#SHAREorg

SHARE
Technology · Connections · Results

RMF – The Latest and Greatest

Stefan Wirag (stefan.wirag@de.ibm.com)
IBM Corporation

Monday, March 12, 2012
Session 10890

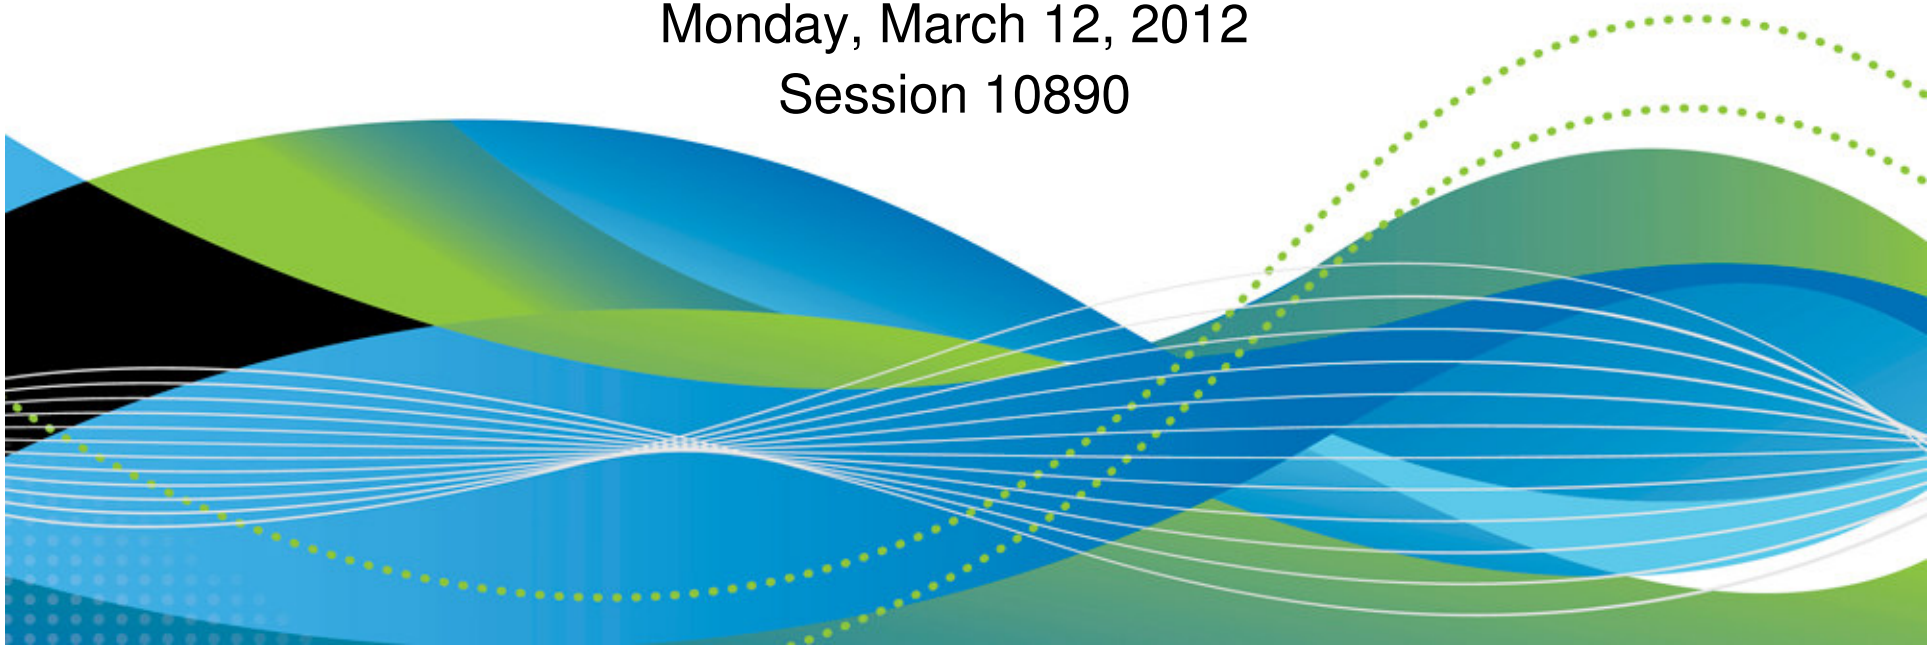

Trademarks

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.

RMF Product Overview

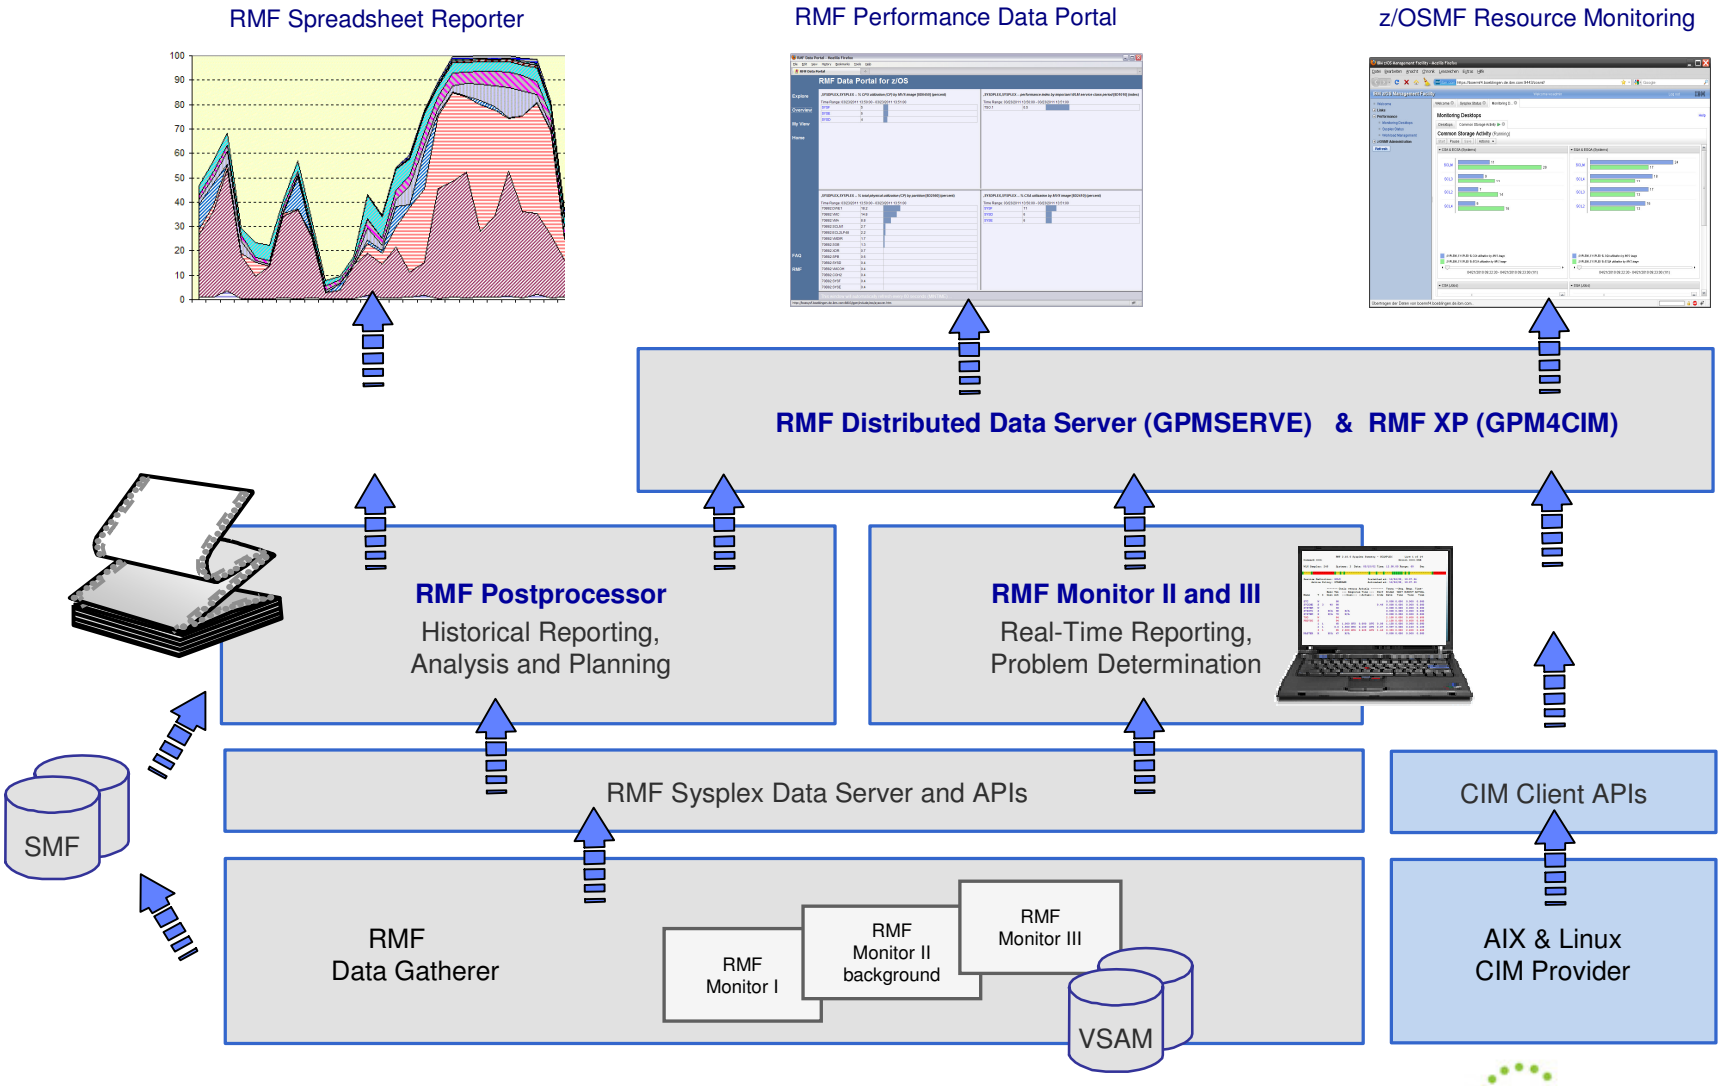

RMF Enhancements

- RMF WLM Reporting Enhancements
  - Enhanced monitoring for transactional workloads with Velocity Goals

z/OS V1.13

- RMF Storage Reporting Enhancements
  - Appropriate reporting of 64-Bit Storage

z/OS V1.13

- Postprocessor XML Formatted Reports
  - State-of-the-art UI for SMF Type 7x Data
  - Standardized Report Format for API Access
  - Remote Reporting via Spreadsheet Reporter

z/OS V1.11

- HTTP API to access Historical data
  - Enable RMF Distributed Data Server to request historical RMF Postprocessor data via HTTP

z/OS V1.12

- Reporting of system serialization delays
  - Detect serialization-related performance issues
  - System suspend lock, GRS enqueue and latch contention information

z/OS V1.13

- In-Ready Work Unit Queue Distribution
  - Identify latent Demand
  - Dispatchable Unit Granularity

z/OS V1.12

- z/OSMF Resource Monitoring
  - Plugin of z/OSMF Management Facility
  - Cross-sysplex performance monitoring from a single point of control

z/OS V1.12

- RMF XP
  - New solution for Cross Platform Performance Monitoring
  - Seamless monitoring for all zBX related platforms

z/OS V1.13

RMF WLM Reporting Enhancements

- Currently WLM reporting does not provide a response time distribution (ended transactions) for workloads with velocity goals
- But it is desirable to have a response time distribution for all transactional workloads, even if they have a velocity goal
  - More data to analyze workload behavior and to detect problems
  - Better support for migration of goal definitions to response time goals
- With z/OS V1.13 the WLM IWMRCOLL interface provides response time distribution data for service class periods with an execution velocity goal
- RMF for z/OS 1.13 exploits the new IWMRCOLL data:
  - RMF Monitor I data gatherer collects the new response time information in SMF 72 subtype 3 record
  - RMF Postprocessor Workload Activity report (WLMGL) displays
    - Response time distribution for response time and execution velocity goals
    - One merged distribution for workloads with response time goals per sysplex
    - One distribution for workloads with execution velocity goals per system in sysplex

RMF WLM Reporting Enhancements …

- Response Time Goal
  Example : Goal = 80% within 2 sec

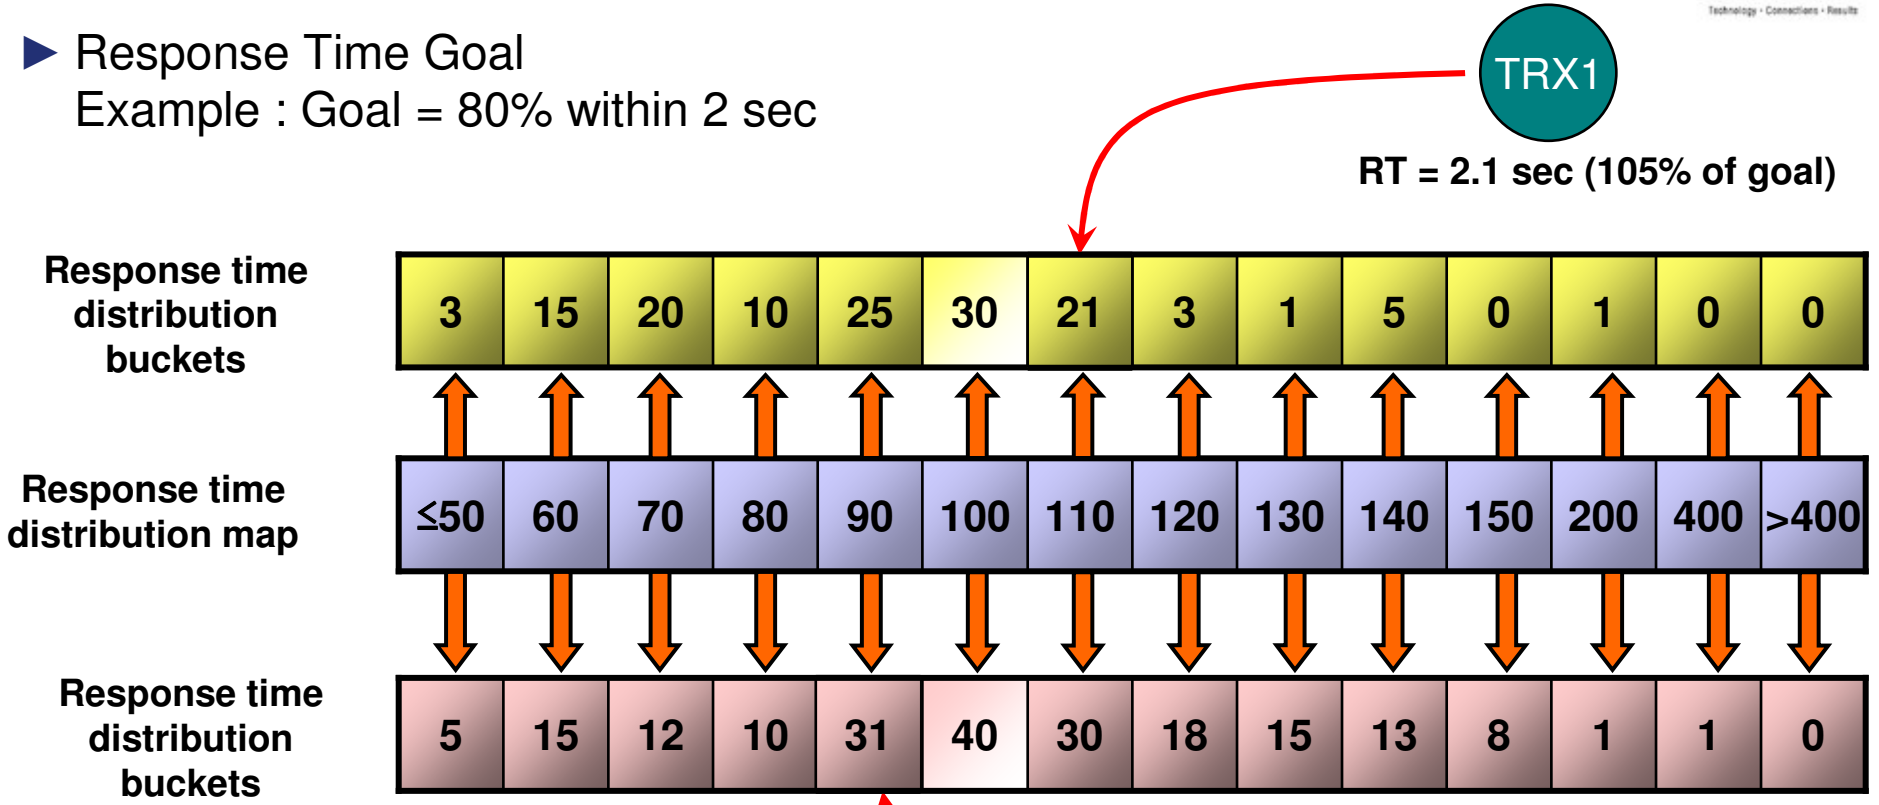

- Execution Velocity Goal
  Example : Velocity = 90%
  Average Response Time = 3.1 sec

$$Execution\ Velocity\ =\ 100 \times \frac{Total\ Using}{Total\ Using\ +\ Total\ Delay}$$

RMF WLM Reporting Enhancements …

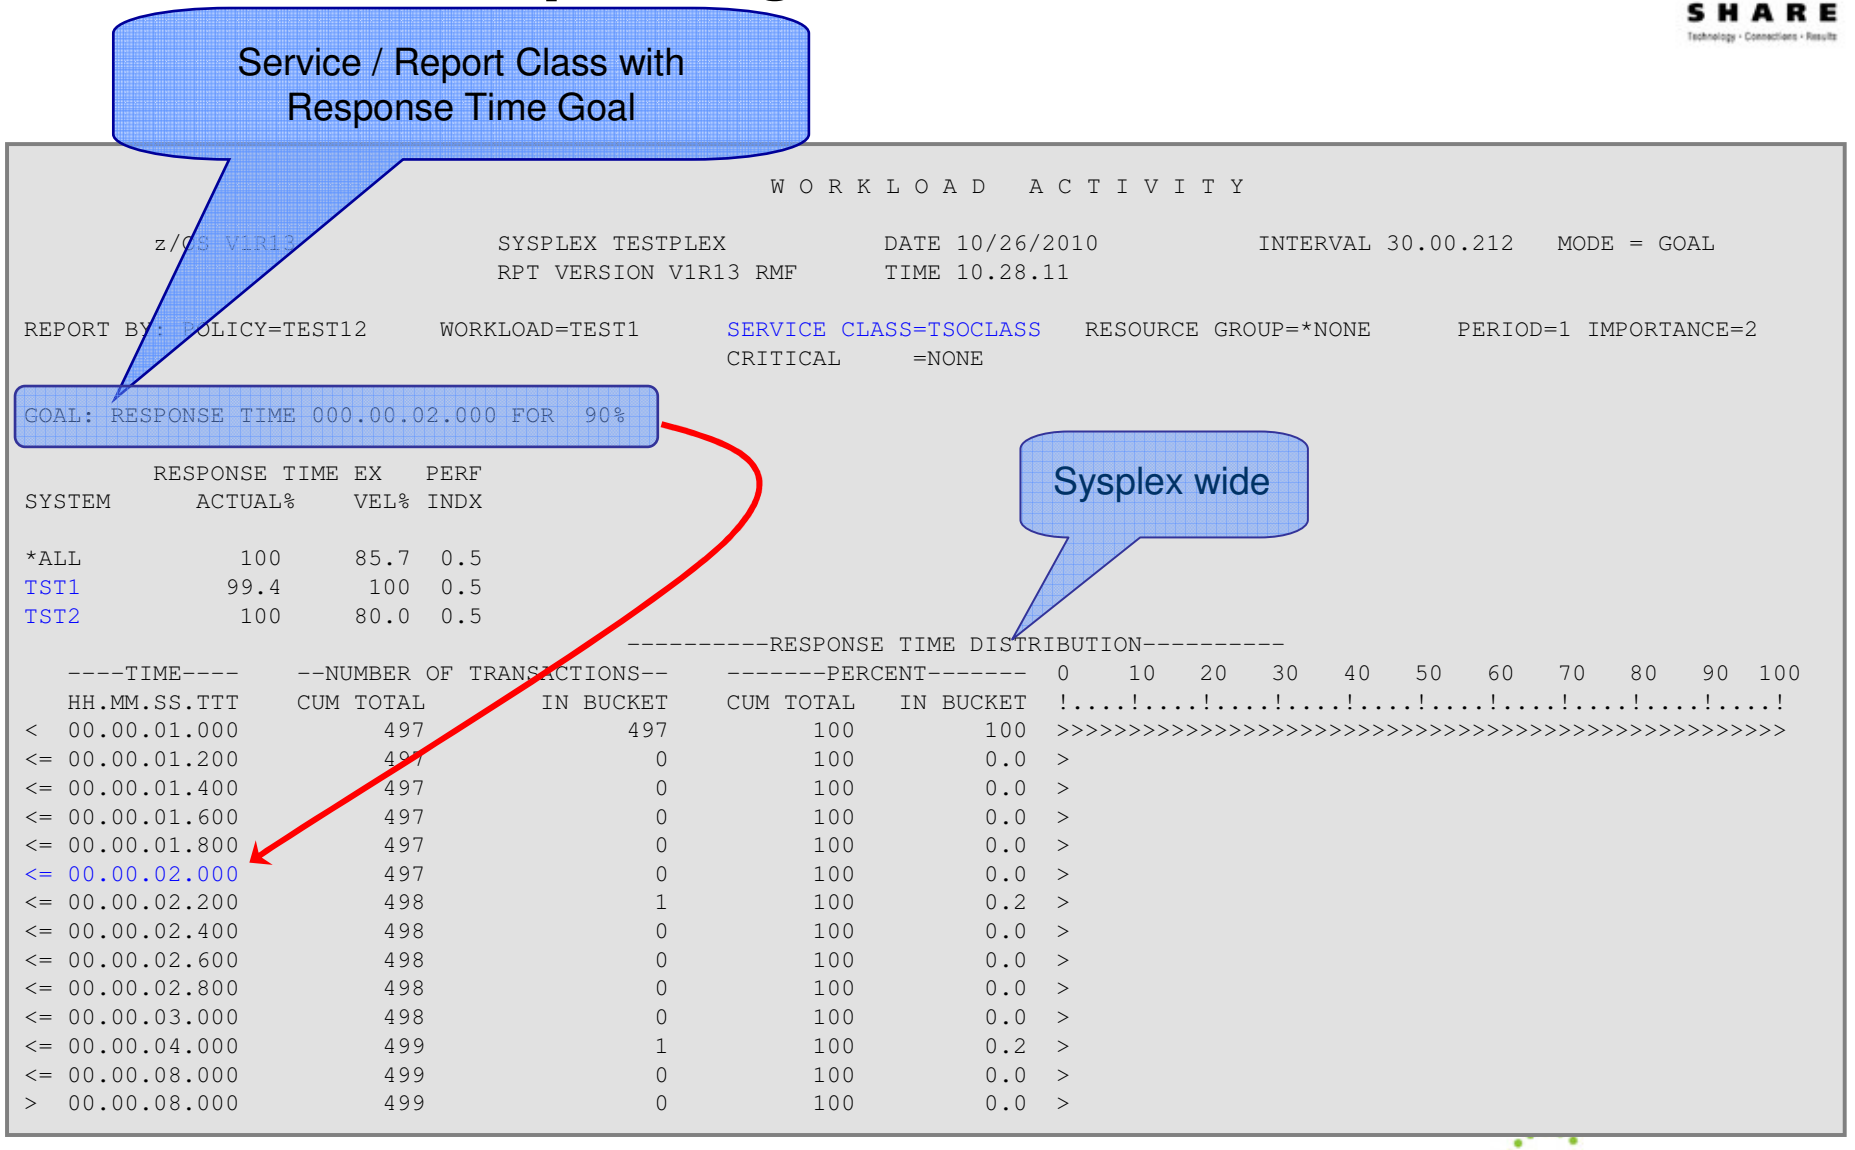


```
                                    W O R K L O A D   A C T I V I T Y

        z/OS V1R13                  SYSPLEX TESTPLEX          DATE 10/26/2010          INTERVAL 30.00.212   MODE = GOAL
                                    RPT VERSION V1R13 RMF     TIME 10.28.11

REPORT BY: POLICY=TEST12     WORKLOAD=TEST1      SERVICE CLASS=TSOCLASS   RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=2
                                                 CRITICAL     =NONE

GOAL: RESPONSE TIME 000.00.02.000 FOR  90%

         RESPONSE TIME EX    PERF
SYSTEM       ACTUAL%    VEL% INDX

*ALL            100     85.7  0.5
TST1           99.4      100  0.5
TST2            100     80.0  0.5

                                          ----------RESPONSE TIME DISTRIBUTION----------
   ----TIME----     --NUMBER OF TRANSACTIONS--    -------PERCENT-------  0    10   20   30   40   50   60   70   80   90  100
   HH.MM.SS.TTT     CUM TOTAL        IN BUCKET    CUM TOTAL   IN BUCKET  !....!....!....!....!....!....!....!....!....!....!
<  00.00.01.000           497              497          100         100  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
<= 00.00.01.200           497                0          100         0.0  >
<= 00.00.01.400           497                0          100         0.0  >
<= 00.00.01.600           497                0          100         0.0  >
<= 00.00.01.800           497                0          100         0.0  >
<= 00.00.02.000           497                0          100         0.0  >
<= 00.00.02.200           498                1          100         0.2  >
<= 00.00.02.400           498                0          100         0.0  >
<= 00.00.02.600           498                0          100         0.0  >
<= 00.00.02.800           498                0          100         0.0  >
<= 00.00.03.000           498                0          100         0.0  >
<= 00.00.04.000           499                1          100         0.2  >
<= 00.00.08.000           499                0          100         0.0  >
>  00.00.08.000           499                0          100         0.0  >
```

