# RMF – The Latest and Greatest

Peter Muench (pmuench@de.ibm.com )
IBM Corporation

Monday, August 8, 2011
Session 10011

# 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

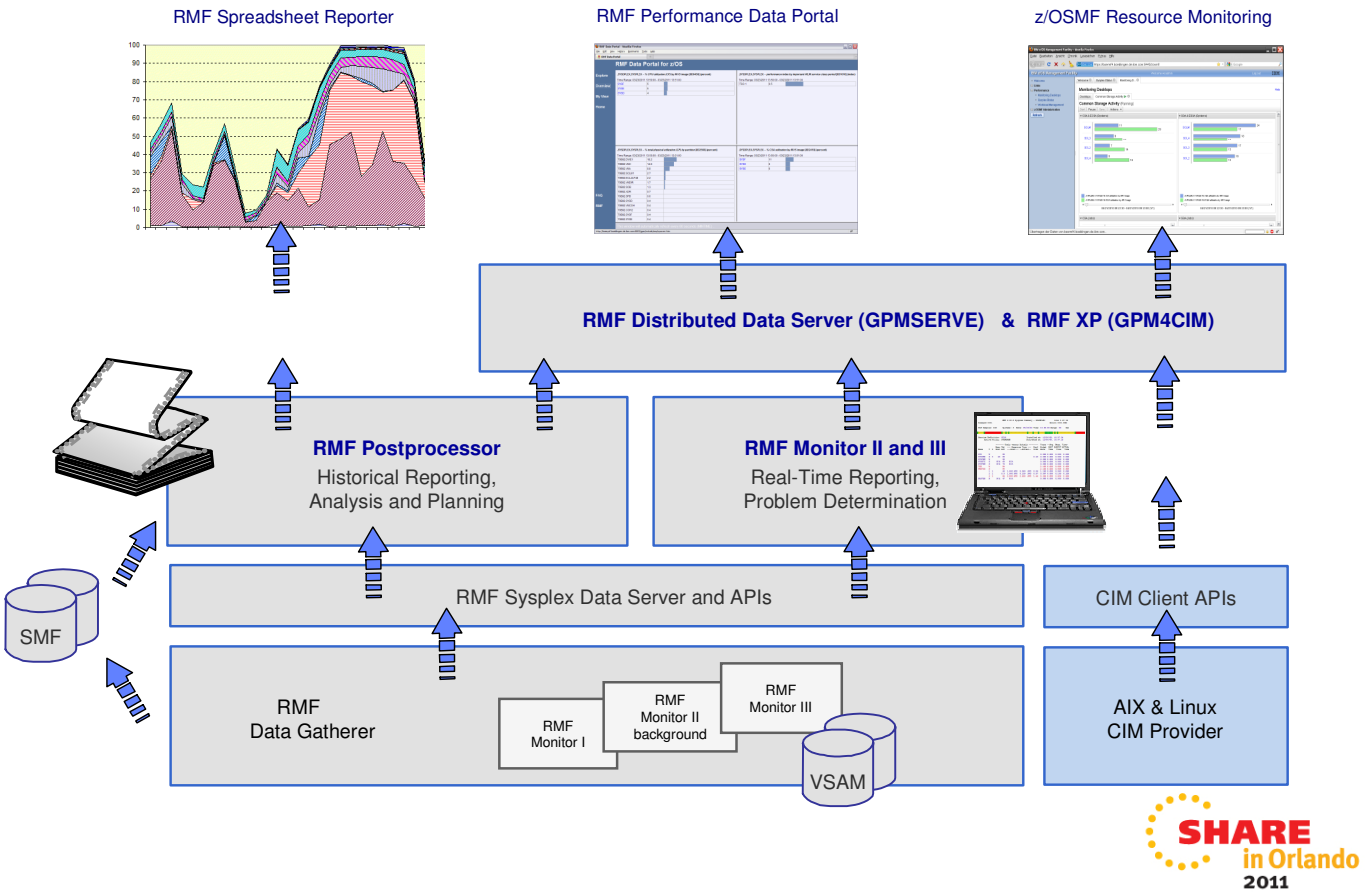


- z/OS Resource Measurement Facility (RMF) is an optional priced feature of z/OS. It supports installations in performance analysis, capacity planning, and problem determination. For these disciplines, different kinds of data collectors are needed:
  - Monitor I long term data collector for all types of resources and workloads. The SMF data collected by Monitor I is mostly used for capacity planning and performance analysis
  - Monitor II snap shot data collector for address space states and resource usage. A subset of Monitor II data is also displayed by the IBM SDSF product
  - Monitor III short-term data collector for problem determination, workflow delay monitoring and goal attainment supervision. This data is also used by the RMF PM Java Client and the RMF Monitor III Data Portal
- Data collected by all three gatherers can be saved persistently for later reporting (SMF records or Monitor III VSAM datasets)
- While Monitor II and Monitor III are realtime reporters, the RMF Postprocessor is the historical reporting function for Monitor I data
- One of the key components for the sysplex wide access of Monitor III data is the RMF Distributed Data Server (DDS). With RMF for z/OS 1.12, DDS supports HTTP requests to retrieve RMF Postprocessor data from a selection of RMF Postprocessor reports. Since the requested data are returned as XML document, a web browser can act as Data Portal to RMF Postprocessor data.
- With z/OS 1.12 there's another exploiter of the RMF DDS data: The z/OSMF Resource Monitoring plugin of the z/OSMF Management Facility.
- RMF for z/OS 1.13 enhances the DDS layer with a new component:
  - RMF XP is the new solution for Cross Platform Performance Monitoring
  - Provides a seamless performance monitoring for all operating systems running on the IBM zEnterprise Bladecenter Extension.

# 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

SHARE in Orlando 2011

In accordance with the availability of new z/OS releases and new hardware functionality, the capabilities of RMF are enhanced consecutively

- Starting with z/OS V1.13, WLM reporting is enhanced to provide response time distributions also for workloads with an execution velocity goal
  - RMF for z/OS 1.13 exploits the new response time distribution data
- Layout and contents of RMF Postprocessor Paging activity report are changed to provide an appropriate reporting of the 64-Bit storage world on today‘s systems
- As an alternative to the standard text format, Postprocessor reports can now be generated in XML format. The new format provides:
  - Browser based, state-of-the-art display of SMF Type 7x data
  - Standardized format for access to performance data via XML parsing
