:
 - Encryption of the incident files, including dumps; transmission of these files to IBM in parallel through FTP to save time
 - To do so, the host and destination must have the z/OS Problem Documentation Upload Utility installed
 - Send additional documentation (attachments) with an incident to an FTP destination
 - Provides free form Notes or comments for each incident
- **Diagnostic data** is improved
 - Logrec Summary Report now based on LOGDATAS
 - Hardware related Logrec reports are captured when I/O-related failing component
 - TYPE C CCH/CRW/SLH
 - TYPE H MIH MISSING INT. HANDLER
 - TYPE O OBR OUTBOARD RECORDS/UNIT CHECKS
 - TYPE X DPS
 - TYPE I IPL
 - TYPE E EOD END OF DAY

Incident Log configuration

CEAPRMxx

```
SNAPSHOT(Y)
HLQ(CEA)
DUMPCAPTURETIME
(
  SLIP(OPERLOG(00:30:00) LOGREC(01:00:00)
    LOGRECSUMMARY(04:00:00))

  DUMP(OPERLOG(00:30:00) LOGREC(01:00:00)
    LOGRECSUMMARY(04:00:00))

  ABEND(OPERLOG(00:30:00) LOGREC(01:00:00)
    LOGRECSUMMARY(04:00:00))
)
COUNTRYCODE(XXX)
BRANCH(XXX)
STORAGE(STORCLAS(STANDARD))
```

- Requires Sysplex Dump Directory to be set up
- Automated diagnostic data capture assumes Operlog or SYSLOG(R12), and Logrec log streams or Logrec data sets (R12)
 - Log Snapshots captured to DASD Log Streams
- Automatic Dump data set allocation recommended
- DAE recommended to be active
- AMATERSE in an APF authorized library
- System REXX used

For more information on z/OSMF and Incident Log

- *z/OSMF Configuration Guide*, SA38-0652
- z/OSMF presentations delivered at this SHARE
- z/OS Hot Topics
 - August 2009 (“Simplification” issue)
 - An introduction to z/OSMF
 - An introduction to the z/OSMF Incident Log
 - Setting up Operlog and Logrec for z/OSMF Incident Log
 - Common Event Adapter
 - Using System Logger for z/OSMF
 - August 2010
 - Easy as z/OSMF
 - A new look for z/OSMF
 - Enhancements to the z/OSMF Incident Log

Related SHARE Sessions


Session Number	Day	Time	Room	Session Title
8696	Tues, Mar. 1	11:00 A.M.	201B	z/OS Problem Determination Update: z/OSMF Incident Log, Runtime Diagnostics, PFA
9061	Tues, Mar. 1	1:30 P.M.	201B	Understanding WebSphere App Server OEM Edition for z/OSMF Sysprogs
8922	Tues, Mar. 1	3:00 P.M.	201D	z/OSMF 1.12 Overview
8658	Tues, Mar. 1	4:30 P.M.	201D	z/OSMF Implementation and Configuration
9035	Wed, Mar. 2	8:00 A.M.	201D	Smarter SVC Dump Processing for Improved z/OS Resiliency
8668	Wed, Mar. 2	11:00 A.M.	201D	Detecting Soft Failures using z/OS PFA
8699	Wed, Mar. 2	6:00 P.M.	207	z/OSMF Roundtable
9039	Thurs, Mar. 3	8:00 A.M.	201B	z/OS Software Deployment
8757	Thurs, Mar. 3	11:00 A.M.	201D	z/OSMF User Experience
9075	Thurs, Mar. 3	4:30 P.M.	208B	z/OSMF Hands on Lab
8859	Fri, Mar. 4	11:00 A.M.	202A	Managing Your Workload with z/OSMF

Summary

Problem Determination Simplification

Built on z/OS's robust (and continually evolving) RAS technology

- Predictive Failure Analysis
- Runtime Diagnostics
- z/OSMF Incident Log
- Base Serviceability functions



- ↻ Machine-speed understanding
- ↻ Better tooling to identify the culprit
- ↻ Integrated Problem data management
- ↻ Enables faster / correct recovery actions