RMF WLM Reporting Enhancements …

Service / Report Class with Execution Velocity Goal

```
                                W O R K L O A D   A C T I V I T Y

      z/OS V1R13              SYSPLEX TESTPLEX          DATE 10/26/2010          INTERVAL 30.00.212   MODE = GOAL
                              RPT VERSION V1R13 RMF     TIME 10.28.11

REPORT BY: POLICY=VICOM2     WORKLOAD=VICOM     SERVICE CLASS=VEL90I2    RESOURCE GROUP=*NONE     PERIOD=1 IMPORTANCE=2
                                                 CRITICAL     =NONE
GOAL: EXECUTION VELOCITY 90.0%     VELOCITY MIGRATION:   I/O MGMT  96.0%     INIT MGMT 44.9%

           RESPONSE TIME EX   PERF
SYSTEM                VEL%  INDX

*ALL         --N/A--    96.0  0.9
TST1                    95.3  0.9
TST2                    96.7  0.9
```

Measurement Interval

Number of midpoint changes

```
                                    ----------RESPONSE TIME DISTRIBUTIONS----------
SYSTEM: TST1 -----INTERVAL: 15.49.248 -----MRT CHANGES:  1 ---
   ----TIME----  -NUMBER OF TRANSACTIONS-  ------PERCENT------
   HH.MM.SS.TTT  CUM TOTAL    IN BUCKET   CUM TOTAL IN BUCKET
<  00.00.05.308          9            9         5.9       5.9
<= 00.00.06.370         45           36        29.4      23.5
<= 00.00.07.431         72           27        47.1      17.6
<= 00.00.08.493         83           11        54.2       7.2
<= 00.00.09.555         90            7        58.8       4.6
<= 00.00.10.617        117           27        76.5      17.6
<= 00.00.11.678        153           36         100      23.5
<= 00.00.12.740        153            0         100       0.0
<= 00.00.13.802        153            0         100       0.0
<= 00.00.14.863        153            0         100       0.0
<= 00.00.15.925        153            0         100       0.0
<= 00.00.21.234        153            0         100       0.0
<= 00.00.42.468        153            0         100       0.0
>  00.00.42.468        153            0         100       0.0

SYSTEM: TST2 -----INTERVAL: 17.01.121 -----MRT CHANGES:  1 ---
   ----TIME----  -NUMBER OF TRANSACTIONS-  ------PERCENT------
   HH.MM.SS.TTT  CUM TOTAL    IN BUCKET   CUM TOTAL IN BUCKET
<  00.00.03.871          0            0         0.0       0.0
<= 00.00.04.645          0            0         0.0       0.0
<= 00.00.05.419          8            8         4.4       4.4
<= 00.00.06.193         48           40        26.2      21.9
<= 00.00.06.967         70           22        38.3      12.0
<= 00.00.07.742         87           17        47.5       9.3
<= 00.00.08.516        101           14        55.2       7.7
<= 00.00.09.290        113           12        61.7       6.6
<= 00.00.10.064        119            6        65.0       3.3
<= 00.00.10.838        160           41        87.4      22.4
<= 00.00.11.613        183           23         100      12.6
<= 00.00.15.484        183            0         100       0.0
<= 00.00.30.968        183            0         100       0.0
>  00.00.30.968        183            0         100       0.0
```

RMF Storage Reporting Enhancements

Rationale

- Structure and metrics in RMF PP Paging Activity report no longer appropriate to present today's system memory in a 64-Bit world:
  - Expanded Storage section removed but layout not adapted
  - Metrics for Physical swapping to auxiliary storage and expanded storage obsolete
  - Increased number of metrics for high virtual storage:
    - Memory Objects in high virtual common, shared or private storage
    - Large Pages (1MB) for high private and common storage

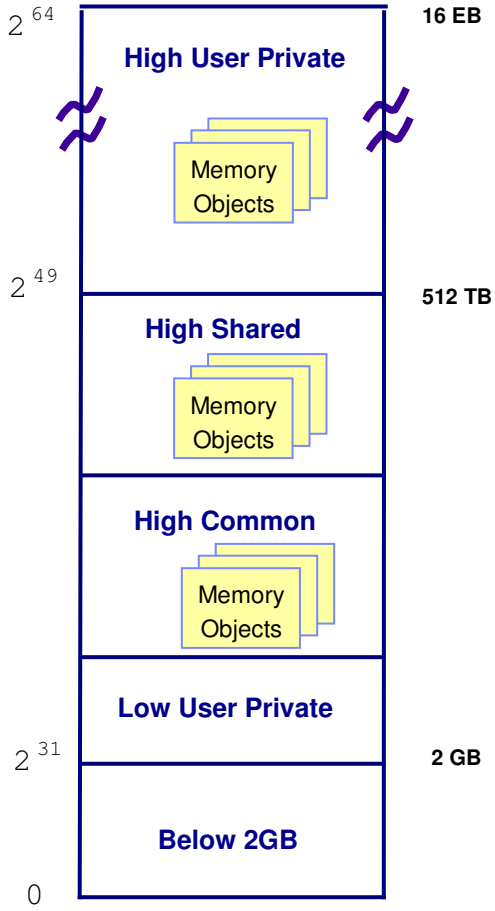

RMF Postprocessor Paging Activity report restructured and enhanced

RMF Storage Reporting Enhancements …

- RMF Enhancements for RMF Monitor I Paging Activity Report
  - Following report sections restructured:
    - Central Storage Paging Rates
    - Central Storage Movement and Request Rates
    - Frame and Slot Counts
  - New Metrics:
    - LFAREA size
    - Central storage frames allocated for high virtual common and shared memory
  - Swap Placement Activity section removed from report
  - New overview condition **LSWAPTOT** for the total logical swap rate. Rate calculated from the accumulated logical swap counts for all swap reasons.
  - New Memory Objects and Frames section to display statistics for high virtual memory objects and frames
  - Report now available in XML format

RMF Storage Reporting Enhancements …

Section headers and MIN/MAX/AVG columns now displayed **as rows**

Frames allocated for high virtual shared and common memory

New

P A G I N G A C T I V I T Y

FRAME AND SLOT COUNTS

P A G I N G A C T I V I T Y

FRAME AND SLOT COUNTS

(181 SAMPLES)

CENTRAL STORAGE FRAMES	TOTAL	AVAILABLE	SQA	LPA	CSA	LSQA	REGIONS+SWA	HV SHARED	HV COMMON
MIN	7,864,316	2,751,600	20,217	6,429	16,601	104,493	4,776,978	77,276	18,387
MAX	7,864,316	2,840,625	20,636	6,438	16,636	105,735	4,865,152	77,276	18,389
AVG	7,864,316	2,784,680	20,465	6,434	16,620	105,347	4,832,283	77,276	18,388

FIXED FRAMES	TOTAL	NUCLEUS	SQA	LPA	CSA	LSQA	REGIONS+SWA	<16 MB	16MB-2GB
MIN	319,129	2,821	18,735	81	15,952	28,812	252,349	56	27,791
MAX	324,274	2,821	19,154	81	15,952	29,482	257,151	71	28,857
AVG	320,866	2,821	18,982	81	15,952	29,314	253,715	56	28,313

SHARED FRAMES	TOTAL SLOTS	CENTRAL STORAGE	FIXED TOT	FIXED BEL	AUX DASD
MIN	81,456	72,227	51	0	0
MAX	81,476	72,247	51	0	0
AVG	81,460	72,231	51	0	0

LOCAL PAGE DATA SET SLOTS	TOTAL	AVAILABLE	BAD	NON-VIO	VIO
MIN	5,399,997	5,399,619	0	20	0
MAX	5,399,997	5,399,977	0	378	0
AVG	5,399,997	5,399,654	0	343	0

COMMON FIXED	1000	2000	1500
SHARED	1000	2000	1500
1 MB	10	20	15

RMF Storage Reporting Enhancements …

- Statistics about Memory Objects and Frames moved from FRAMES AND SLOT COUNTS section to new MEMORY OBJECTS AND FRAMES section

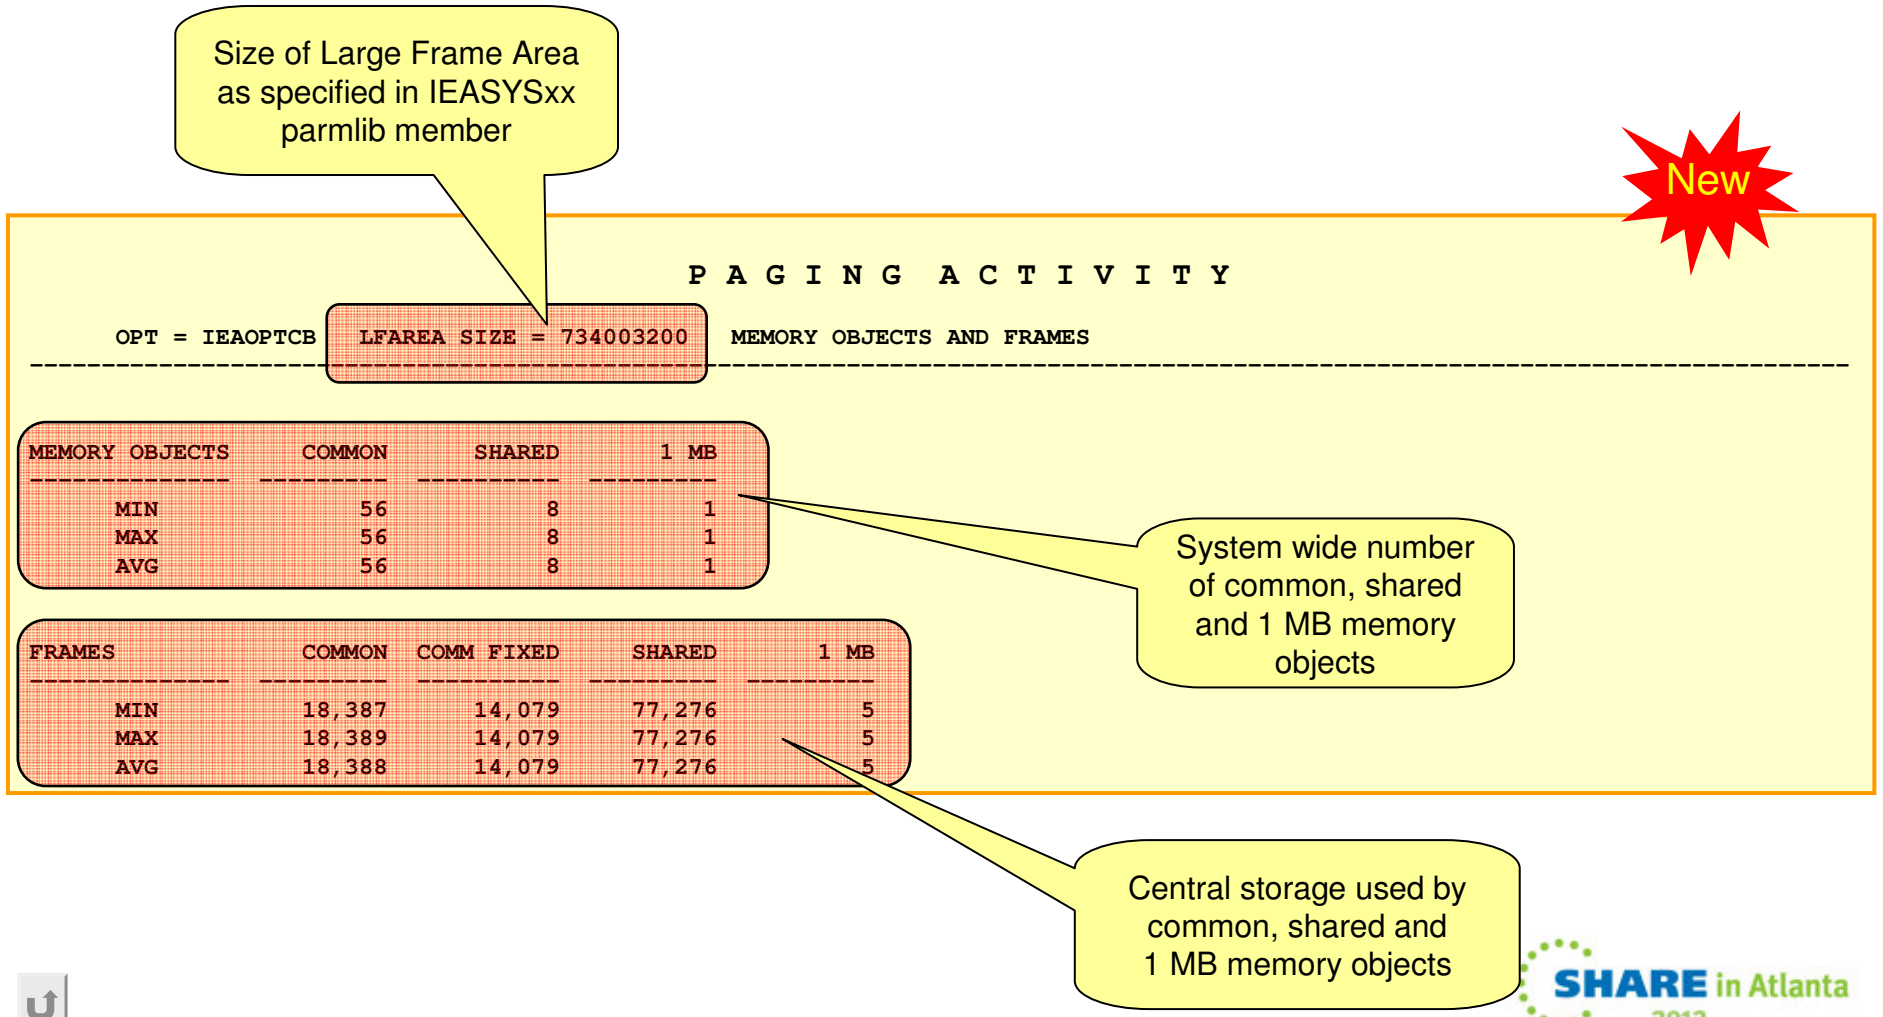

Postprocessor XML Formatted Reports

Rationale:

- RMF Postprocessor reports are limited to a page width of 132 characters
- No state-of-the-art display capability of Postprocessor reports
- No easy access to RMF Postprocessor data for application programs
  - ⇨ cumbersome to parse the text output
  - ⇨ each report has its own layout

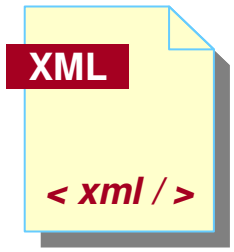

RMF Postprocessor reports can now be generated in XML Format

Postprocessor XML Formatted Reports

- The following single-system reports are available in XML format (both Interval and Duration reports)
  - ➡ CPU Activity report (including Partition Data, LPAR Cluster and Group Capacity reports)
  - ➡ CRYPTO Activity report
  - ➡ FICON Director Activity report
  - ➡ OMVS Kernel Activity report
  - ➡ ESS Disk Systems Activity report
  - ➡ DEVICE Activity report — z/OS 1.12
  - ➡ PAGING Activity report
  - ➡ Serialization Delay report (SDELAY) — z/OS 1.13
- With z/OS 1.12, the first sysplex report is available in XML format
  - ➡ Workload Activity report (WLMGL) — z/OS 1.12
- Overview reports are also available in XML format
- Summary and Exception reports are not available in XML format
- Interval reports based on data collected during a Monitor II background session are not available in XML format

Postprocessor XML Formatted Reports

- The generation of Postprocessor reports in XML format is controlled by the new ddnames XPRPTS, XPXSRPTS and XPOVWRPT
- If the XML output is routed to permanent data sets rather than to SYSOUT, define the data set with RECFM=VB and LRECL between 256 and 8192. Specify an appropriate BLKSIZE.

ddname	Contents	Allocations	Notes
XPRPTS	Combined **single-system** report in XML format	One ddname for one data set to contain all single system reports for each interval during the session.	There is no dynamic allocation of this ddname, you have to define it explicitly if you want to get all reports in XML format into one data set or output class. If you define this ddname, no MFRnnnnn files are created. If you define this ddname and PPRPTS, no XML output in file XPRPTS is created.
XPOVWRPT	Combined Overview report in XML format	One ddname for one data set to contain all overview reports for each system included in the input data.	There is no dynamic allocation of this ddname, you have to define it explicitly if you want to get all overview reports in XML format into one data set or output class. If you define this ddname, no PPORPnnn files are created.
XPXSRPTS z/OS 1.12	Combined **sysplex-wide** report in XML format	One ddname for one data set to contain all sysplex reports for each interval included in the input data.	There is no dynamic allocation of this ddname, you have to define it explicitly if you want to get all reports in XML format into one data set or output class. If you define this ddname, no MFRnnnnn files are created. If you define this ddname and PPXSRPTS, no XML output in file XPXSRPTS is created.

Spreadsheet Reporter – XML Support

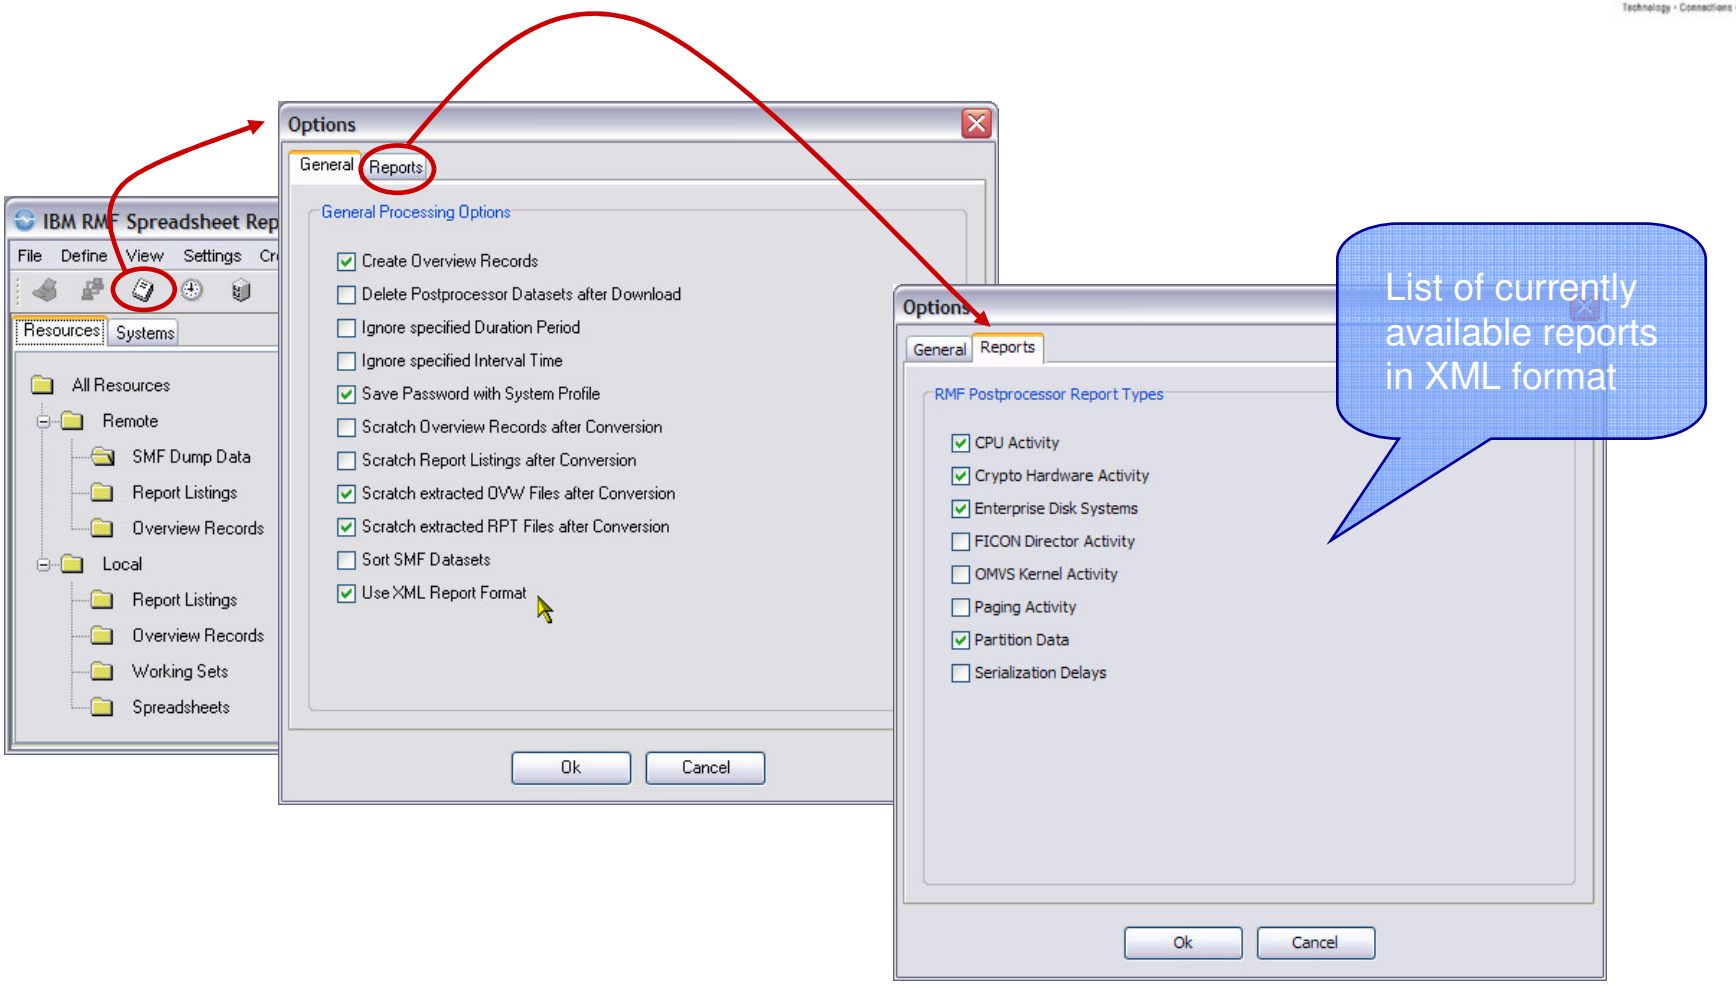

Spreadsheet Reporter – XML Support

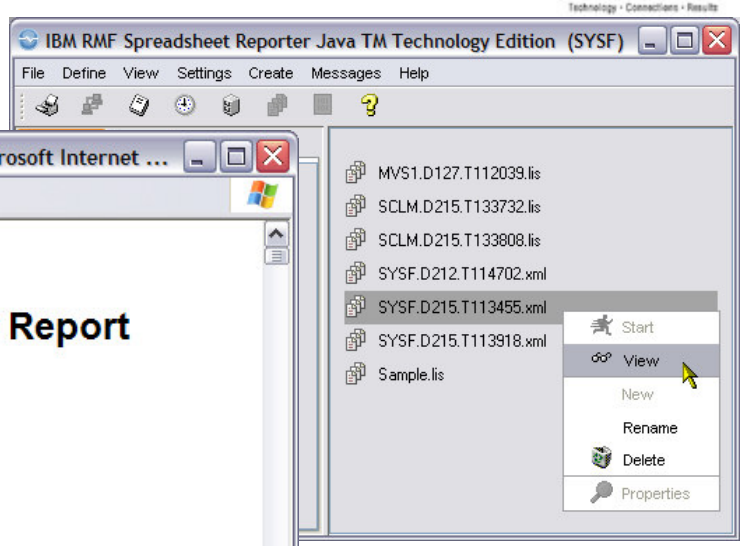

C:\Documents and Settings\Administrator\Application Data\RMF\RMF Spreadsheet Reporter\RmfListin - Microsoft Internet ...