- The RMF Distributed Data Server (DDS) is extended by the capability to grant instant and easy access to RMF long-term historical data by using HTTP and XML.
- With the increasing workload on z/OS systems, resource contention can become more and more a factor that impacts the overall performance.
  - RMF for z/OS 1,13 collects and reports System suspend lock, GRS enqueue and latch contention data
- The in-ready queue distribution is now reported on work unit granularity
- z/OS Resource Monitoring is a new plugin of the z/OSMF Management Facility. The Web browser based user interface allows a cross-sysplex performance monitoring from a single point of control:
  - Enterprise-wide health check of all z/OS sysplexes.
  - Performance monitoring with graphically display of RMF Monitor III metrics as well as Linux metrics by means of customizable views.
- RMF XP is the new solution for Cross Platform Performance Monitoring
  - Provides a seamless performance monitoring for all operating systems running on the IBM zEnterprise Bladecenter Extension.

# 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

- Today, WLM reports response time distributions for workloads with a response time goal in the IWMWRCAA answer area of IWMRCOLL interface
- Starting with z/OS V1.13, WLM reporting is enhanced to provide response time distributions also for workloads with an execution velocity goal
- RMF for z/OS 1.13 exploits the new response time distribution data for transactions with an execution velocity goal:
  - RMF Monitor I collects the new interface data in the SMF 72.3 records. The Service/Report Class Period Data section of the SMF 72.3 record is extended with new fields.
  - RMF Postprocessor (PP) Workload Activity report (WLMGL) now reports response time distributions for response time and execution velocity goals:
    - The response time distribution for service/report class periods with a response time goal is a **sysplex-wide view** with the combined SMF72.3 data from all systems in the sysplex
    - For service/report class periods with excecution velocity goals, there is one response time distribution per system in the sysplex.
  - The RMF PP **WLMGL** report option is enhanced with the new suboption **RTD|NORTD** that can be used to display/suppress the response time distribution section in the workload activity report:
    - The suboption can be specified together with suboptions SCPER and RCPER
    - It is ignored if specified together with other suboptions than SCPER and RCPER
    - Default: RTD
    - Example: SYSRPTS(WLMGL(SCPER(VEL90I1,VEL90I3),NORTD))
      -> Creates RMF PP WLMGL report for service classes VEL90I1 and VEL90I3 **without** response time distribution section
  - The new response time distribution for service/report class periods with execution velocity goal helps the customer to obtain better analysis data and to migrate their goal definitions to a meaningful response time goal.