File Edit View Favorites Tools Help

RMF Postprocessor Duration Report [SYSD] : CPU Activity Report

z/OS V1R9 Start: 07/31/2009-08.15.00 Interval: 000:44:59 hours

V1R11 RMF End: 07/31/2009-09.00.00 Cycle: 1000 milliseconds

CPU ACTIVITY

CPU	2097
Model	709
H/W Model	E26
Sequence Code	0000000000070B82
HiperDispatch	NO

CPU Number	CPU Type	Time% Online	Time% LPAR Busy	Time% MVS Busy	Time% Parked	LOG PROC Share%	HiperDispatch Priority	I/O Interrupts Rate	I/O Interrupts% via TPI
0	CP	100.00	3.19	3.15	------	28.7		9.81	0.25
1	CP	100.00	3.14	3.11	------	28.7		13.34	0.61
TOTAL/AVERAGE	**CP**		**3.17**	**3.13**		**57.4**		**23.14**	**0.46**
3	AAP	100.00	0.00	0.00	------	14.7			
TOTAL/AVERAGE	**AAP**		**0.00**	**0.00**		**14.7**			
2	IIP	100.00	0.00	0.00	------	12.8			
TOTAL/AVERAGE	**IIP**		**0.00**	**0.00**		**12.8**			

Done My Computer

XML Support – Postprocessor XML Toolkit

- New Postprocessor XML Toolkit shipped with RMF for z/OS 1.13
- Simplifies display of RMF Postprocessor XML reports in a web browser

⇨ Toolkit Installation:

1. Download member ERBXMLTK as binary file erbxmltk.msi
2. Install MSI Package

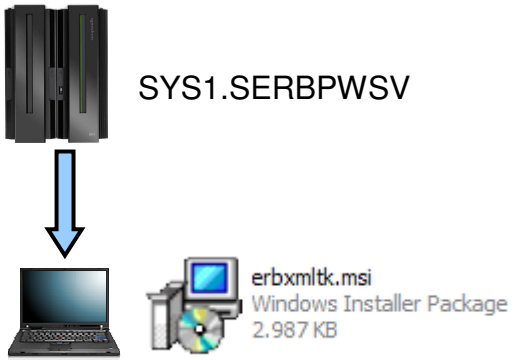

Default Installation Directory

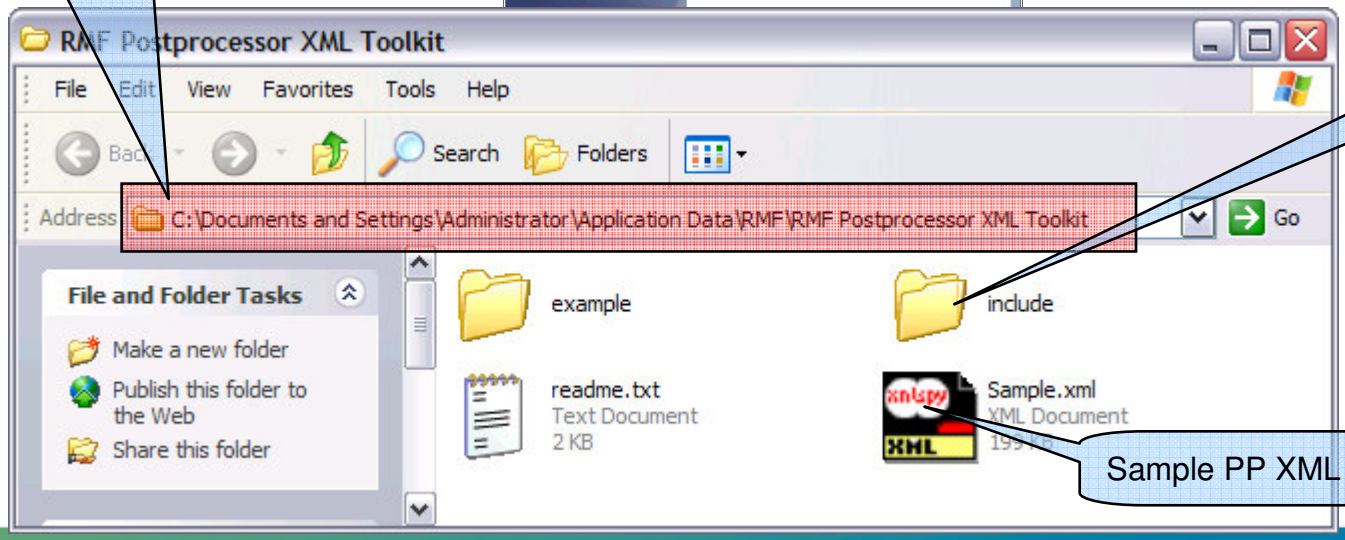

⇨ XSL stylesheet files
⇨ Java script file
⇨ Bitmap files

Sample PP XML report

XML Support – Postprocessor XML Toolkit …

➡ Toolkit Usage:

1. Download Postprocessor XML report into Toolkit directory

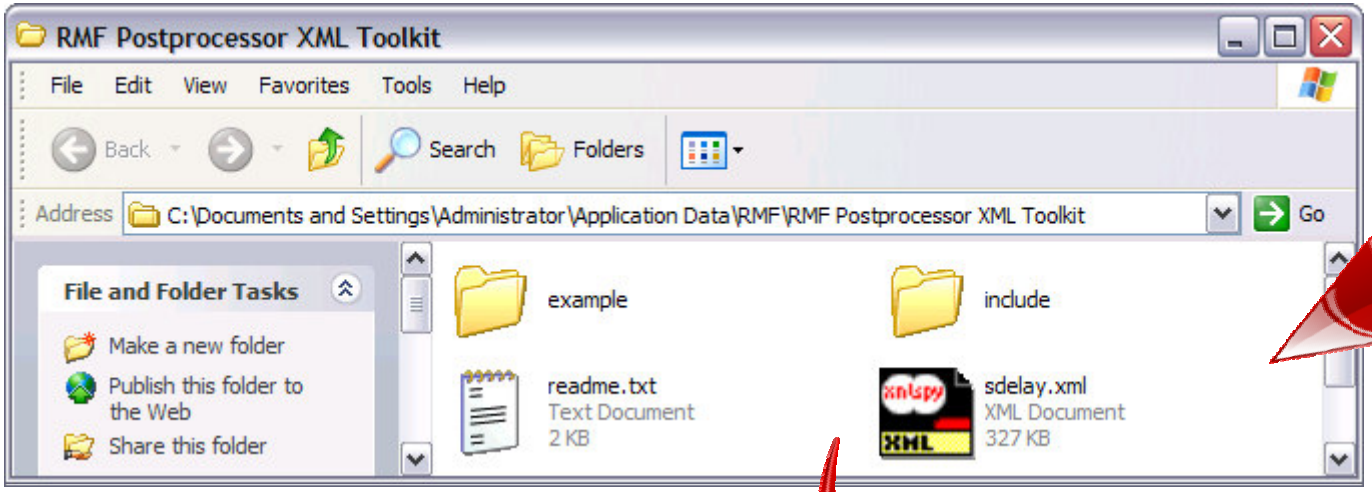

BPMU.SDELAY.LISTING

Download in ASCII format

2. Open XML report with web browser

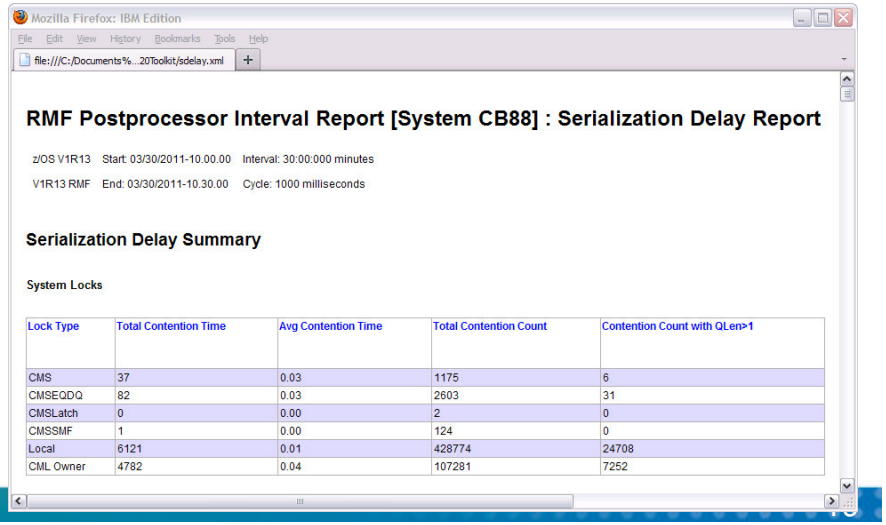

Mozilla Firefox: IBM Edition

RMF Postprocessor Interval Report [System CB88] : Serialization Delay Report

z/OS V1R13 Start: 03/30/2011-10.00.00 Interval: 30:00:000 minutes
V1R13 RMF End: 03/30/2011-10.30.00 Cycle: 1000 milliseconds

Serialization Delay Summary

System Locks

Lock Type	Total Contention Time	Avg Contention Time	Total Contention Count	Contention Count with QLen>1
CMS	37	0.03	1175	6
CMSEQDQ	82	0.03	2603	31
CMSLatch	0	0.00	2	0
CMSSMF	1	0.00	124	0
Local	6121	0.01	428774	24708
CML Owner	4782	0.04	107281	7252

HTTP API to access Historical data

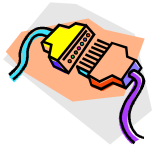

- Application programs can use Distributed Data Server (DDS) HTTP API to retrieve RMF Postprocessor XML reports
- All RMF Postprocessor XML formatted reports supported
- Web browser can be used as Postprocessor Data Portal

http://boesysf:8803/gpm/rmfpp.xml?reports="CPU"&date=20100702,20100702 - Microsoft Internet Explorer

File Edit View Favorites Tools Help

RMF Postprocessor Interval Report [SYSD] : CPU Activity Report

z/OS V1R10 Start: 07/02/2010-11.45.00 Interval: 14:59:999 minutes

V1R12 RMF End: 07/02/2010-12.00.00 Cycle: 1000 milliseconds

CPU ACTIVITY

CPU	2097
Model	722
H/W Model	E40
Sequence Code	0000000000070B82
HiperDispatch	NO
CPC Capacity	1631
Change Reason	N/A

CPU Number	CPU Type	Time% Online	Time% LPAR Busy	Time% MVS Busy	Time% Parked	LOG PROC Share%	HiperDispatch Priority	I/O Interrupts Rate	I/O Interrupts% via TPI
0	CP	100.00	3.50	3.47	------	32.2		4.61	0.19
1	CP	100.00	3.52	3.48	------	32.2		5.07	0.46
TOTAL/AVERAGE	CP		3.51	3.47		64.4		9.68	0.33

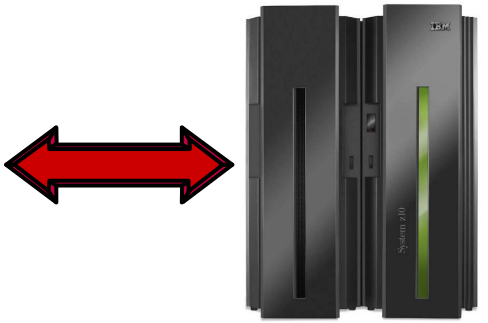

HTTP API to access Historical data …

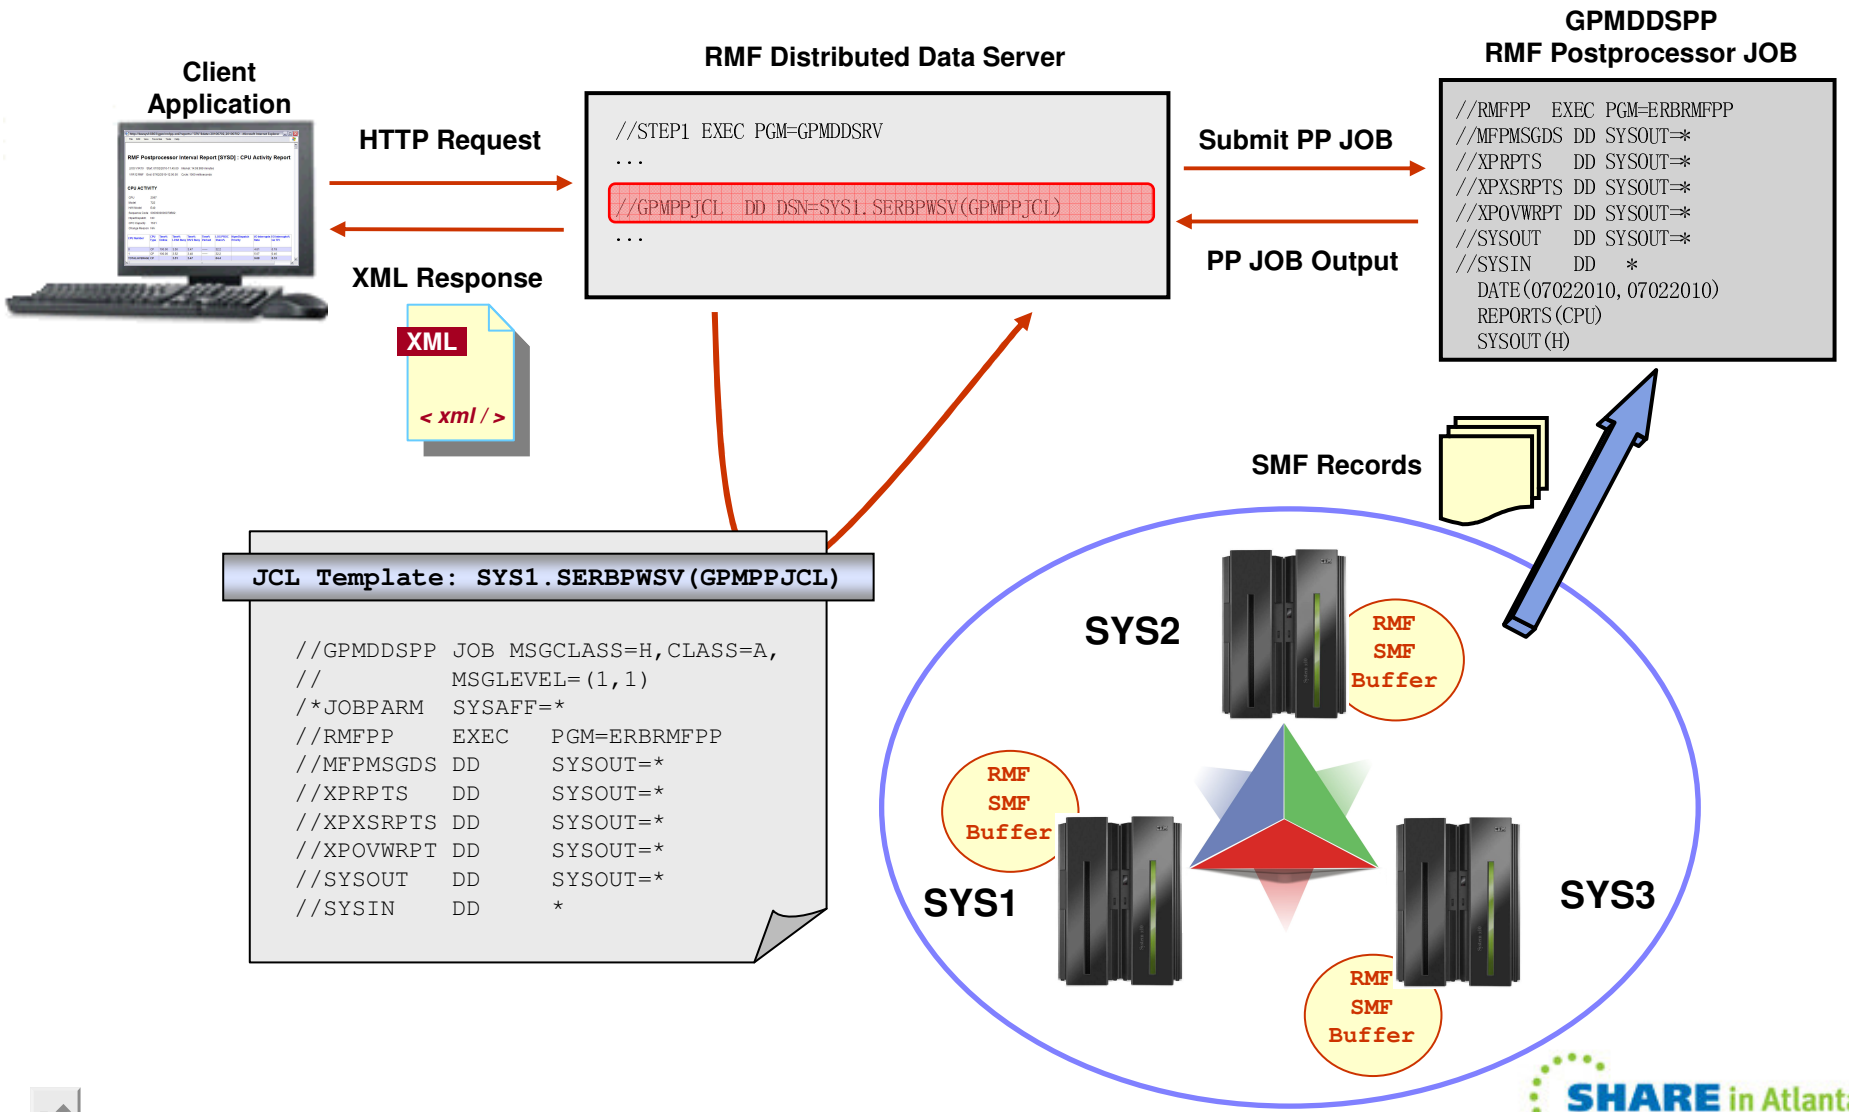

Reporting of system serialization delays

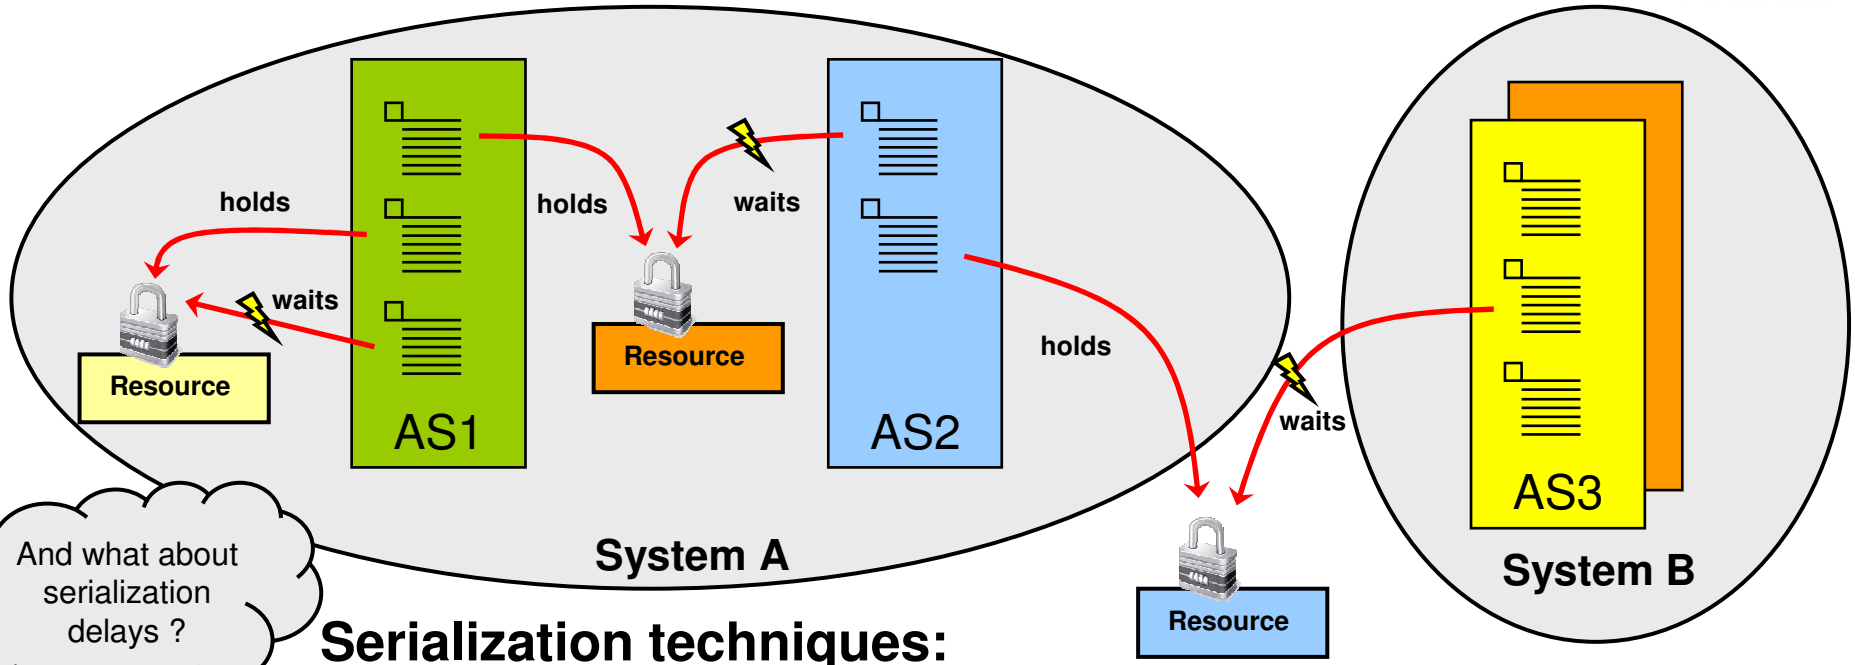

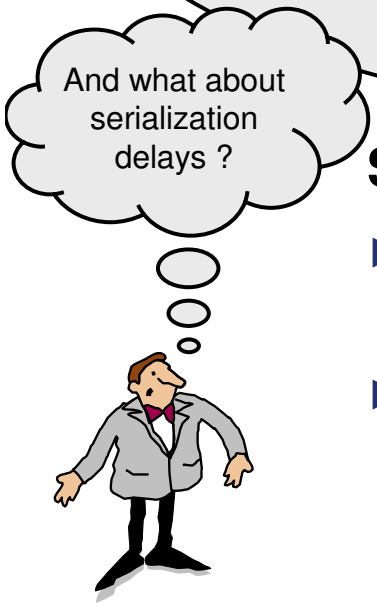

Serialization techniques:

- SETLOCK service for system locks:
  - Serializes system resources within an address space (local lock) or a single MVS system (global lock)
- GRS services:
  - ENQ service
    - Serializes resources with scope JOB STEP, SYSTEM or SYSTEMS (multi system)
  - Latch service
    - High speed serialization service
    - Serializes resources within a single address space or across several address spaces in a single MVS system

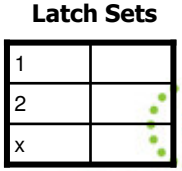

Reporting of system serialization delays …

- RMF for z/OS 1.13 collects and displays system-wide contention information and contention information on address space level
- Reported in new RMF Postprocessor (PP) Serialization Delay Report
- Data about following lock types collected and reported:
  - System Suspend lock types:
    - CMS
    - CMSEQDQ
    - CMSLatch
    - CMSSMF
    - LOCAL
    - CML Lock Owner and
    - CML Lock Requestor
  - GRS lock types:
    - GRS Latch locks
    - GRS Enqueue Step
    - GRS Enqueue System and
    - GRS Enqueue Systems locks

New Serialization Delay Report helps the customer to analyze performance problems due to resource contention situations

Reporting of system serialization delays …

- Contention data collected by RMF Monitor III and stored in new SMF 72 subtype 5 record
  - Data gathering controlled by setting in SMF Parmlib member SMFPRMxx
    - Use TYPE/NOTYPE in Parmlib member to enable/disable data collection

- New RMF Postprocessor (PP) Serialization Delay Report:
  - New suboption SDELAY | NOSDELAY for Postprocessor REPORTS option
  - Only available in XML format
  - XML report can be created and displayed
    - via RMF Spreadsheet Reporter
    - via RMF Data Portal for z/OS
    - via RMF PP Job & Postprocessor XML Toolkit

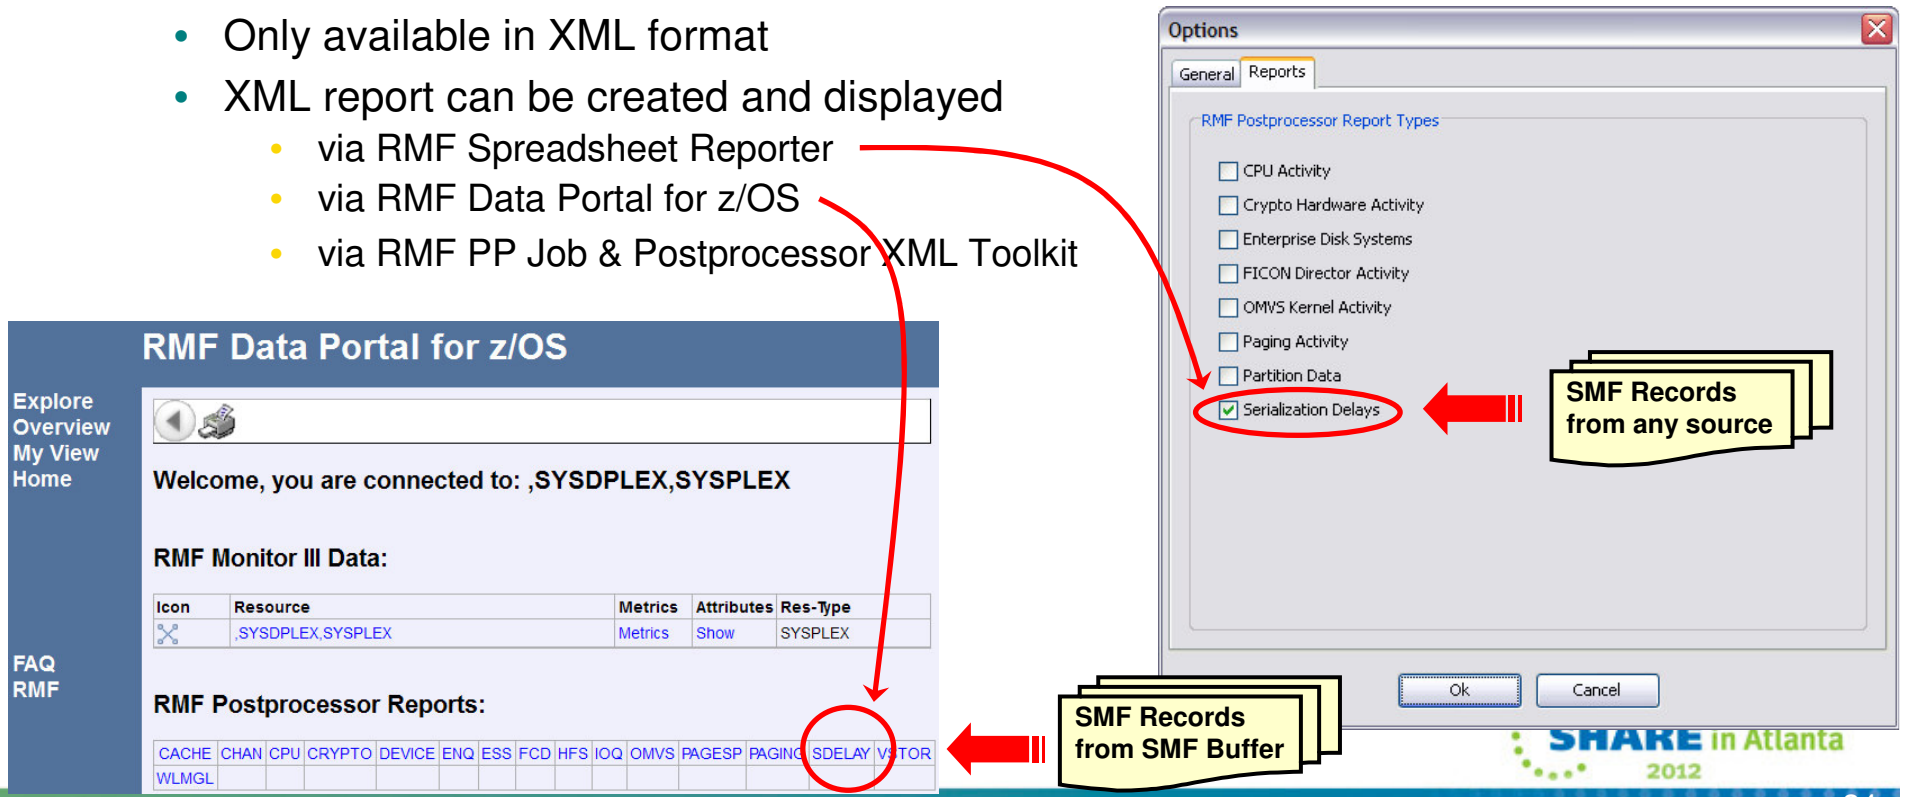

Serialization delays: System Summary

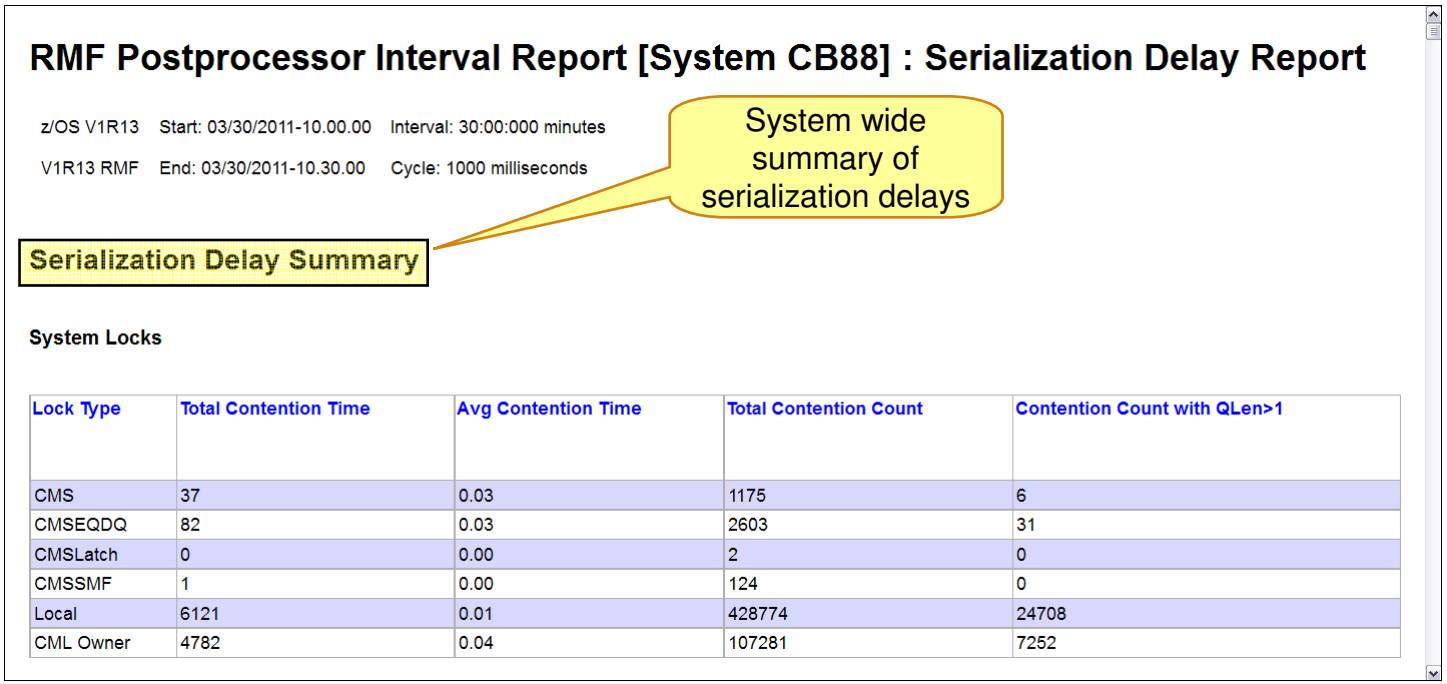

RMF Postprocessor Interval Report [System CB88] : Serialization Delay Report

z/OS V1R13 Start: 03/30/2011-10.00.00 Interval: 30:00:000 minutes

V1R13 RMF End: 03/30/2011-10.30.00 Cycle: 1000 milliseconds

System wide summary of serialization delays

Serialization Delay Summary

System Locks

Lock Type	Total Contention Time	Avg Contention Time	Total Contention Count	Contention Count with QLen>1
CMS	37	0.03	1175	6
CMSEQDQ	82	0.03	2603	31
CMSLatch	0	0.00	2	0
CMSSMF	1	0.00	124	0
Local	6121	0.01	428774	24708
CML Owner	4782	0.04	107281	7252

- System wide summary of system suspend locks
- Locks types:
  - CMS Lock, CMS EnqueueDequeue Lock, CMS Latch Lock and CMS SMF Lock
  - LOCAL Lock and CML Lock Owner
- Total and average contention time in milliseconds
- Number of times a work unit was suspended
- Number of times a unit of work was suspended and there was at least one other unit of work suspended for the same lock.

Serialization delays: System Summary …

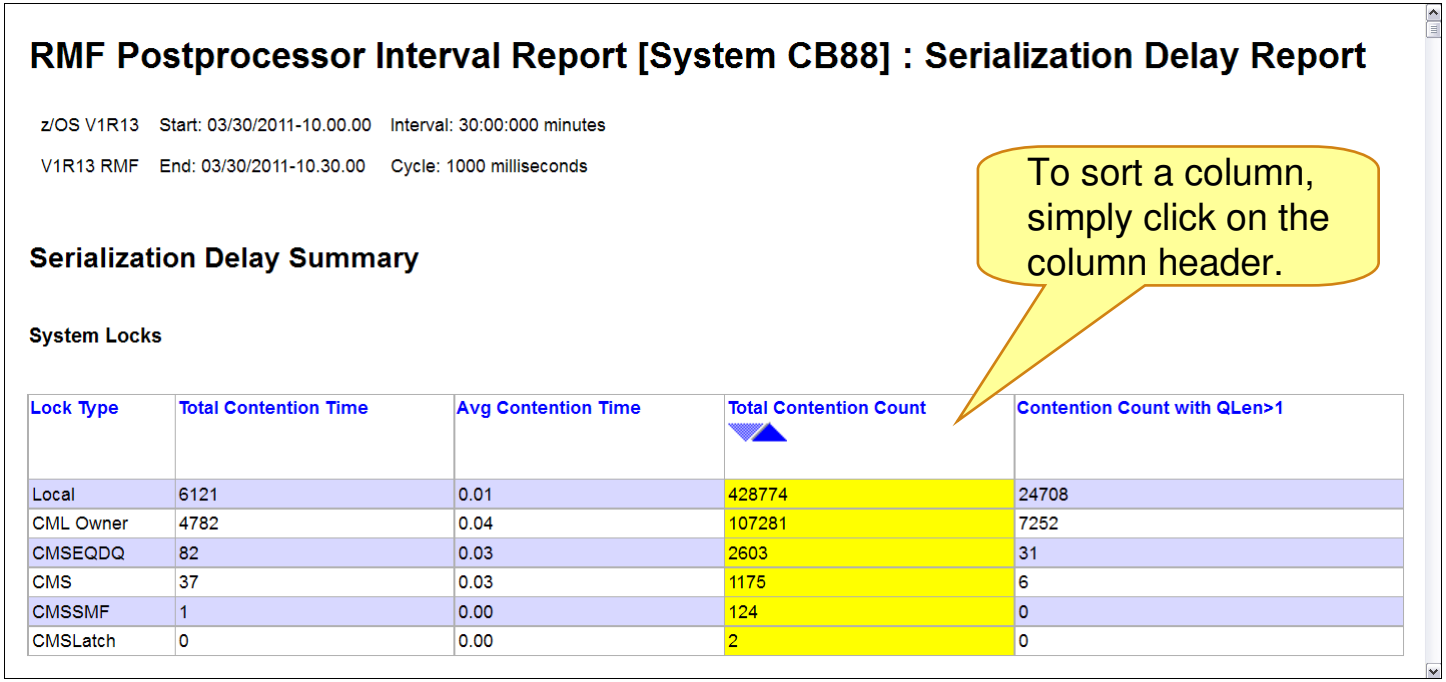

RMF Postprocessor Interval Report [System CB88] : Serialization Delay Report

z/OS V1R13 Start: 03/30/2011-10.00.00 Interval: 30:00:000 minutes

V1R13 RMF End: 03/30/2011-10.30.00 Cycle: 1000 milliseconds

To sort a column, simply click on the column header.

Serialization Delay Summary

System Locks

Lock Type	Total Contention Time	Avg Contention Time	Total Contention Count	Contention Count with QLen>1
Local	6121	0.01	428774	24708
CML Owner	4782	0.04	107281	7252
CMSEQDQ	82	0.03	2603	31
CMS	37	0.03	1175	6
CMSSMF	1	0.00	124	0
CMSLatch	0	0.00	2	0

- Sort capability. Just click on a column header of your choice and the report is sorted according to the values of the selected column
- The color of this column turns to yellow to indicate the currently active sort criteria
- By default, the sort order is descending. A second click on the column header switches to ascending sort order

Serialization delays: System Summary …

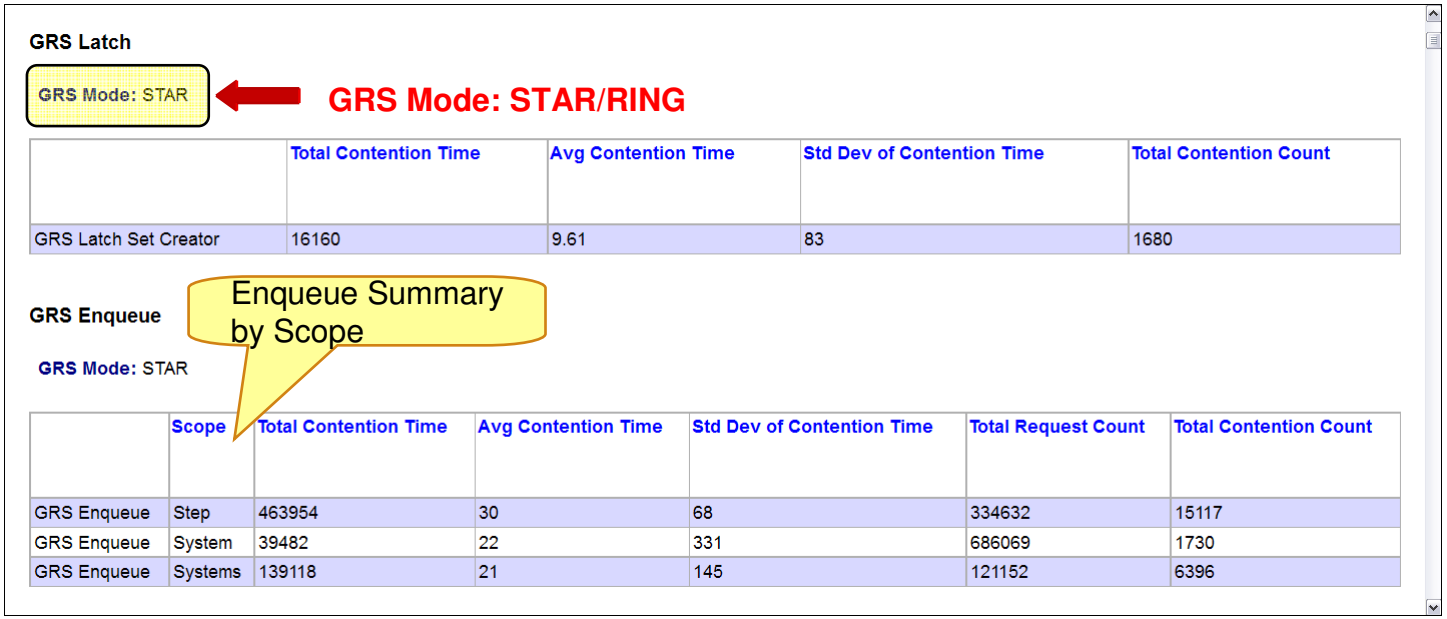

GRS Latch

GRS Mode: STAR ← **GRS Mode: STAR/RING**

	Total Contention Time	Avg Contention Time	Std Dev of Contention Time	Total Contention Count
GRS Latch Set Creator	16160	9.61	83	1680

GRS Enqueue

Enqueue Summary by Scope

GRS Mode: STAR

	Scope	Total Contention Time	Avg Contention Time	Std Dev of Contention Time	Total Request Count	Total Contention Count
GRS Enqueue	Step	463954	30	68	334632	15117
GRS Enqueue	System	39482	22	331	686069	1730
GRS Enqueue	Systems	139118	21	145	121152	6396

- System wide summary of GRS Latch and ENQ requests
- Statistics reported for Latch and ENQ requests:
  - Total and average contention time in milliseconds
  - The standard deviation of the total contention time in milliseconds
  - The total number of suspended requests.
- Additional statistics reported for GRS ENQ requests:
  - The scope of an ENQ request (STEP, SYSTEM, SYSTEMS)
  - The total number of GRS ENQ requests

Serialization delays: CMS Lock Details

Serialization Delay Details

Details of serialization delays for up to 20 address spaces

CMS Lock Details

Address Space ID	Job Name	Service Class Name	Service Class Period	CMS - Total Contention Time	CMS - Avg Contention Time	CMS - Total Contention Count	CMS - Contention Count with QLen>1	CMSEQDQ - Total Contention Time	CMSEQDQ - Avg Contention Time	CMSEQDQ - Total Contention Count	CMSEQDQ - Contention Count with QLen>1	CMSLatch - Total Contention Time	CMSLatch - Avg Contention Time	CMSLatch - Total Contention Count	CMSLatch - Contention Count with QLen>1	CMSSMF - Total Contention Time	CMSSMF - Avg Contention Time	CMSSMF - Total Contention Count	CMSSMF - Contention Count with QLen>1
0113	T015004	TSOHIGH	2	13	0.03	340	1	14	0.03	414	1					0	0.00	4	0
0103	T015042	TSOLOW	2	8	0.02	279	0	5	0.01	335	4					0	0.00	4	0
0104	T015149	TSOLOW	2	4	0.02	138	2	9	0.03	259	2								
0110	T015134	TSOLOW	2	2	0.02	77	1	8	0.01	439	8					0	0.00	3	0
00D8	CICS3A21	STCHI	1													0	0.00	31	0
00D6	CICS3A22	STCHI	1													0	0.00	21	0
0028	RMF	SYSSTC	1													0	0.00	1	0

- Serialization details on address space level for CMS lock types:
  - CMS Lock
  - CMS EnqueueDequeue Lock
  - CMS Latch Lock
  - CMS SMF Lock
- Contention data for work unit(s) running in the reported address space:
  - Total and average contention time in milliseconds
  - Number of times a work unit was suspended
  - Number of times a unit of work was suspended and there was at least one other unit of work suspended for the same lock.

Serialization delays: CML and Local Lock Details

CML and Local Lock Details

Address Space ID	Job Name	Service Class Name	Service Class Period	CML Lock Owner - Total Contention Time	CML Lock Owner - Avg Contention Time	CML Lock Owner - Total Contention Count	CML Lock Owner - Contention Count with QLen>1	Local Lock - Total Contention Time	Local Lock - Avg Contention Time	Local Lock - Total Contention Count	Local Lock - Contention Count with QLen>1	CML Lock Requestor - Total Contention Time	CML Lock Requestor - Avg Contention Time	CML Lock Requestor - Total Contention Count	CML Lock Requestor - Contention Count with QLen>1
000B	CONSOLE	SYSTEM	1	3323	0.06	47817	6032	698	0.05	12190	2675				
0018	IXGLOGR	SYSTEM	1	0	0.00	2	0	584	0.07	7593	122				
0113	T015004	TSOHIGH	2					542	0.03	16860	46	77	0.01	4765	81
0059	CATALOG	SYSTEM	1	44	0.01	2877	241	418	0.01	32046	1923				
0050	NFSCLNTC	SYSSTC	1	1212	0.02	41616	624	309	0.03	10152	338				
004D	NETVIEW	SYSSTC	1									797	0.07	10638	1181
0047	AUTOEMCS	STCLOW	1									776	0.07	11064	1243
00F4	T016030	TSOLOW	2									48	0.02	1708	28
00F6	T016082	TSOLOW	1									47	0.02	1605	30

- Serialization details on address space level for local locks
- Contention data for LOCAL lock requests and two views of CML lock requests:

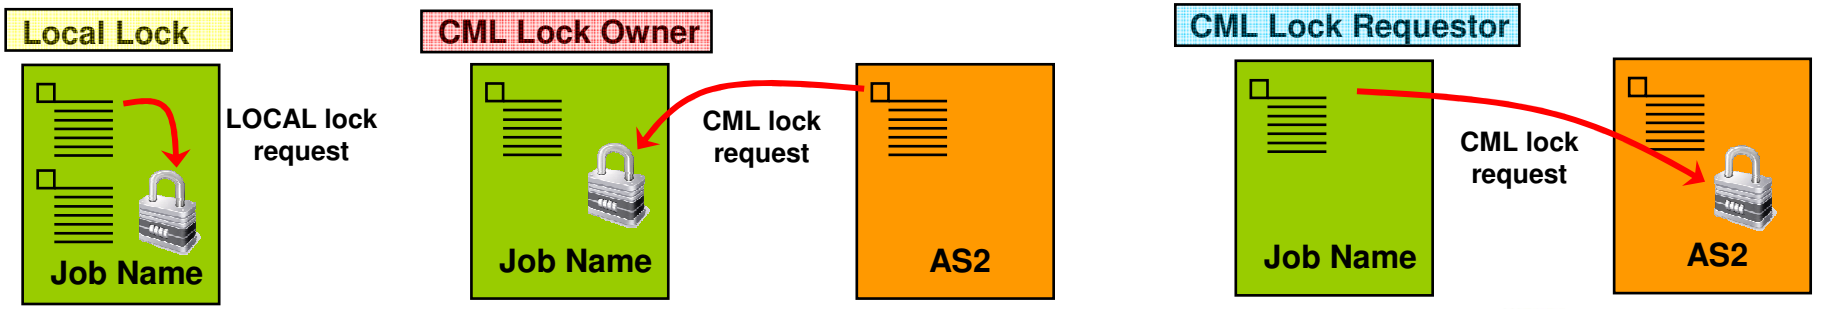

Local lock requested by work unit **running** in **same** AS

Local lock requested by work unit **running** in **another** AS

Local lock **form another** AS requested

Serialization delays: GRS Latch Details

GRS Latch Details

GRS Mode: STAR

Address Space ID	Job Name	Service Class Name	Service Class Period	Latch Set Creator - Total Contention Time	Latch Set Creator - Avg Contention Time	Latch Set Creator - Std Dev Contention Time	Latch Set Creator - Total Contention Count	Latch Requestor - Total Contention Time	Latch Requestor - Avg Contention Time	Latch Requestor - Std Dev Contention Time	Latch Requestor - Total Contention Count
0059	CATALOG	SYSTEM	1	11169	22	36	502	11169	22	36	502
0051	PFA	STCLOW	1	3264	3264		1	3264	3264		1
0018	IXGLOGR	SYSTEM	1	1672	1.81	16	921	5	1.25	1.12	4
0010	OMVS	SYSTEM	1	26	0.11	0.20	223				
0043	RRS	STCHI	1	20	20		1	20	20		1
000C	WLM	SYSTEM	1	5	0.31	0.31	16	5	0.33	0.15	15
00D8	CICS3A21	STCHI	1					856	1.90	15	449
00D4	CICS3A23	STCHI	1					519	2.22	19	233
00D6	CICS3A22	STCHI	1					289	1.24	13	233

- Serialization details on address space level for GRS Latch obtain requests
- Two views of GRS Latch obtain requests:
  - Latch Set Creator statistics:
    - Latch obtain requests against latch sets created by this address space
  - Latch Requestor statistics:
    - Latch obtain requests issued from this address space
- Reported statistics:
  - Total and average contention time in milliseconds
  - Standard deviation of the total contention time.
  - Number of times a latch obtain request was suspended

Latch Sets

1	
2	
x	

Serialization delays: GRS Enqueue Details

GRS Enqueue Details

GRS Mode: STAR

Address Space ID	Job Name	Service Class Name	Service Class Period	ENQ STEP - Total Contention Time	ENQ STEP - Avg Contention Time	ENQ STEP - Std Dev Contention Time	ENQ STEP - Request Count	ENQ STEP - Contention Count	ENQ SYSTEM - Total Contention Time	ENQ SYSTEM - Avg Contention Time	ENQ SYSTEM - Std Dev Contention Time	ENQ SYSTEM - Request Count	ENQ SYSTEM - Contention Count	ENQ SYSTEMS - Total Contention Time	ENQ SYSTEMS - Avg Contention Time	ENQ SYSTEMS - Std Dev Contention Time	ENQ SYSTEMS - Request Count	ENQ SYSTEMS - Contention Count
0059	CATALOG	SYSTEM	1	461440	30	68	136473	14975	541	6.07	13	97894	89	132295	21	147	98760	6033
004A	PFA7	OMVSLOW	2	1808	62	164	42	29										
011A	T015161	TSOLOW	1						10411	5205	69	8968	2					
0108	T015083	TSOLOW	1						10397	5198	86	5133	2					
0100	T015200	TSOLOW	1						5229	2614	3695	6523	2	24	4.80	3.35	134	5
0054	SMS	SYSSTC	1						1019	56	103	727	18	89	4.45	2.80	354	20
0008	SMSPDSE	SYSTEM	1						734	1.02	0.55	1297	717					
0104	T015149	TSOLOW	2						59	6.55	5.57	75473	9	83	4.61	2.09	2461	18
0113	T015004	TSOHIGH	2						58	4.14	3.01	164026	14	181	4.30	3.88	2191	42
00FC	T016094	TSOLOW	2											16	8.00	4.36	67	2

- Serialization details on address space level for GRS ENQ requests
- Grouped by GRS ENQ request scope: Step, System and Systems
  - Total and average contention time in milliseconds
  - Standard deviation of the total contention time.
  - Total number of GRS ENQ requests
  - Number of times a GRS ENQ request was suspended

Work Unit Queue Distribution

- Problem
  - RMF reporting for CPU contention was based on address spaces level
  - Did not consider multiple work units (WEBs) within one address space
  - Performance analysts need statistics about the CPU contention on WEB granularity
- Solution
  - The RMF CPU Activity reporting is enhanced
- Benefit
  - The enhanced CPU Activity reporting helps the customer to obtain information about the In-Ready distribution based on WEBs and the number of work units per CPU type

Work Unit Queue Distribution: Monitor I CPU Report

SYSTEM ADDRESS SPACE AND WORK UNIT ANALYSIS

---------NUMBER OF ADDRESS SPACES---------			
QUEUE TYPES	MIN	MAX	AVG
IN	550	1,008	594.8
IN READY	4	438	22.7
OUT READY	0	1	0.0
OUT WAIT	0	0	0.0
LOGICAL OUT RDY	0	628	10.3
LOGICAL OUT WAIT	178	634	589.4
ADDRESS SPACE TYPES			
BATCH	281	284	282.0
STC	708	765	736.4
TSO	97	98	97.9
ASCH	0	1	0.0
OMVS	43	97	68.0

Work unit count on CPU type level

---------NUMBER OF WORK UNITS----------			
CPU TYPES	MIN	MAX	AVG
CP	444	888	555.5
AAP	22	33	28.8
IIP	0	0	0.0

----------DISTRIBUTION OF IN-READY WORK UNIT QUEUE----------

NUMBER OF WORK UNITS	(%)	0 10 20 30 40 50 60 70 80 90 100
<= N	55.5	>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
= N + 1	4.4	>>>
= N + 2	4.0	>>>
= N + 3	3.7	>>
<= N + 5	7.1	>>>>
<= N + 10	14.7	>>>>>>>>
<= N + 15	8.0	>>>>>
<= N + 20	1.9	>
<= N + 30	0.2	>
<= N + 40	0.0	
<= N + 60	0.0	
<= N + 80	0.0	
<= N + 100	0.0	
<= N + 120	0.0	
<= N + 150	0.0	
> N + 150	0.0	

16 Buckets representing the work unit count with regard to the number of online processors

N = NUMBER OF PROCESSORS ONLINE UNPARKED (22.4 ON AVG)

Work Unit Queue Distribution: Spreadsheet Reporter

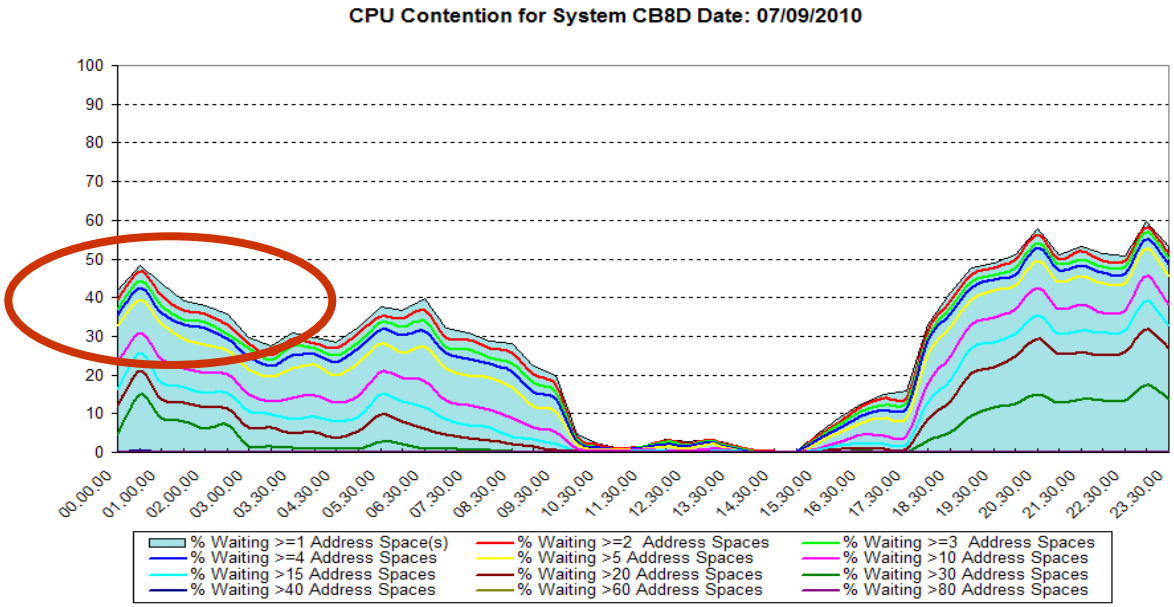

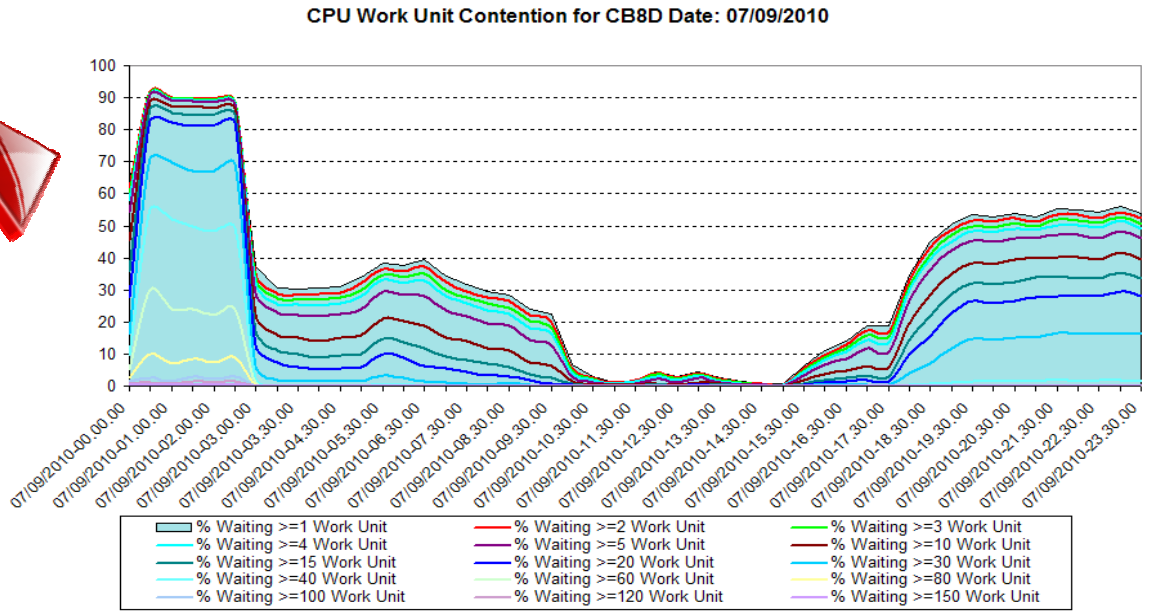

IBM z/OSMF Management Facility

z/OS application, browser access

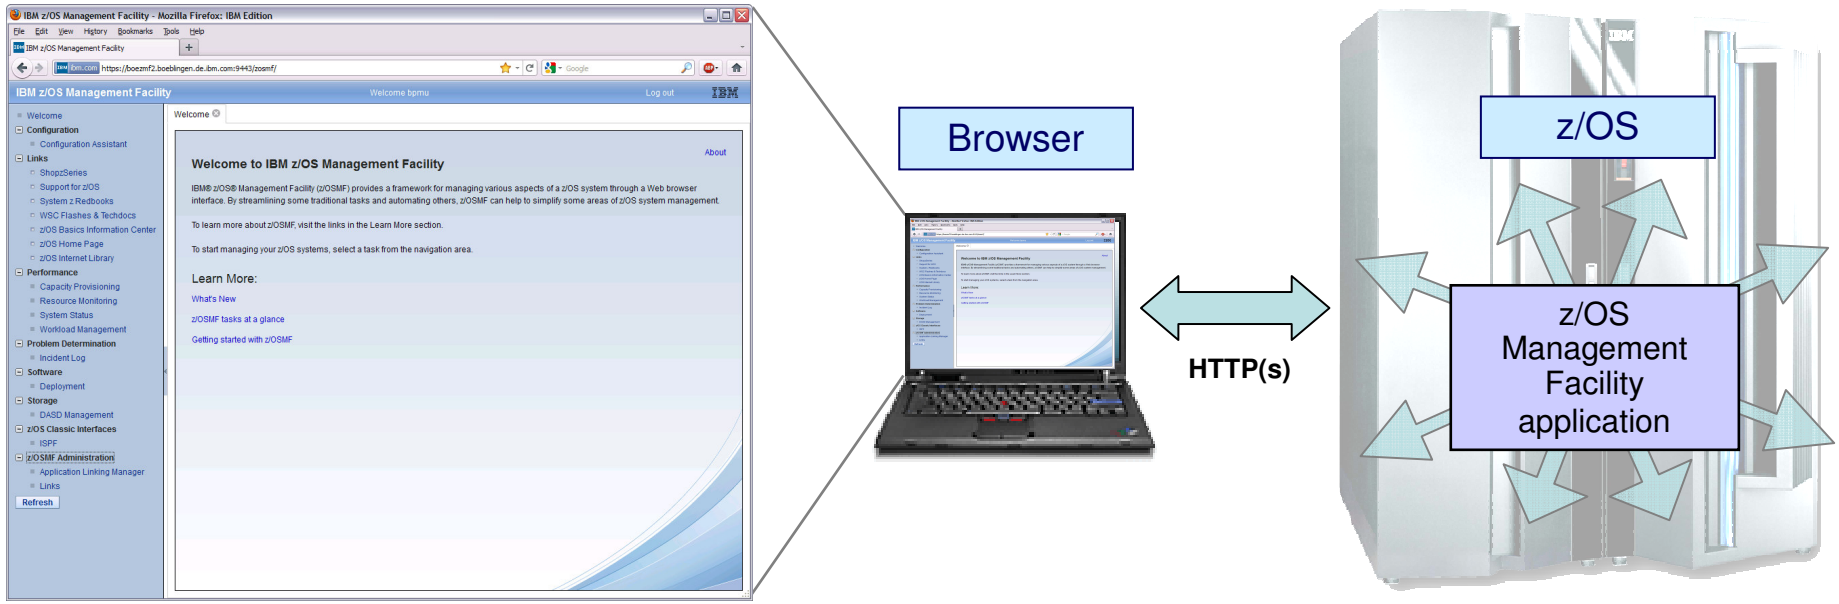

- *z/OS Management Facility is a Web 2.0 application on z/OS*
  - Manages z/OS from z/OS
  - Browser communicates with z/OSMF via secure connection, anywhere, anytime

IBM z/OSMF Management Facility

Resource Monitoring – System Status Task

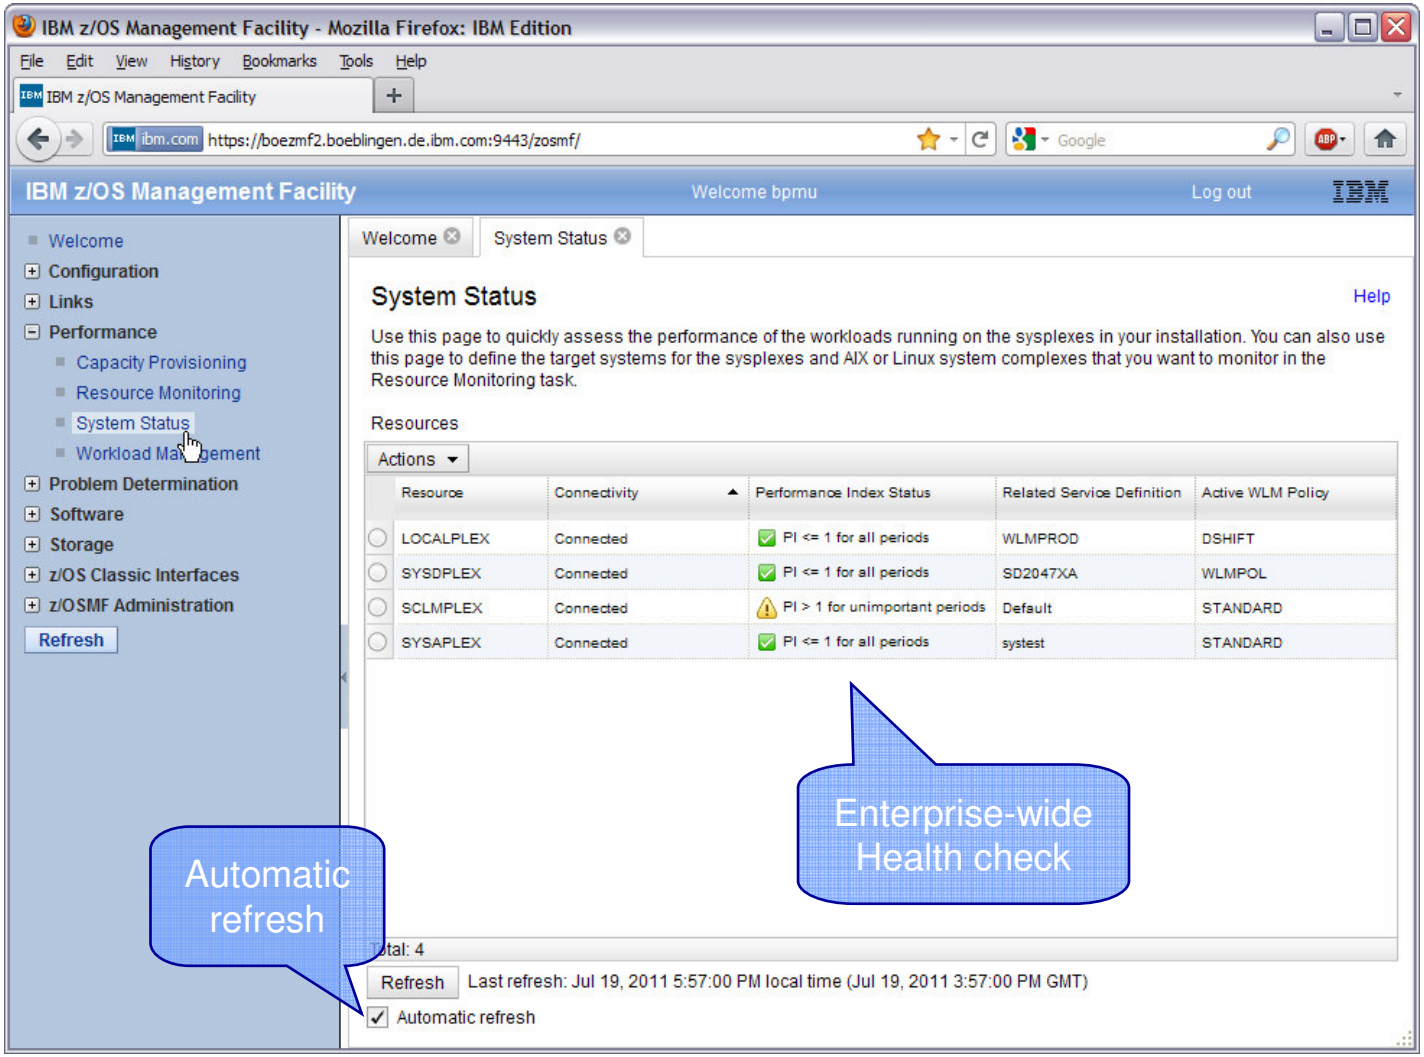

IBM z/OSMF Management Facility …

Resource Monitoring – Monitoring Dasboards

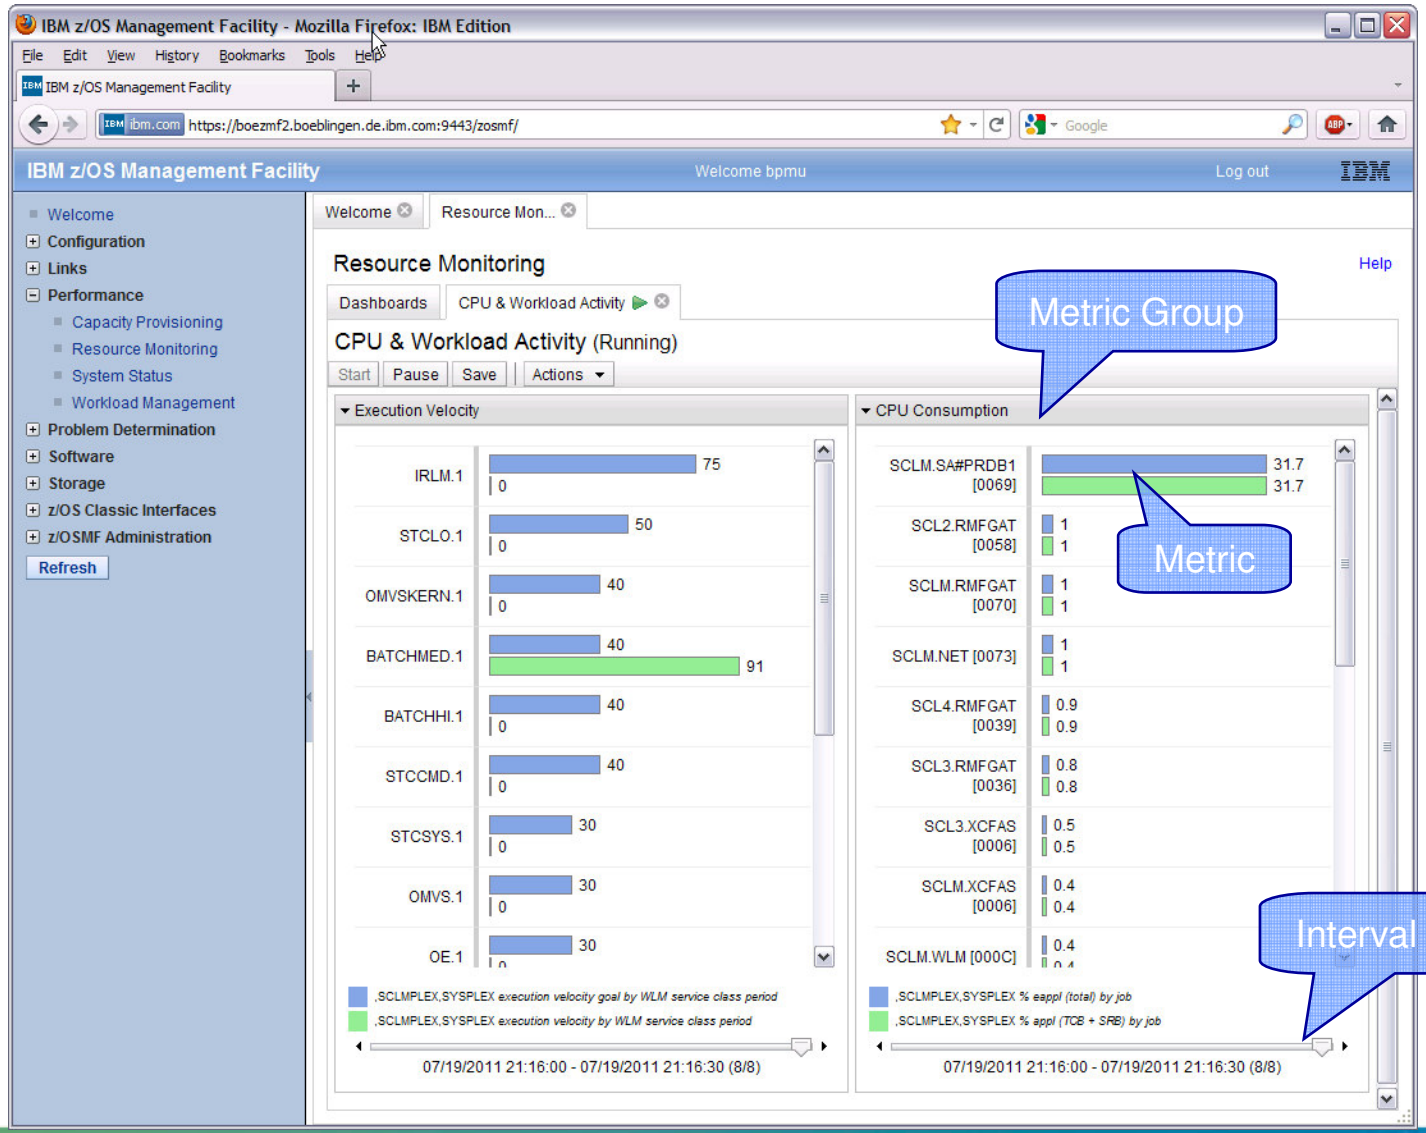

RMF XP – What is it ?