# RMF WLM Reporting Enhancements …

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

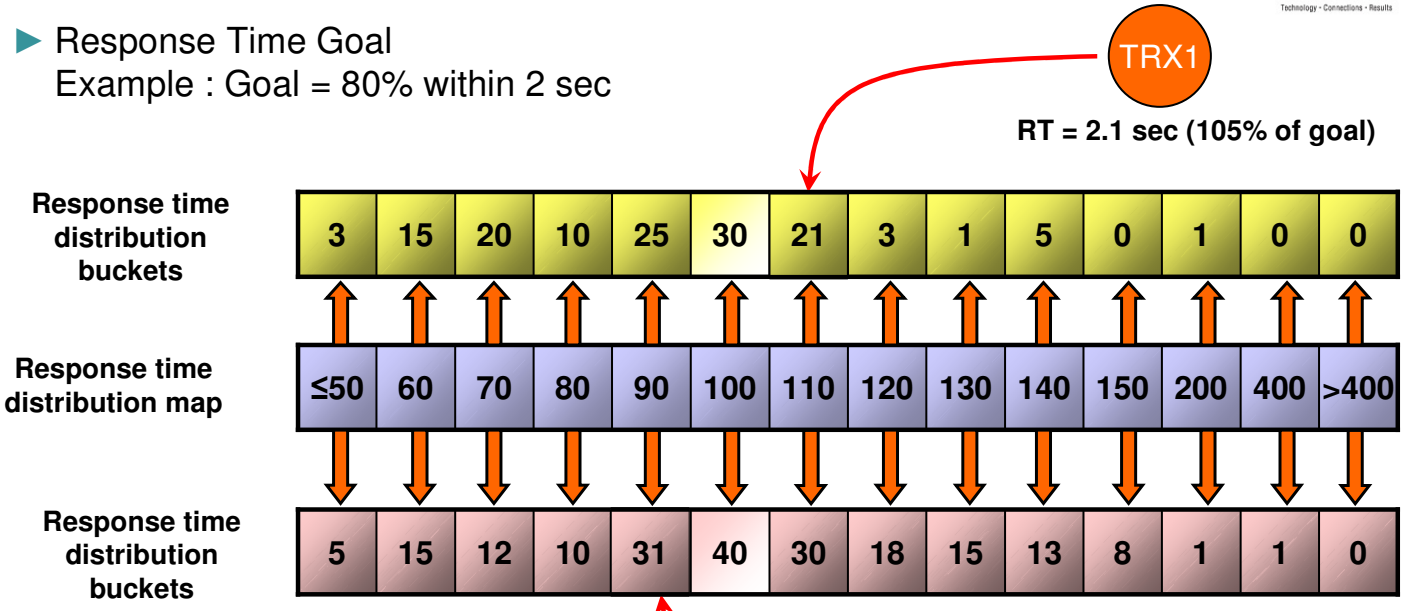


- Execution Velocity Goal
  Example : Velocity = 90%
  Average Response Time = 3.1 sec