- RMF XP is the solution for
  Cross Platform Performance Monitoring

- RMF XP supports the Operating Systems running on
  - x Blades
  - p Blades

✓ Available with z/OS V1R13 RMF and
z/OS V1R12 RMF (APAR OA36030)

RMF XP – Basic Idea

- The Common Information Model (aka CIM) instrumentation is available for almost all operating systems of this planet
- RMF has the infrastructure already in place to
  - combine performance data from multiple systems to a Sysplex wide view
  - display performance data by means of state-of-the-art graphical frontends

Isn‘t it a good idea to bring those well-proven things together ?

We thought it is and we created the RMF XP

RMF XP – Component Overview

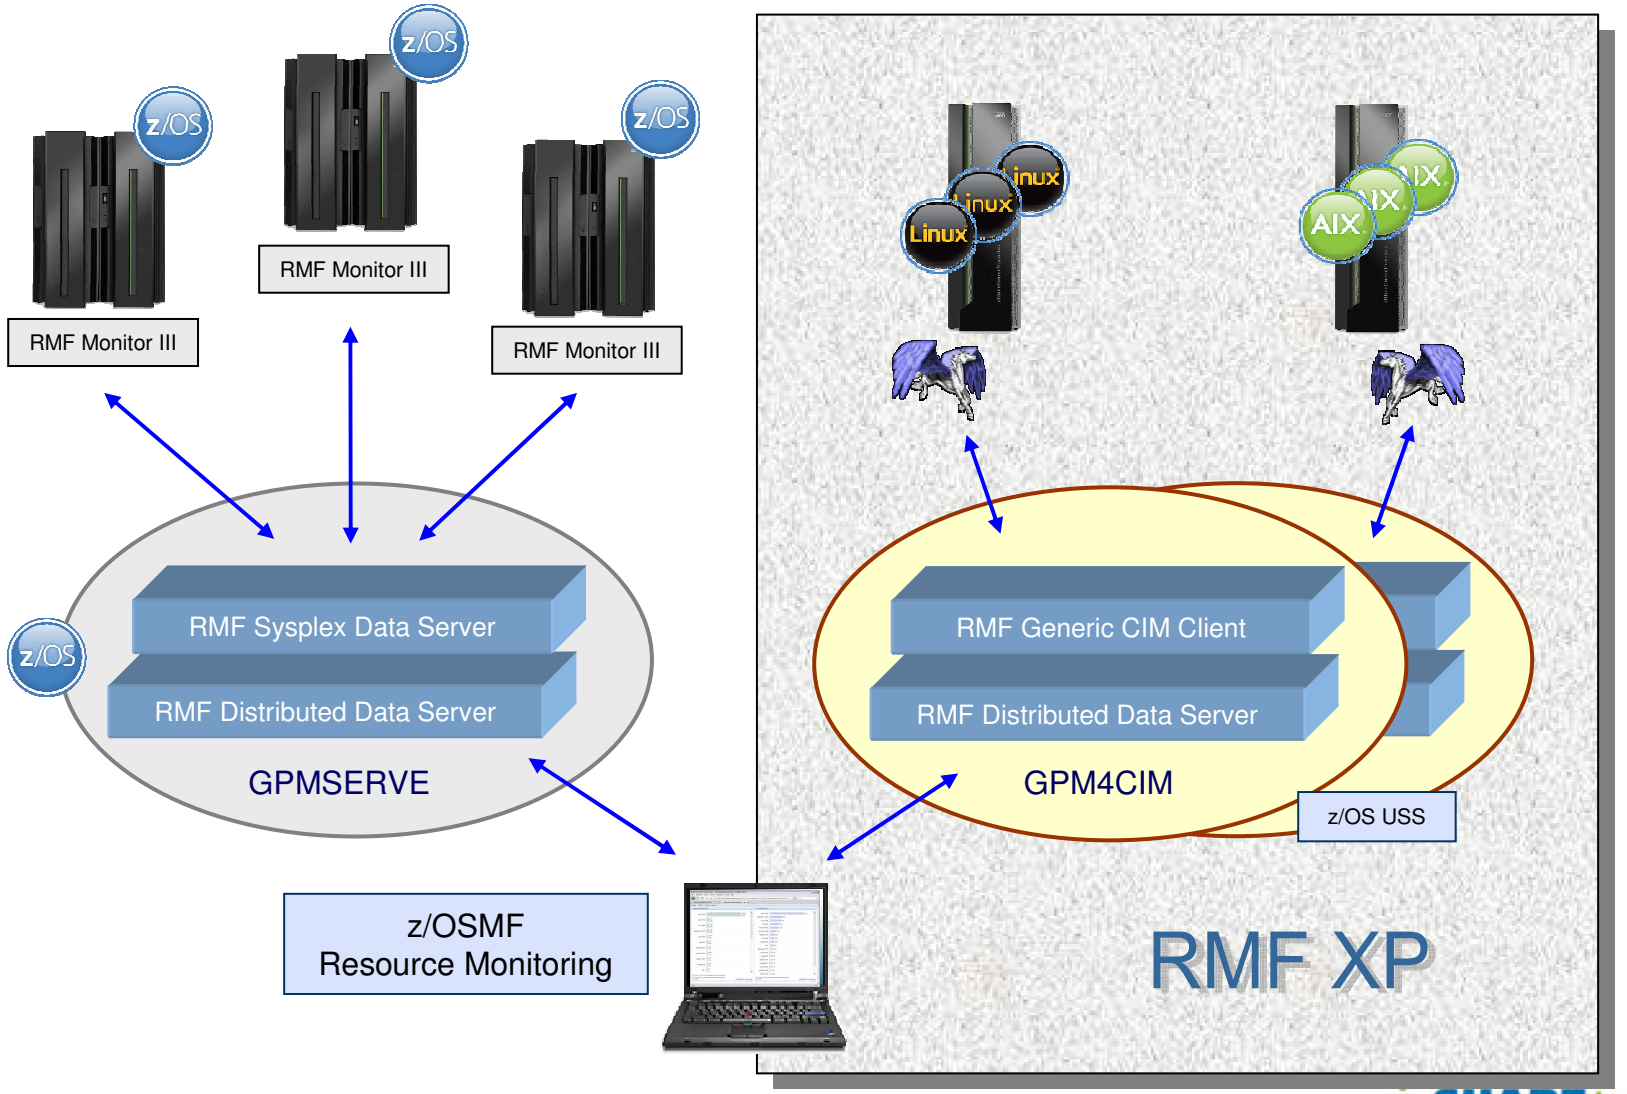

RMF XP – Linux Topology

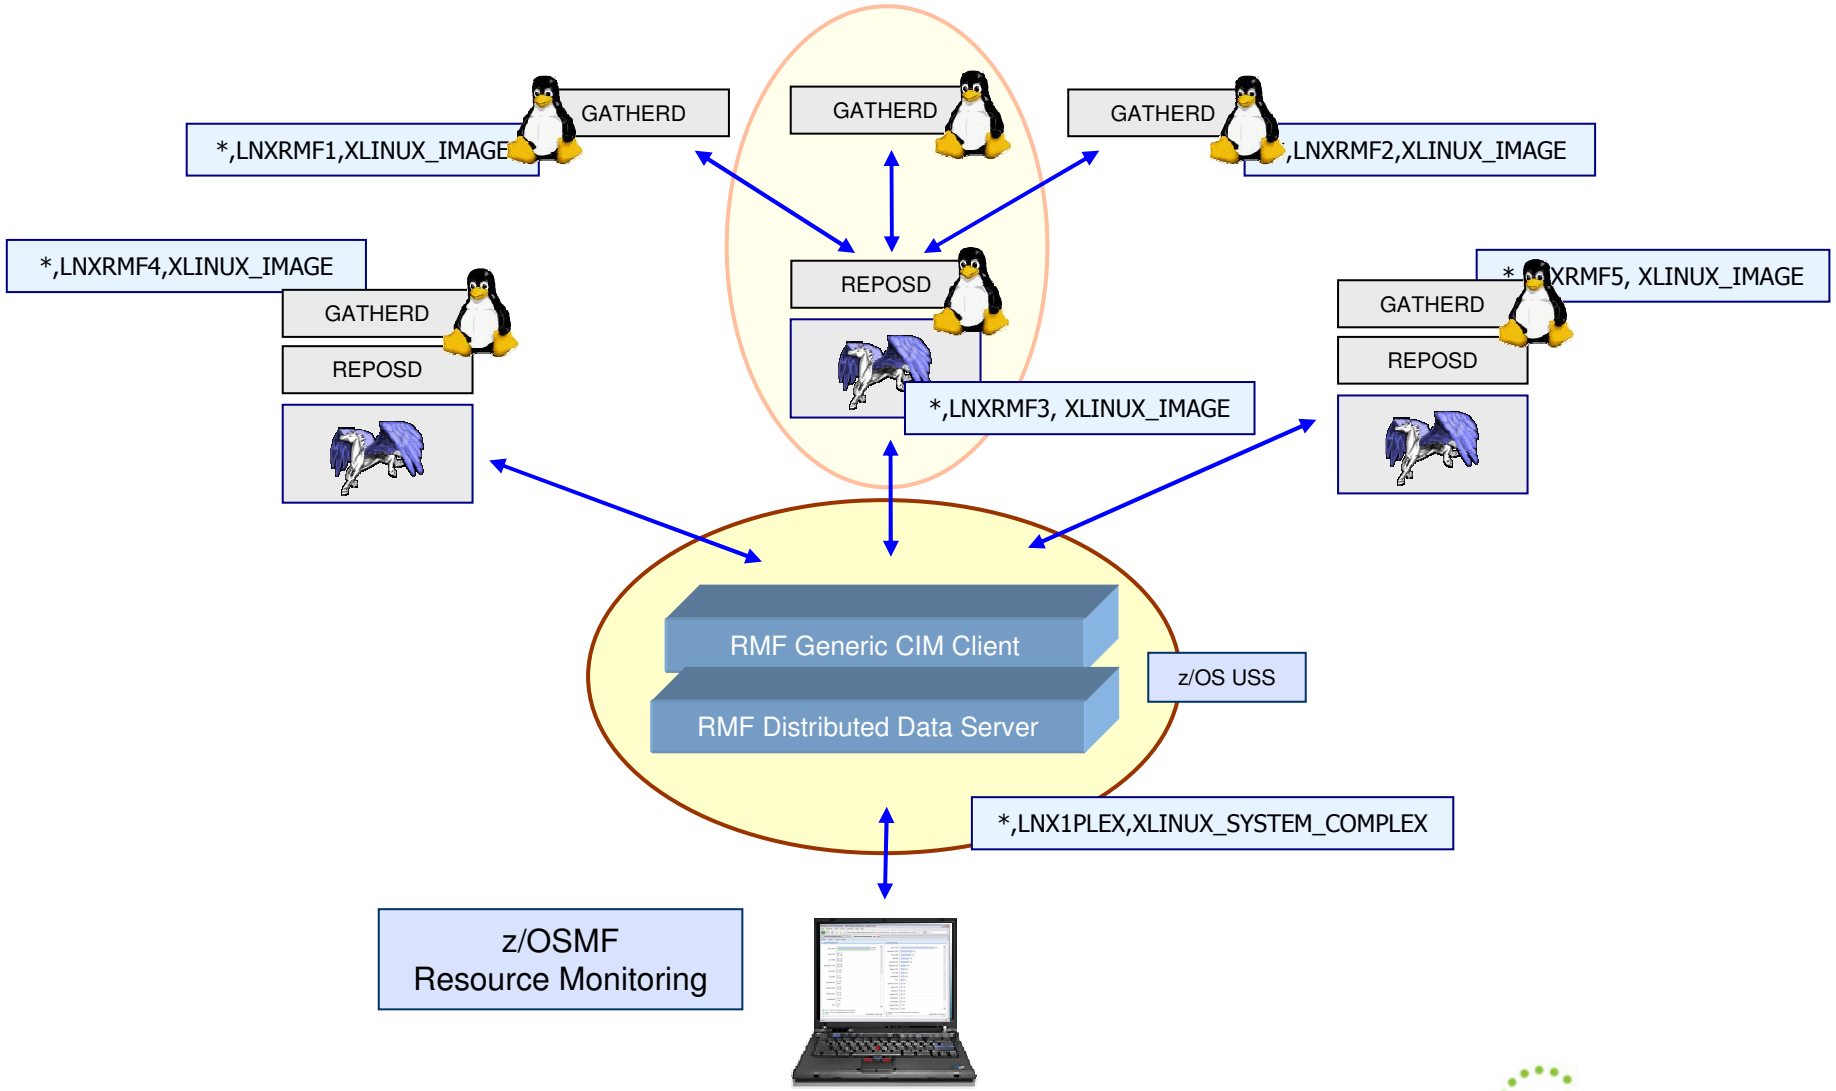

RMF XP – AIX Topology

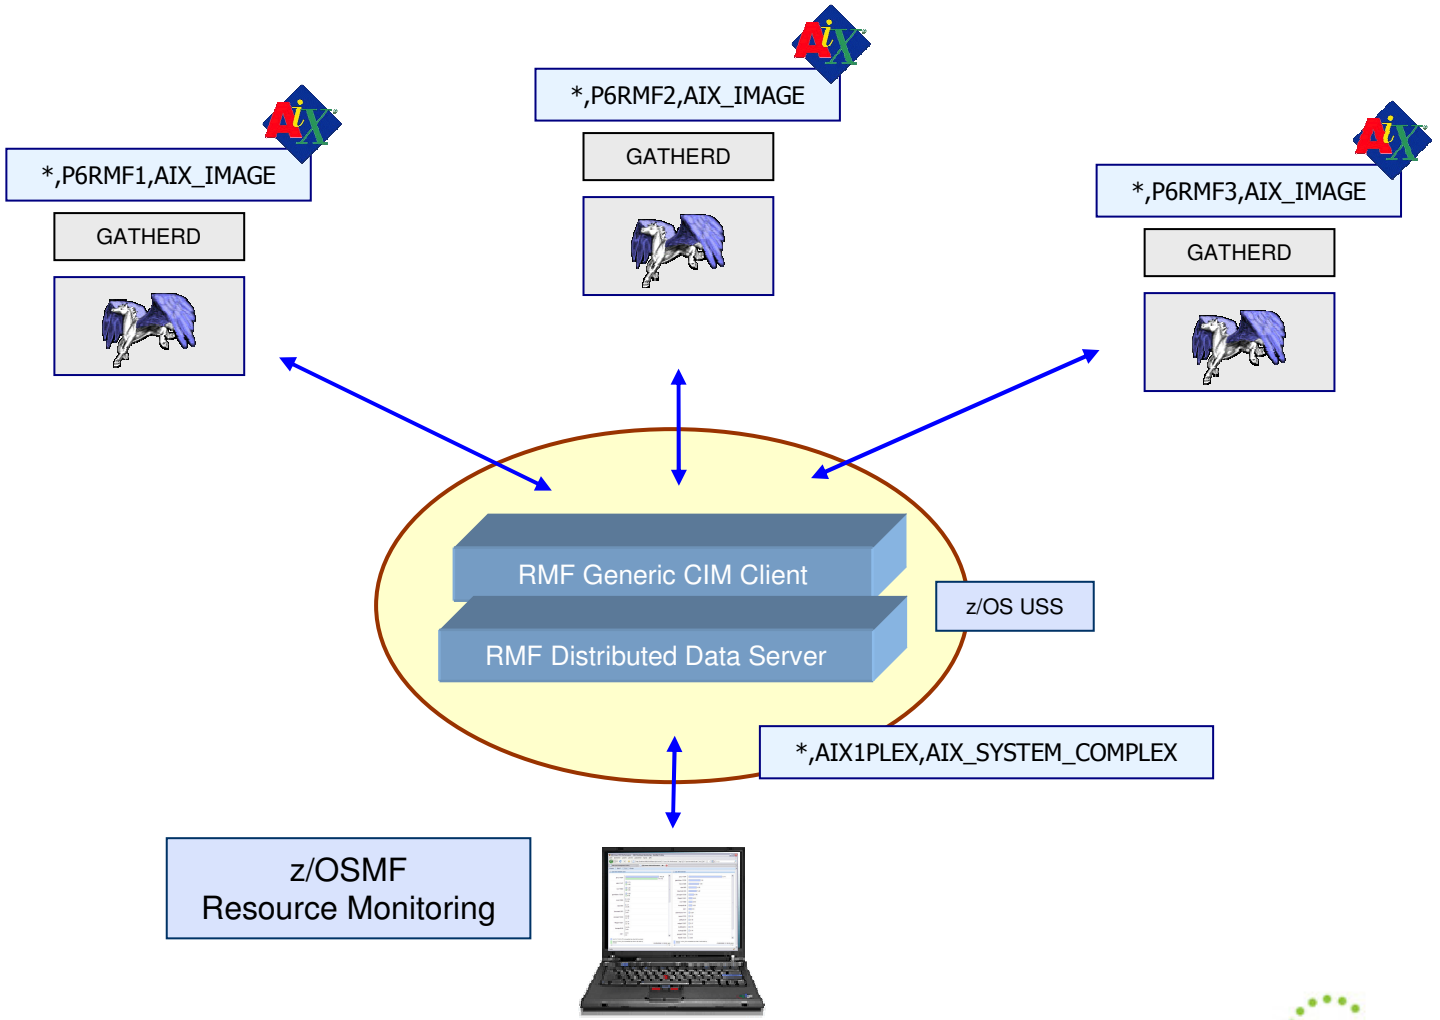

RMF XP - Invocation

- Started Task: SYS1.PROCLIB(GPM4CIM)
- Runs in USS Environment via BPXBATCH
- Multiple instances can run in parallel: one STC per platform
  - S GPM4CIM.GPM4A,OS=A
  - S GPM4CIM.GPM4X,OS=X
  - S GPM4CIM.GPM4Z,OS=Z

```
//GPM4CIM  PROC OS=X
//STEP1    EXEC PGM=BPXBATCH,TIME=NOLIMIT,REGION=0M,
//       PARM='PGM /usr/lpp/gpm/bin/gpm4cim cfg=/etc/gpm/gpm4&OS..cfg'
//STDENV   DD   PATH='/etc/gpm/gpm4cim.env'
//STDOUT   DD   PATH='/var/gpm/logs/gpm4cim&OS..out',
//         PATHOPTS=(OWRONLY,OCREAT,OTRUNC),
//         PATHMODE=(SIRUSR,SIWUSR,SIRGRP)
//STDERR   DD   PATH='/var/gpm/logs/gpm4cim&OS..trc',
//         PATHOPTS=(OWRONLY,OCREAT,OTRUNC),
//         PATHMODE=(SIRUSR,SIWUSR,SIRGRP)
//SYSPRINT DD   SYSOUT=*
//SYSOUT   DD   SYSOUT=*
//         PEND
```

RMF XP - Configuration Definition

```
//GPM4CIM  PROC OS=A
//STEP1    EXEC PGM=BPXBATCH,TIME=NOLIMIT,REGION=0M,
//       PARM='PGM /usr/lpp/gpm/bin/gpm4cim cfg=/etc/gp /gpm4&OS..cfg'
//STDENV   DD   PATH='/etc/gpm/gpm4cim.env'
//STDOUT   DD   PATH='/var/gpm/logs/gpm4cim&OS..out',
//         PATHOPTS=(OWRONLY,OCREAT,OTRUNC),
//         PATHMODE=(SIRUSR,SIWUSR,SIRGRP)
//STDERR   DD   PATH='/var/gpm/logs/gpm4cim&OS..trc',
//         PATHOPTS=(OWRONLY,OCREAT,OTRUNC),
//         PATHMODE=(SIRUSR,SIWUSR,SIRGRP)
//SYSPRINT DD   SYSOUT=*
//SYSOUT   DD   SYSOUT=*
//         PEND
```

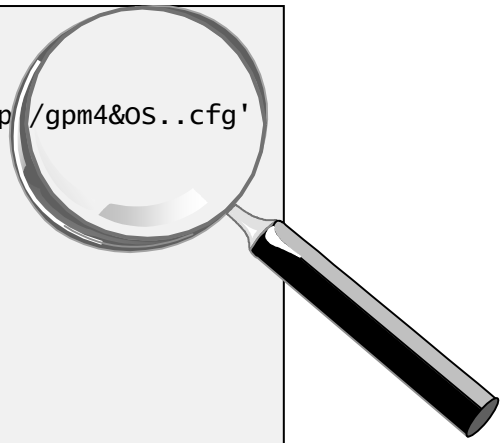

```
MAXSESSIONS_HTTP(20)                          /* MaxNo of concurrent HTTP requests */
HTTP_PORT(8805)                               /* Port number for HTTP requests     */
HTTP_ALLOW(*)                                 /* Mask for hosts that are allowed   */
HTTP_NOAUTH(*)                                /* No server can access without auth.*/
INTERVAL(300)                                /* Length of the monitoring interval */
AIX_COMPLEX(WEBPLEX)                          /* Name of system complex            */
AIX_IMAGE(p6rmf1.boeblingen.de.ibm.com:5988)  /* Hostname of member                */
AIX_IMAGE(p6rmf2.boeblingen.de.ibm.com:5988)
```