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

- For each service/report class period, WLM maintains 14 response time distribution buckets with the number of transactions that completed within a particular time.
- Workloads with response time goals:
  - Distribution is built around **response time goal** as **defined** in WLM service definition
  - The ***response time goal*** for the service class period is split into 14 response time buckets where:
    - bucket 1 covers the gap from 0 to half the goal
    - buckets 2 to 11 cover the gap between half the goal to 1.5 times the goal evenly divided
    - bucket 12 covers two times the goal
    - bucket 13 covers four times the goal
    - bucket 14 covers the gap from four times the goal to infinity
- Workloads with execution velocity goal
  - Distribution is built around **average response time** as **measured** by WLM
  - The ***average of all response times (midpoint)*** for the service/report class period is split into 14 response time buckets where:
    - bucket 1 covers the gap from 0 to half the midpoint
    - buckets 2 to 11 cover the gap between half to 1.5 times the midpoint evenly divided
    - bucket 12 covers two times the midpoint
    - bucket 13 covers four times the midpoint
    - bucket 14 covers the gap from four times the midpoint to infinity

# RMF WLM Reporting Enhancements …

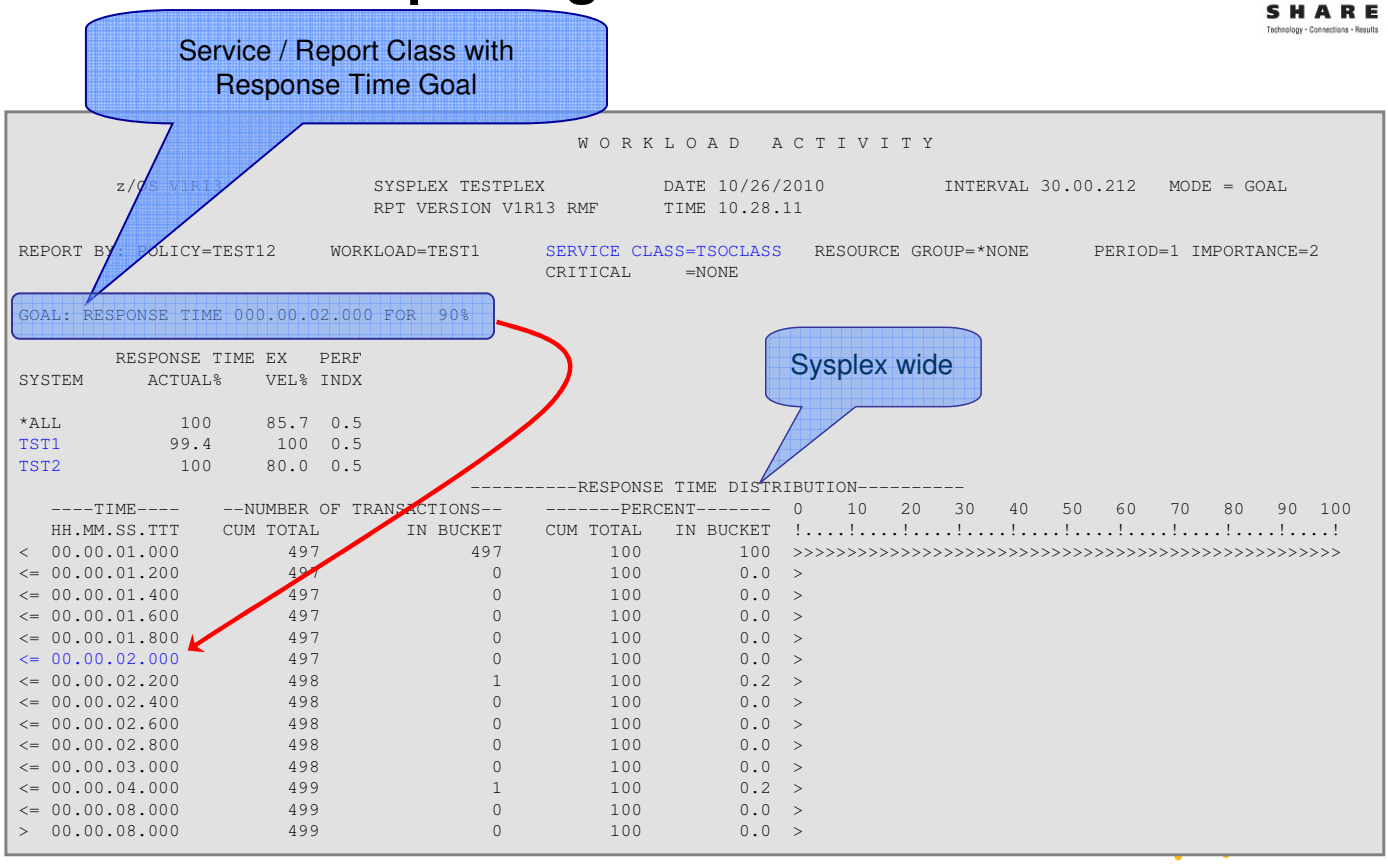


```
                                    W O R K L O A D   A C T I V I T Y

         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  >
```

**Service / Report Class with Response Time Goal:**

- The response time distribution in the RMF PP WLMGL report is a sysplex-wide view created by merging the SMF 72.3 data from all systems in a sysplex.

# RMF WLM Reporting Enhancements …

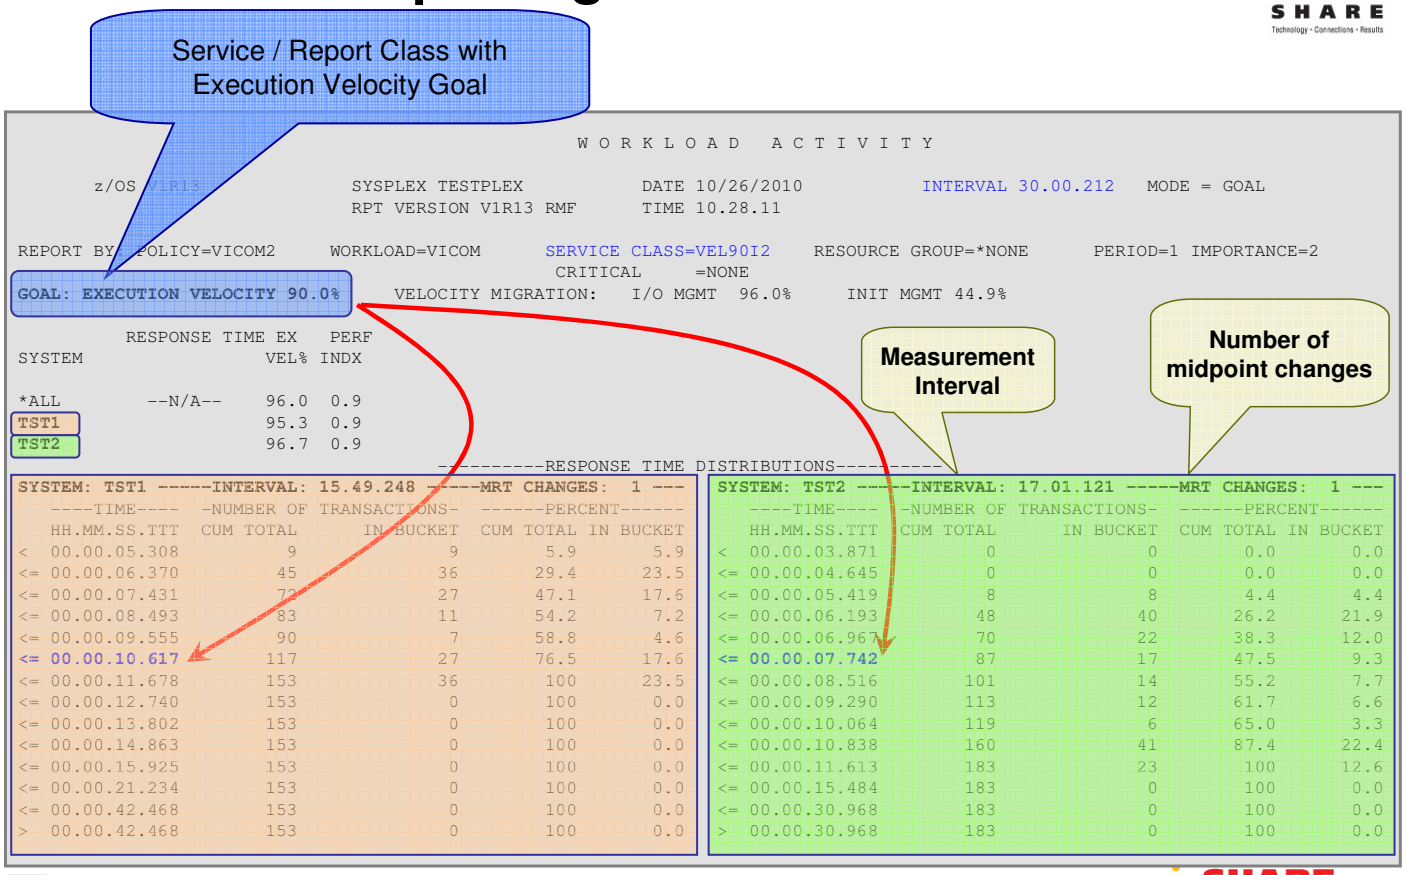


**Service / Report Class with Execution Velocity Goal:**

- The RMF PP WLMGL report displays one response time distribution per system.
- Response time distribution is build around the average response time **measured** by WLM.
- Average repsonse time is calculated from the response time of ended transactions on the system.

# 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

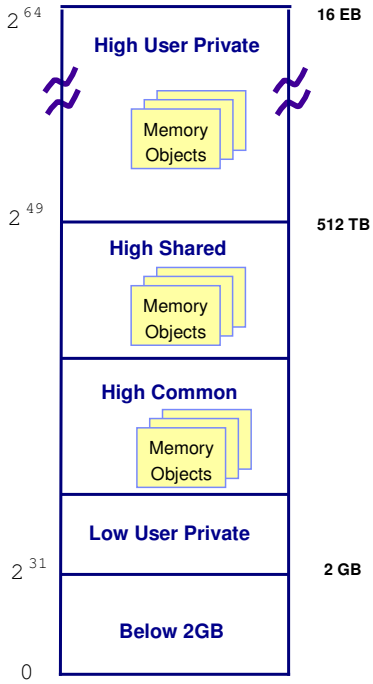


⇨ RMF Postprocessor Paging Activity report restructured and enhanced

- Several reasons to restructure the RMF Postprocessor Paging Activity
  - Section with expanded storage data was removed from the report, but layout not adapted.
  - Beginning with z/OS 1.8, physical swapping of address spaces to AUX is no longer done so that the physical swap metrics are obsolete
  - Current report layout does not allow a proper reporting of an increasing number of metrics for high virtual memory.

# 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 Postprocessor Paging Activity report restructured and enhanced:
  - Report sections restructured:
    - Central Storage Paging Rates
    - Central Storage Movement and Request Rates
    - Frame and Slot Counts
  - New metrics for 64-Bit storage:
    - Size of large frame area (LFAREA)
    - Central storage frames for memory objects in high virtual common and shared storage.
  - The ***Swap Placement Activity*** section is no longer provided.
    - Related overview conditions are not supported when data for z/OS V1R13 or later is processed.
    - As replacement of **logical swap** statistics. RMF provides the new overview condition **LSWAPTOT** to report the accumulated number of all logical swaps.
  - New Memory Objects and Frames section with statistics about high virtual memory objects and frames.
  - Paging Activity report now available in XML format.