RMF XP – Platform specific Resource Models

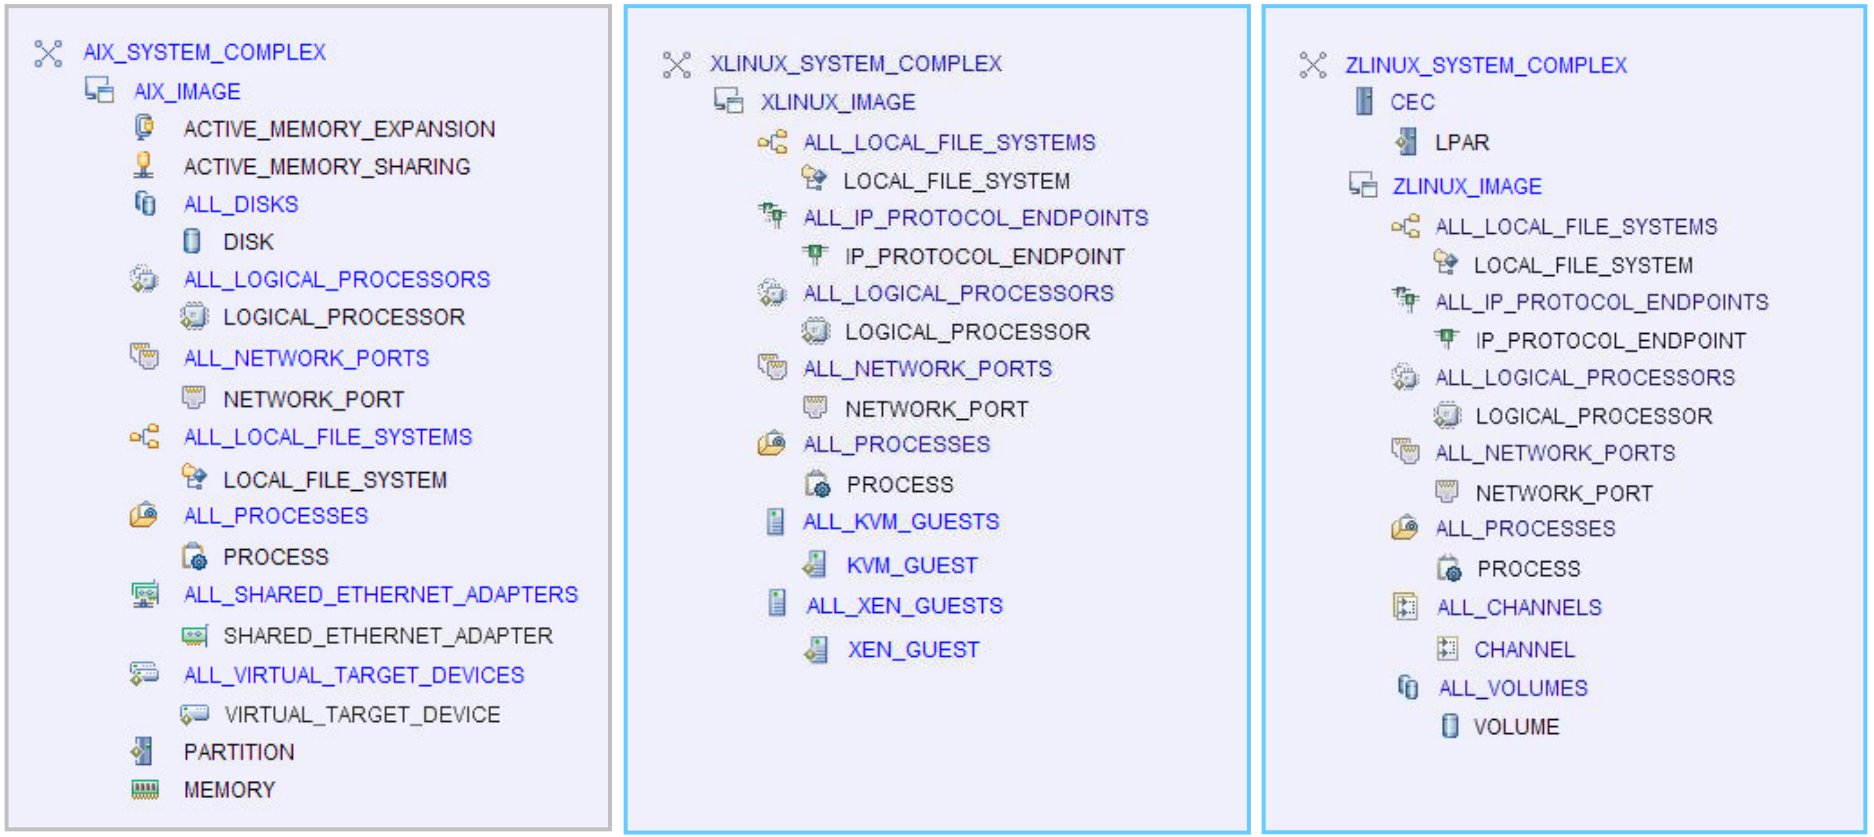

RMF XP – Security

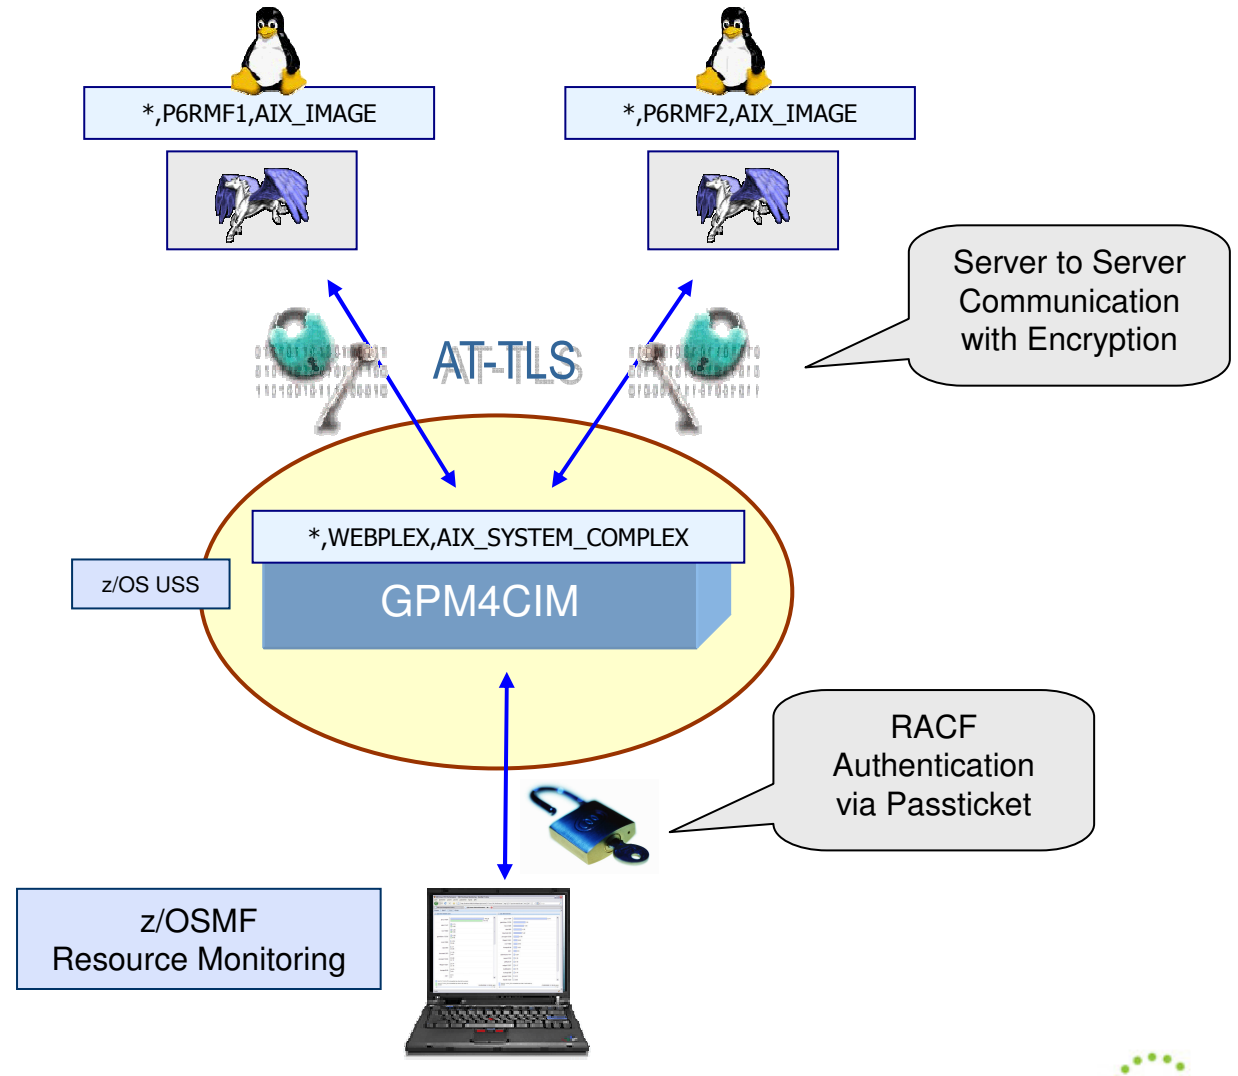

RMF XP - zIIP Exploitation

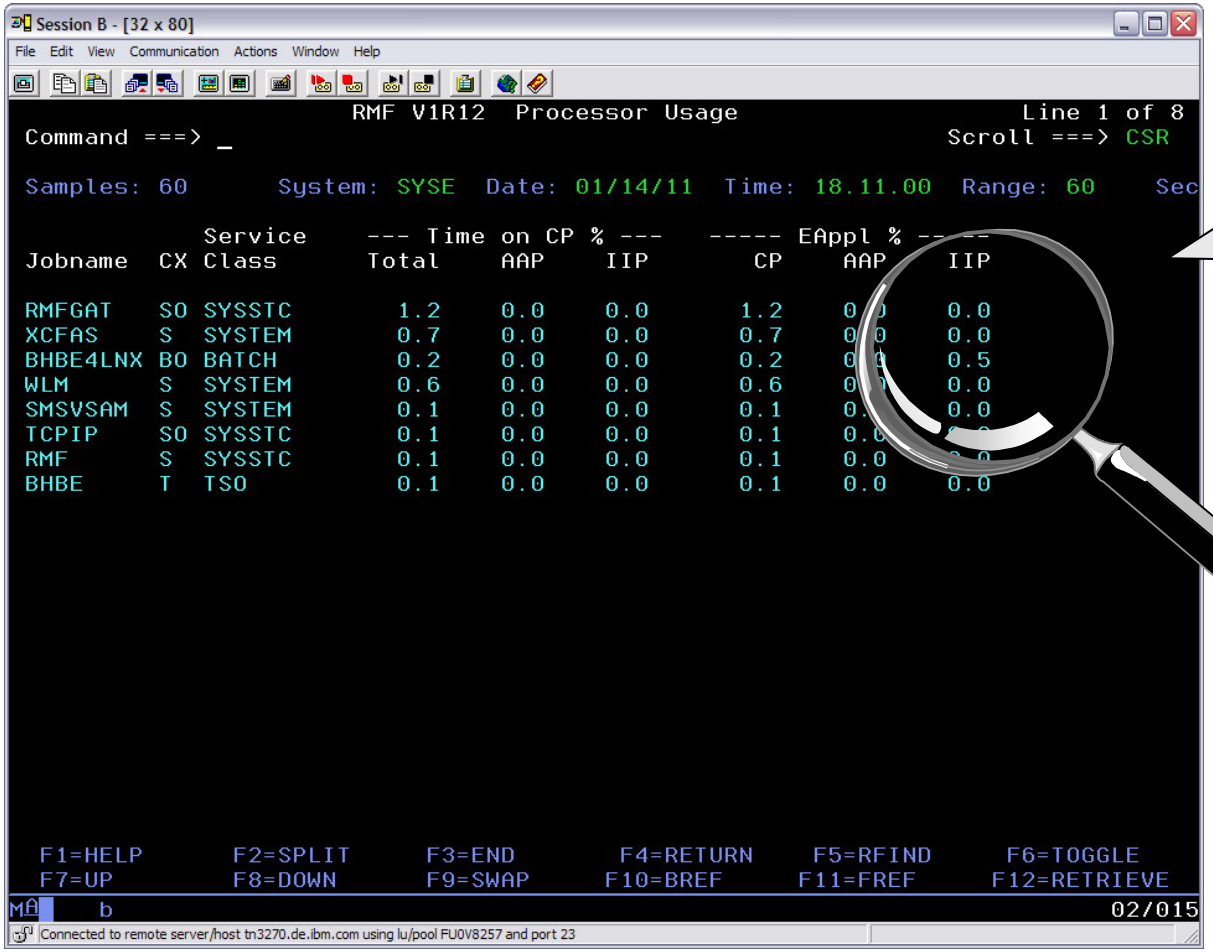

Up to 70% CPU utilization can be offloaded to zIIPs !

RMF XP – Resource Tree

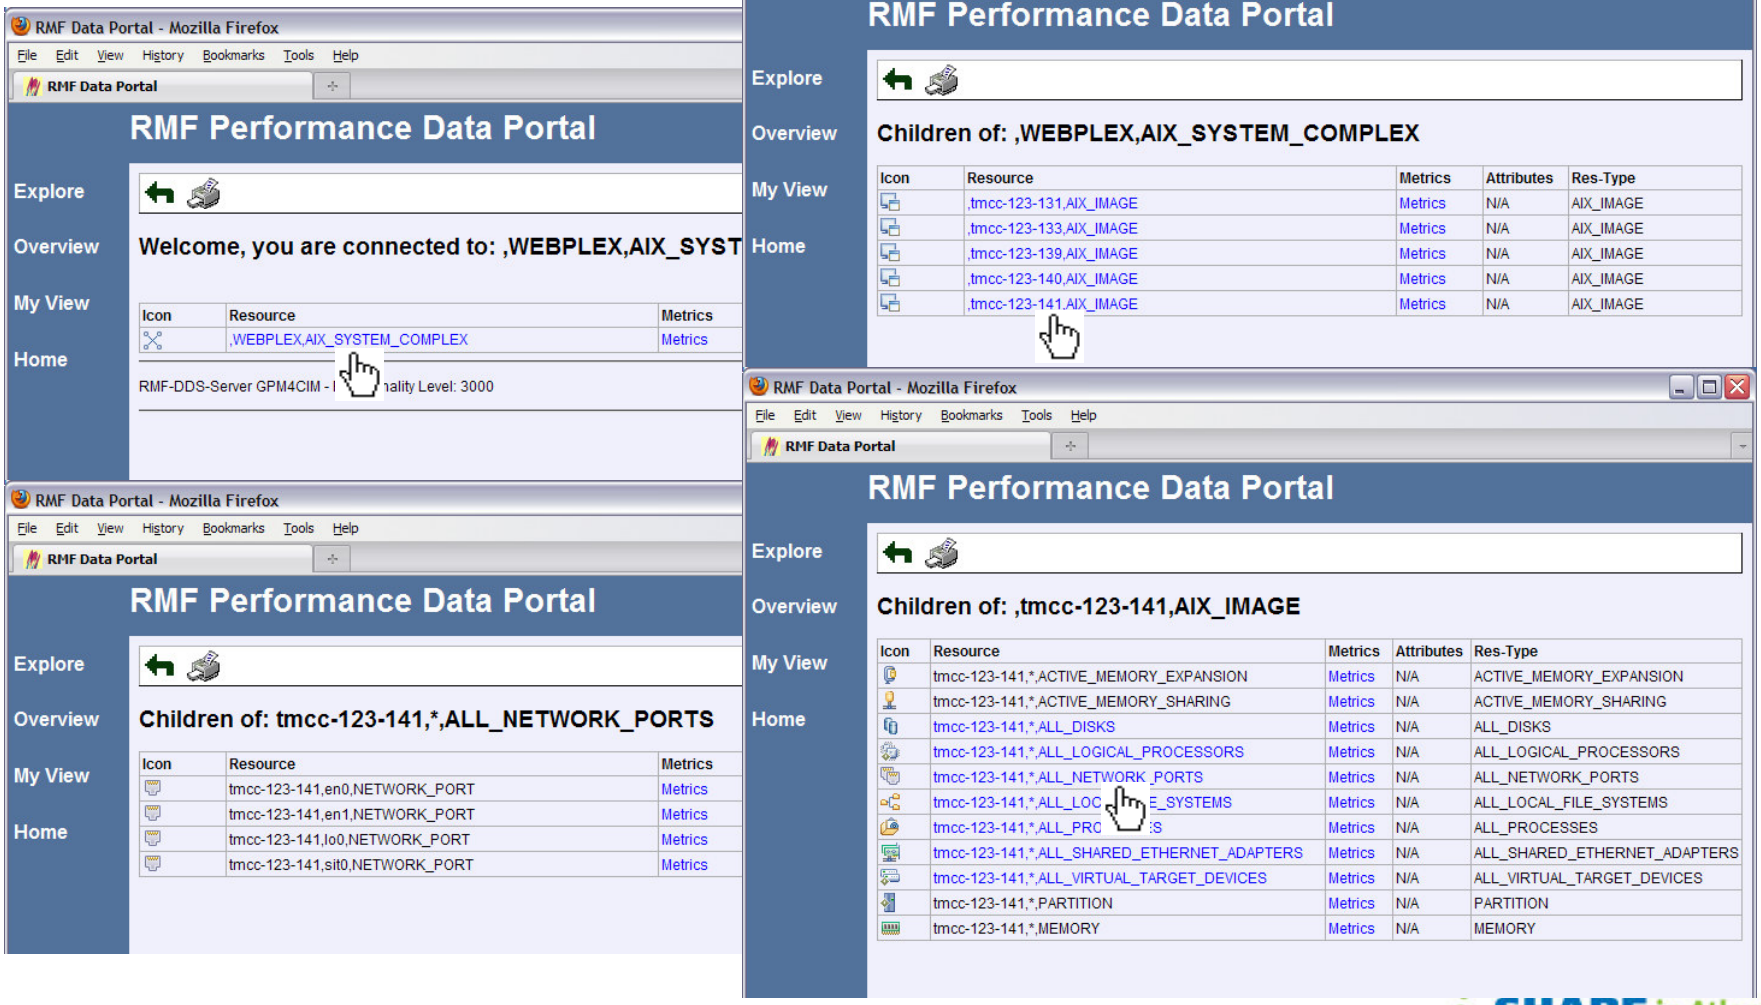

RMF XP – Metrics

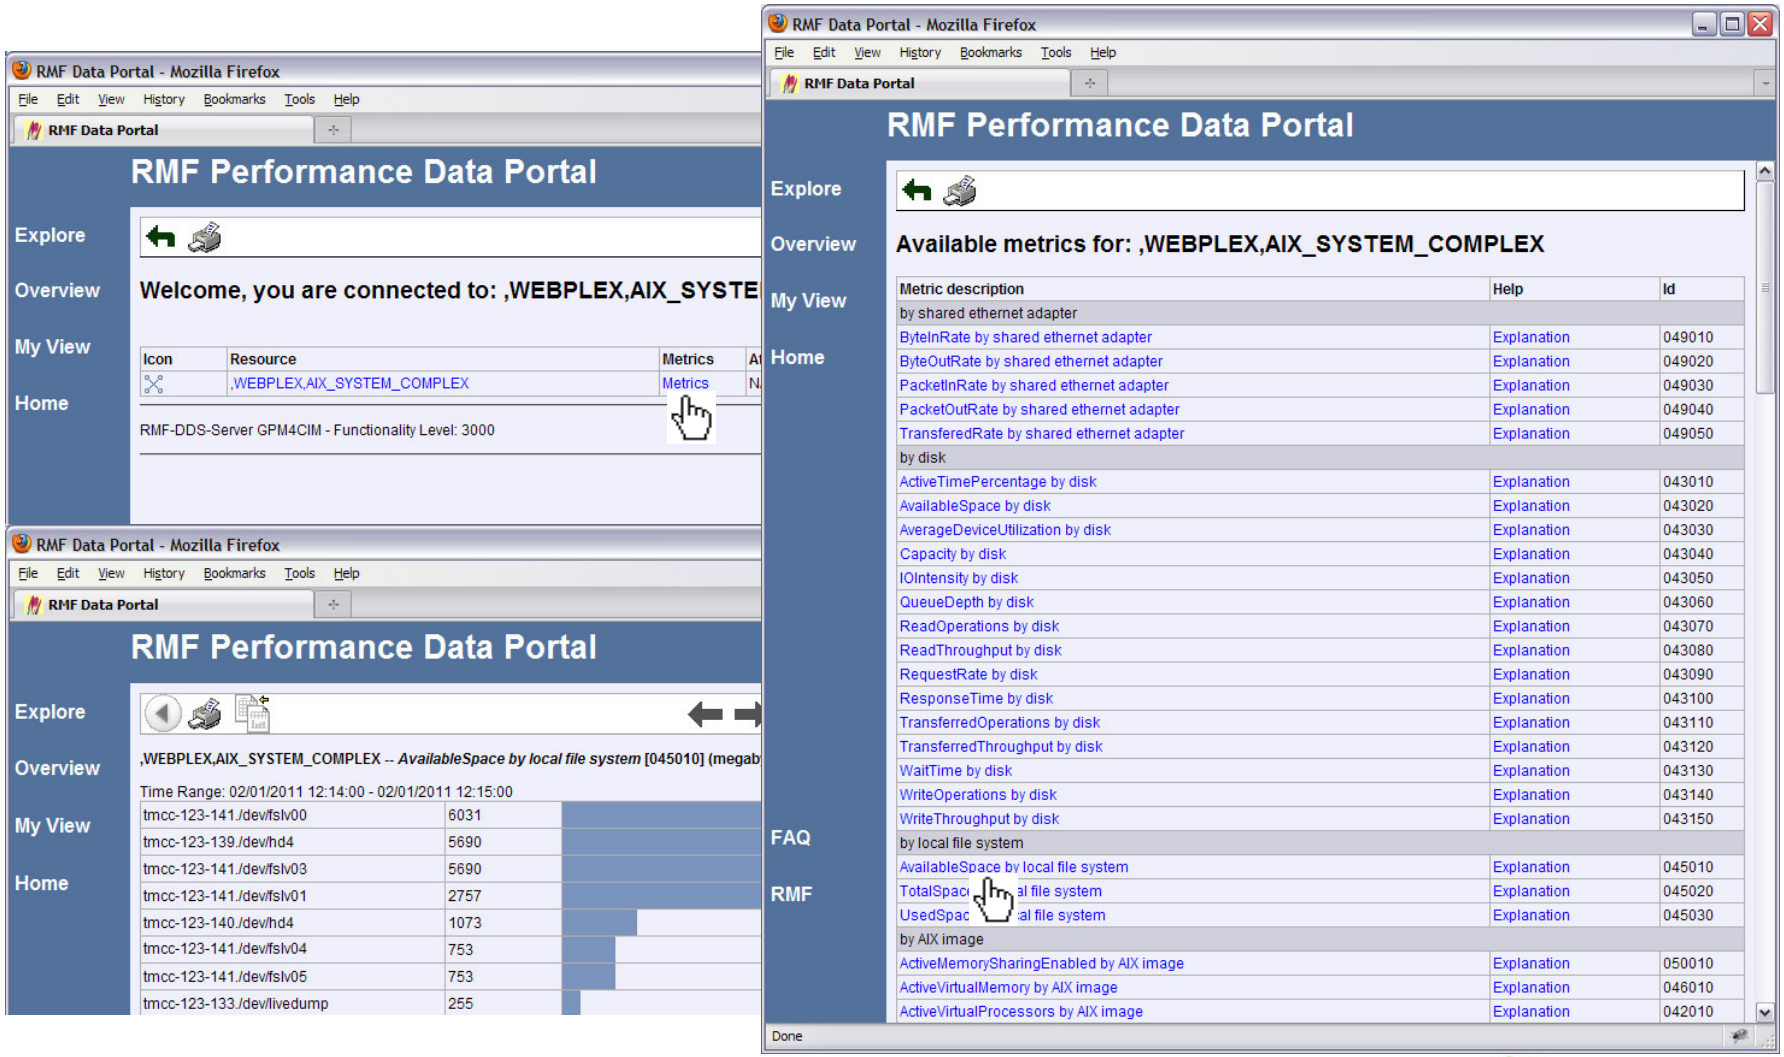

RMF XP – Metric Values

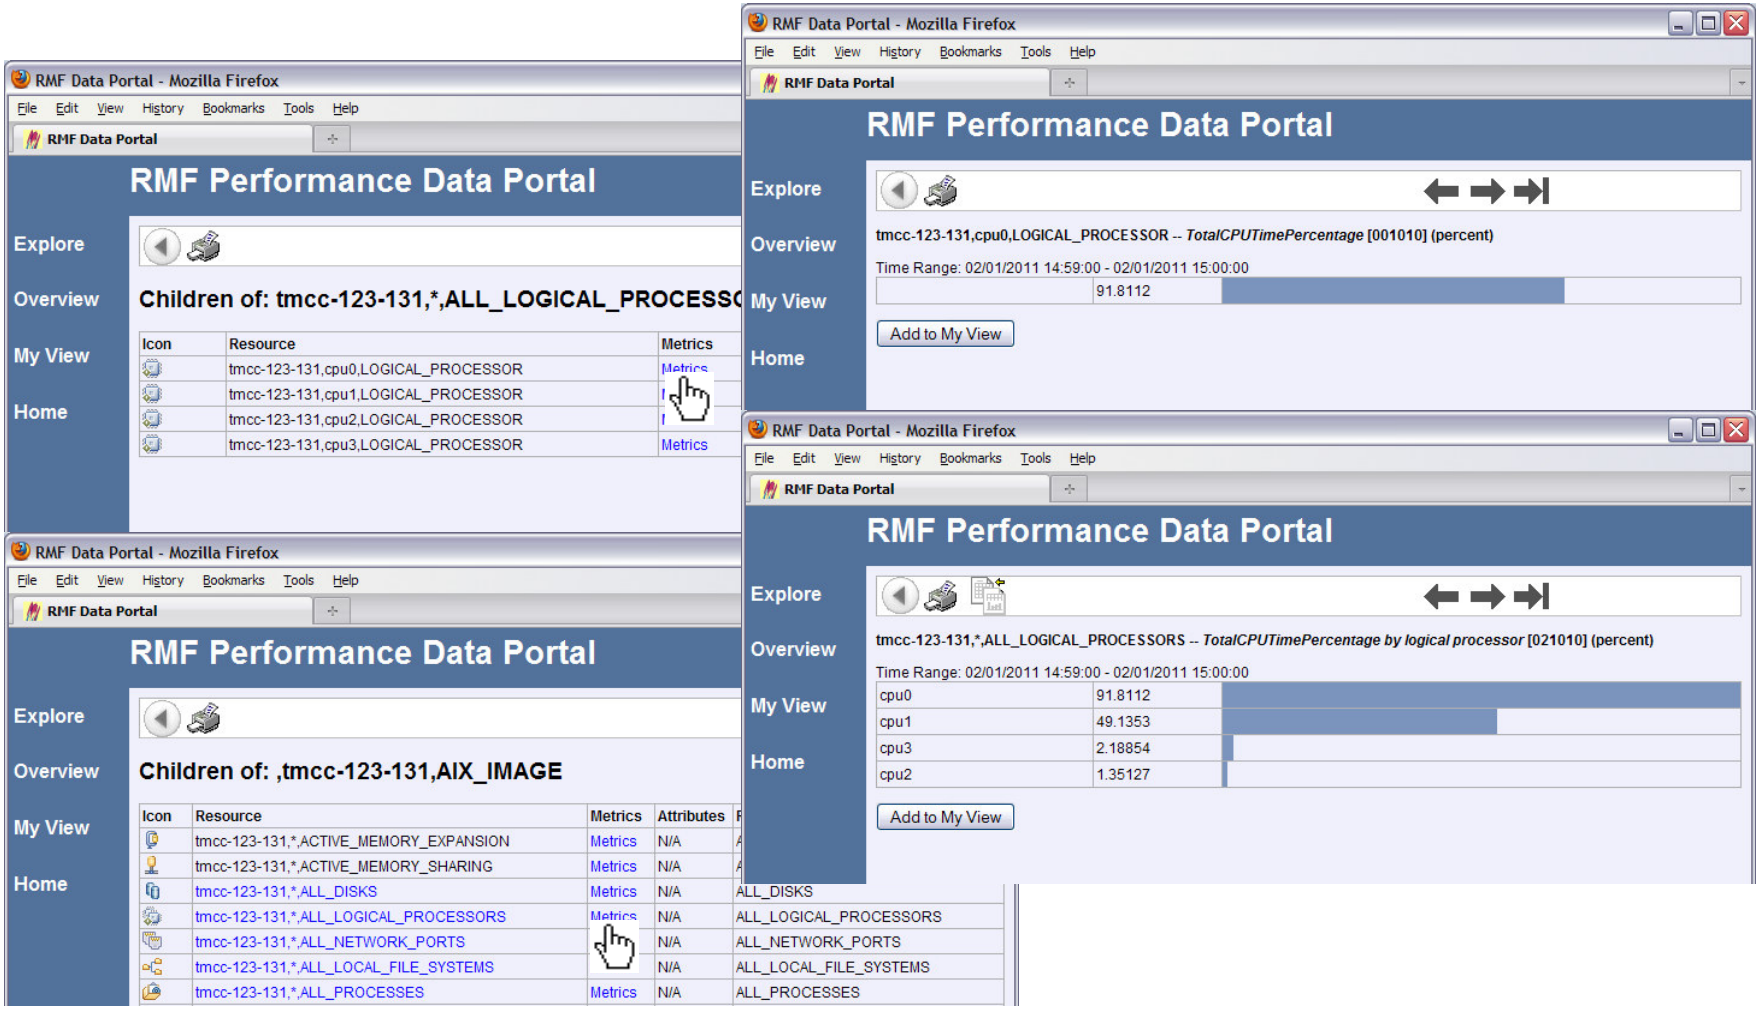

RMF XP – Metric Scope

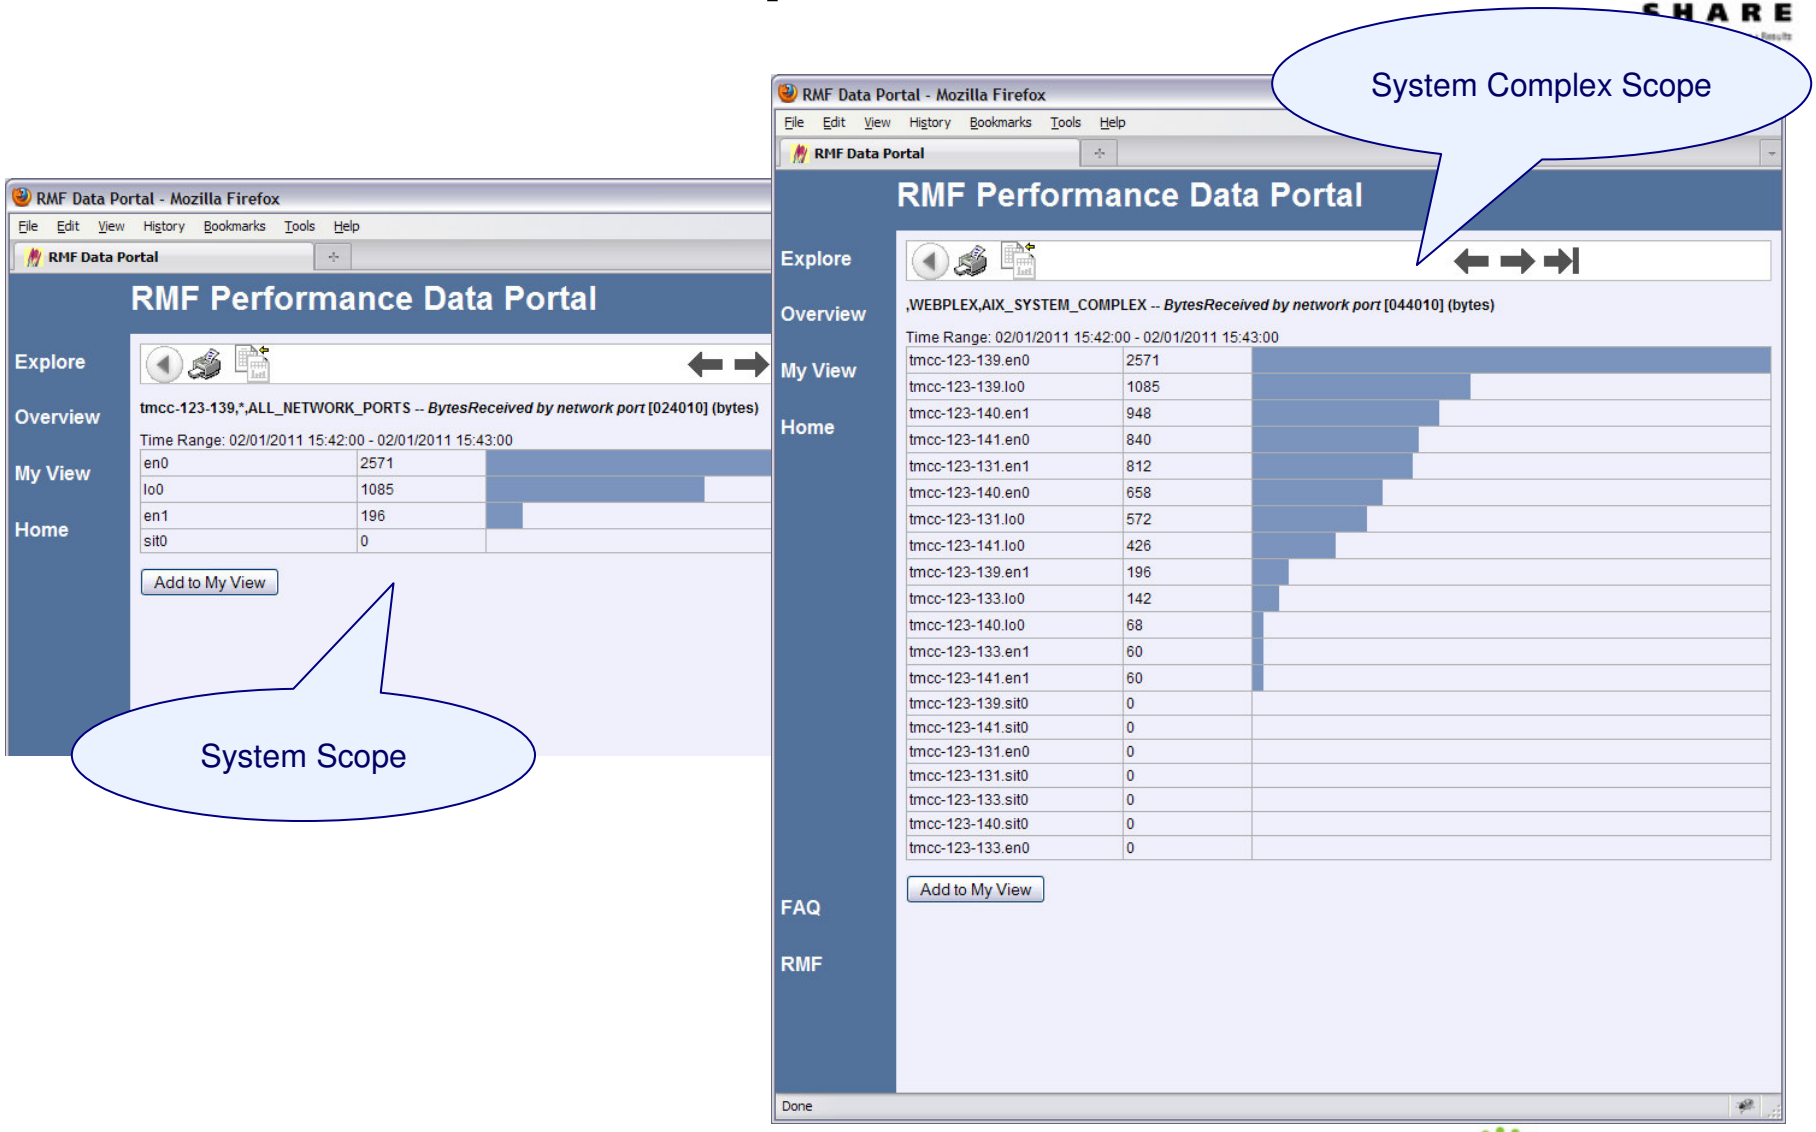

RMF XP – z/OSMF Integration

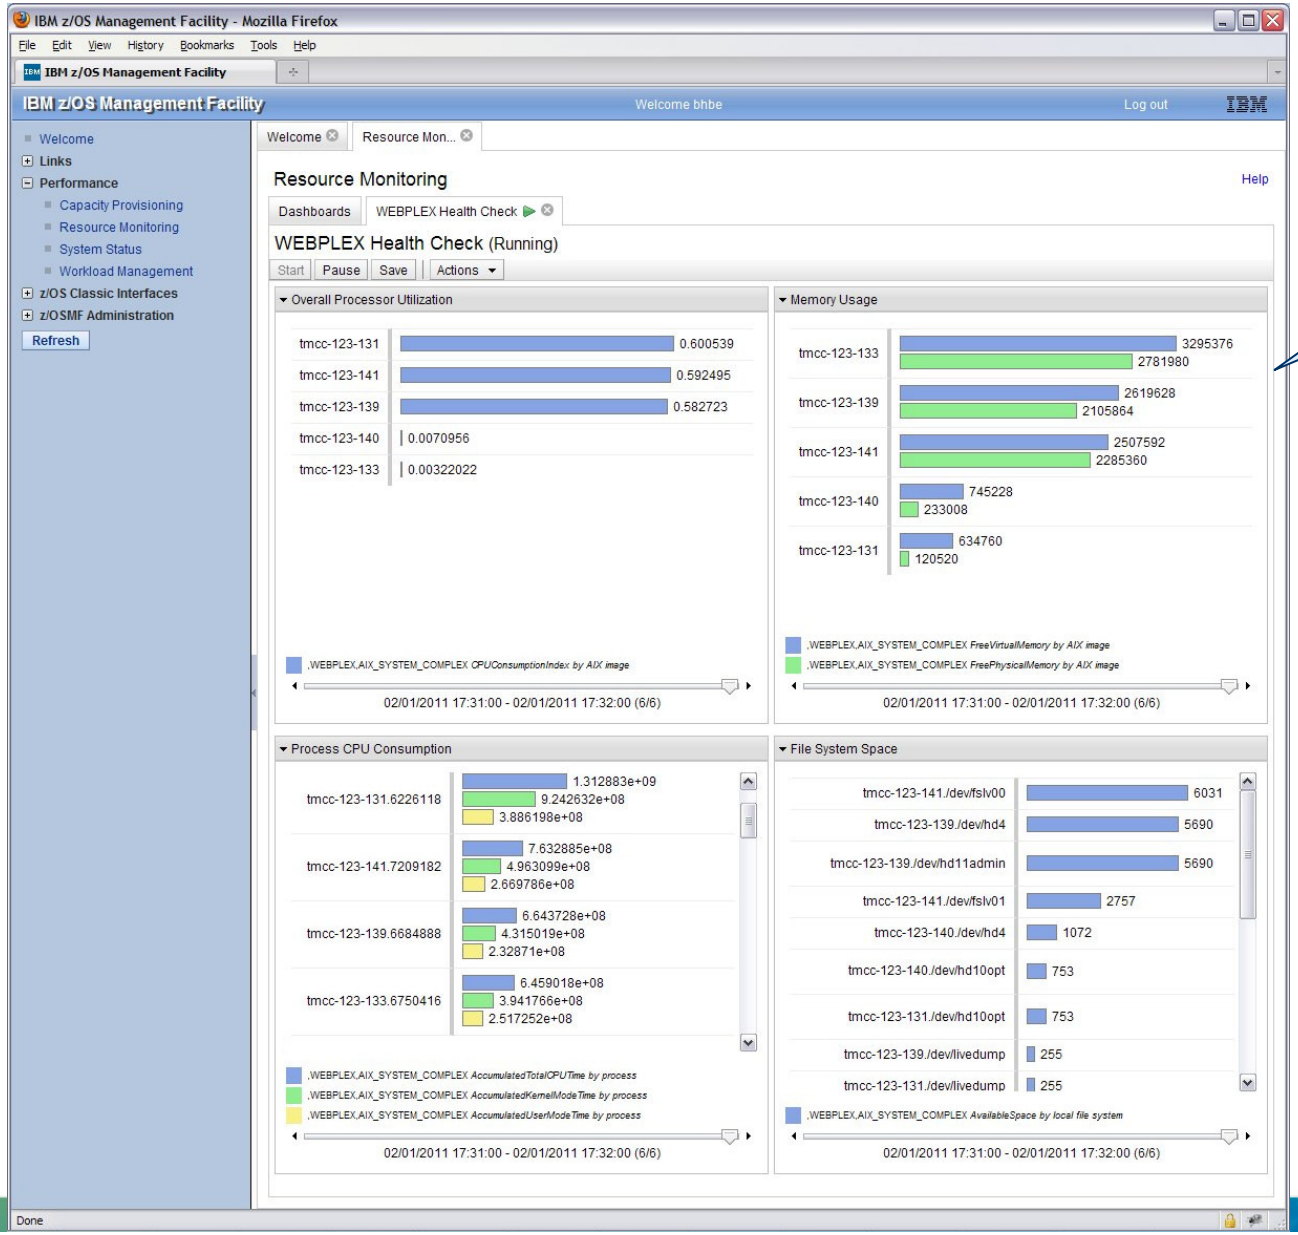

Health Check:

- ✓ Processor
- ✓ Memory
- ✓ Filesystem

RMF XP - Summary

- Seamless performance monitoring solution for z/OS and distributed platforms
- z/OS as management platform for distributed environments
- Easy to setup, almost no customization needed
- Two graphical frontends
  - Instant access via web browser
  - z/OSMF with advanced capabilities
- Available with z/OS V1R13 RMF and z/OS V1R12 RMF (APAR OA36030)

RMF XP -

HOT TOPICS

z/OS Hot Topics

Learn about key z/OS topics and features

SHARE

Technology · Connections · Results

Now performing!

RMF XP

BY KARIN GENTHER, HARALD BENDER, AND JÜRGEN BAUMANN

Do you want to keep track of one or more IBM zEnterprise BladeCenter Extension (zBX) and performance? Then Resource Measurement Facility Cross Platform (RMF XP) is your choice for cross platform monitoring!

RMF XP provides an integrated performance monitoring solution for heterogeneous environments by currently supporting the operating systems:

- AIX®
- Linux on System x®
- Linux on System z.

Hence, with RMF XP, you can monitor all operating systems which can run on

Performance data at a glance!

The core component of RMF XP is the GPM4CIM server. Similar to the existing Distributed Data Server (DDS) for z/OS, the GPM4CIM server receives HTTP requests and sends back responses as structured XML documents. Because the GPM4CIM started task runs in the z/OS UNIX System Services environment, at least one z/OS system is necessary to run the RMF XP component.

No rehearsal

To start the GPM4CIM server from the console, RMF provides the procedure GPM4CIM as a member in SYS1.PROCLIB, as the JCL example in the GPM4CIM PROC shows:

- The log and trace output is written to the files specified with the STDOUT and STDERR DD cards.
- The 'cfg=' program parameter in the PARM statement points to the GPM4CIM configuration file.
- Different platforms are distinguished by the variable added to the OS statement:
- OS=A (AIX on System p®)
- OS=X (Linux on System x)
- OS=Z (Linux on System z).

Now available on the web! http://publibfp.dhe.ibm.com/epubs/pdf/eoz2n1d0.pdf

RMF XP – Want to see more ?

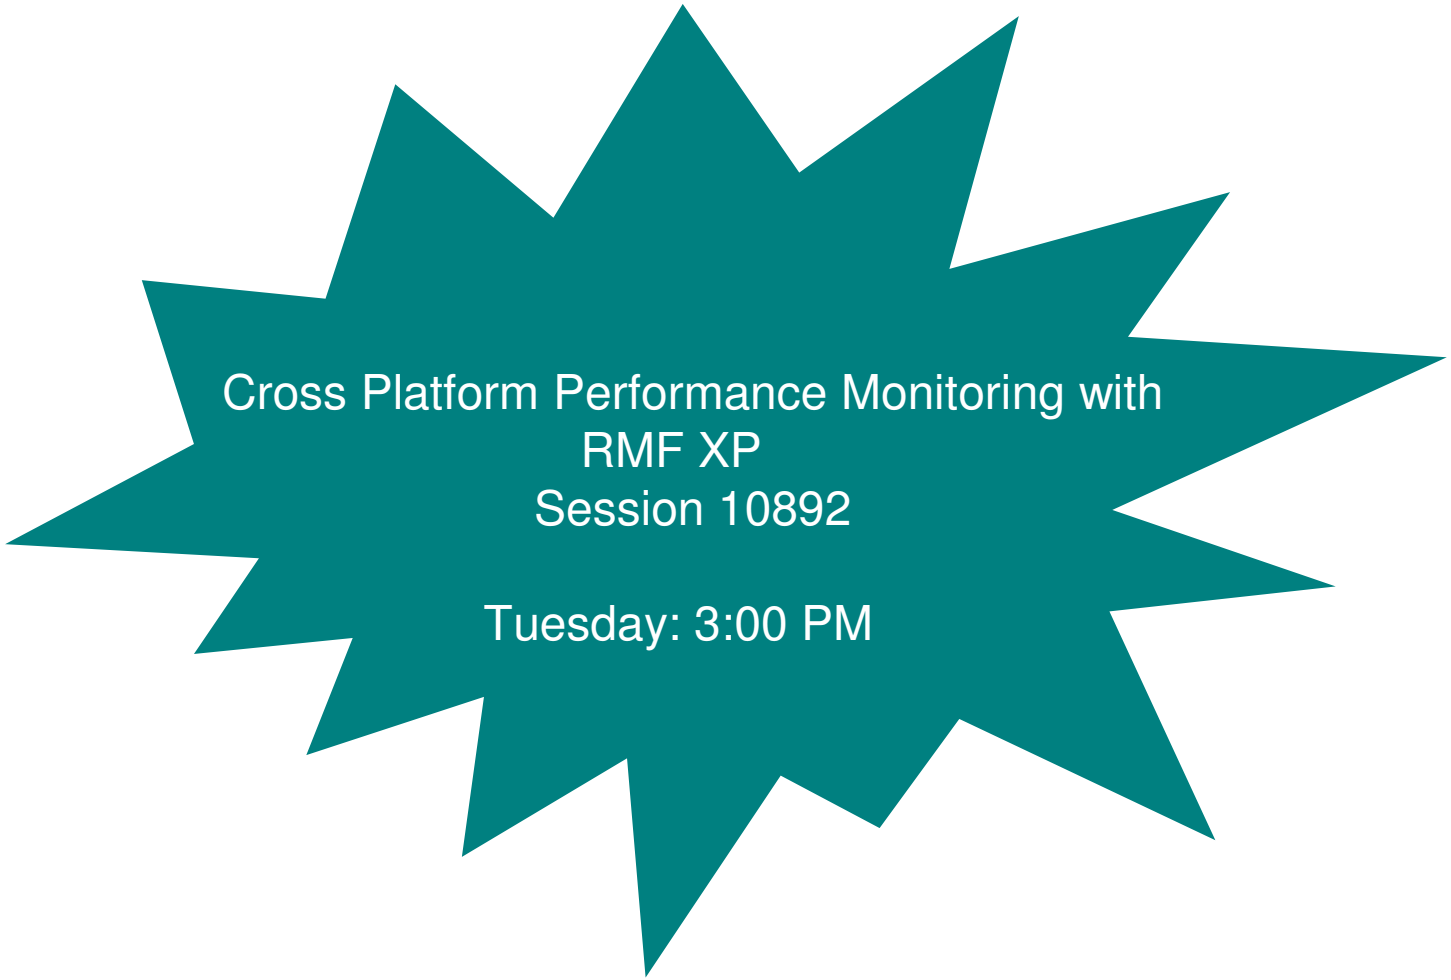

Information and Tools

RMF homepage: www.ibm.com/systems/z/os/zos/features/rmf/

- Product information, newsletters, presentations, ...
- Downloads
  - Spreadsheet Reporter
  - RMF PM Java Edition
  - RMF data collector for Linux

RMF email address: rmf@de.ibm.com

Documentation and news:

- RMF Performance Management Guide, SC33-7992
- RMF Report Analysis, SC33-7991
- RMF User's Guide, SC33-7990
- Latest version of PDF files can be downloaded from:

www.ibm.com/systems/z/os/zos/bkserv/r13pdf/#rmf

RMF Redbook !!!

IBM
SG24-6645-00
Effective zSeries Performance Monitoring using Resource Measurement Facility (RMF)
Review of the traditional facilities
Learn about all the new features and how their setup
How to use RMF for performance monitoring
Pierre Cassier
Raimo Korhonen
Peter Mailand
Michael Teuffel
ibm.com/redbooks
Redbooks

Appendix: Function Reference

Function	Availability
RMF WLM Reporting Enhancements	z/OS 1.13 RMF
RMF Storage Reporting Enhancements	z/OS 1.13 RMF
Postprocessor XML Formatted Reports	z/OS 1.11 RMF
HTTP API to access Historical data	z/OS 1.12 RMF
Reporting of system serialization delays	z/OS 1.13 RMF
In-Ready Work Unit Queue Distribution	z/OS 1.12 RMF
IBM z/OSMF Management Facility: Resource Monitoring Plugin	z/OS 1.12 zOSMF
RMF XP	z/OS 1.13 RMF z/OS 1.12 RMF (APAR OA36030)