# RMF Storage Reporting Enhancements …

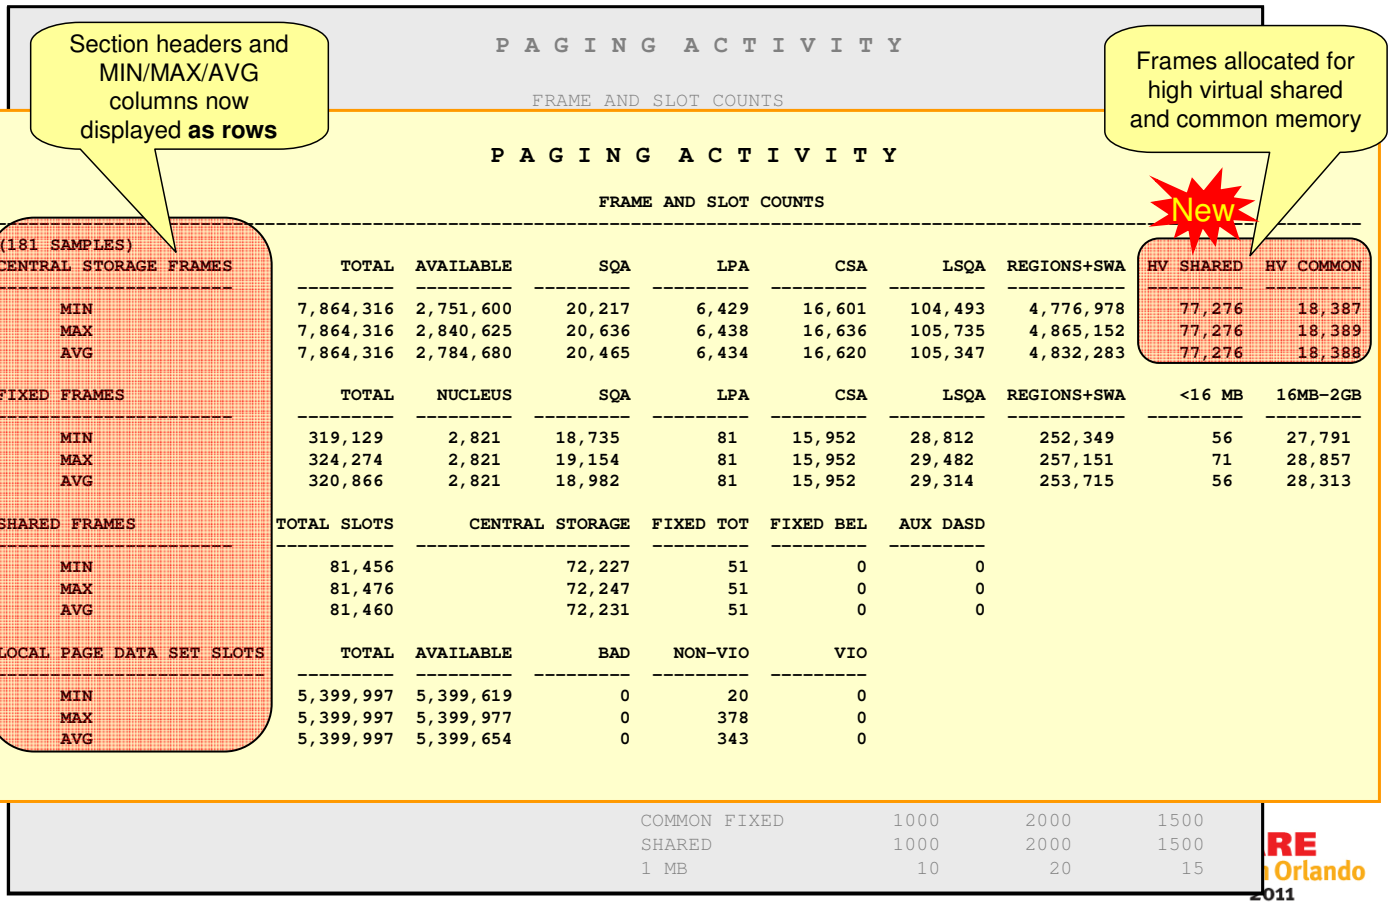

Section headers and MIN/MAX/AVG columns now displayed **as rows**

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 |

- Layout of the Postprocessor **Paging Activity report** is changed to display the section headers and MIN/MAX/AVG columns as **rows**.
- The screenshot shows the new layout of the FRAME AND SLOT COUNTS section.
- CENTRAL STORAGE FRAMES statistics are enhanced with new metrics ‚HV SHARED' and ‚HV COMMON':
  - HV SHARED: MIN, MAX and AVG number of frames allocated for high virtual shared memory
  - HV COMMON: MIN, MAX and AVG number of frames allocated for high virtual common memory

# RMF Storage Reporting Enhancements …

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

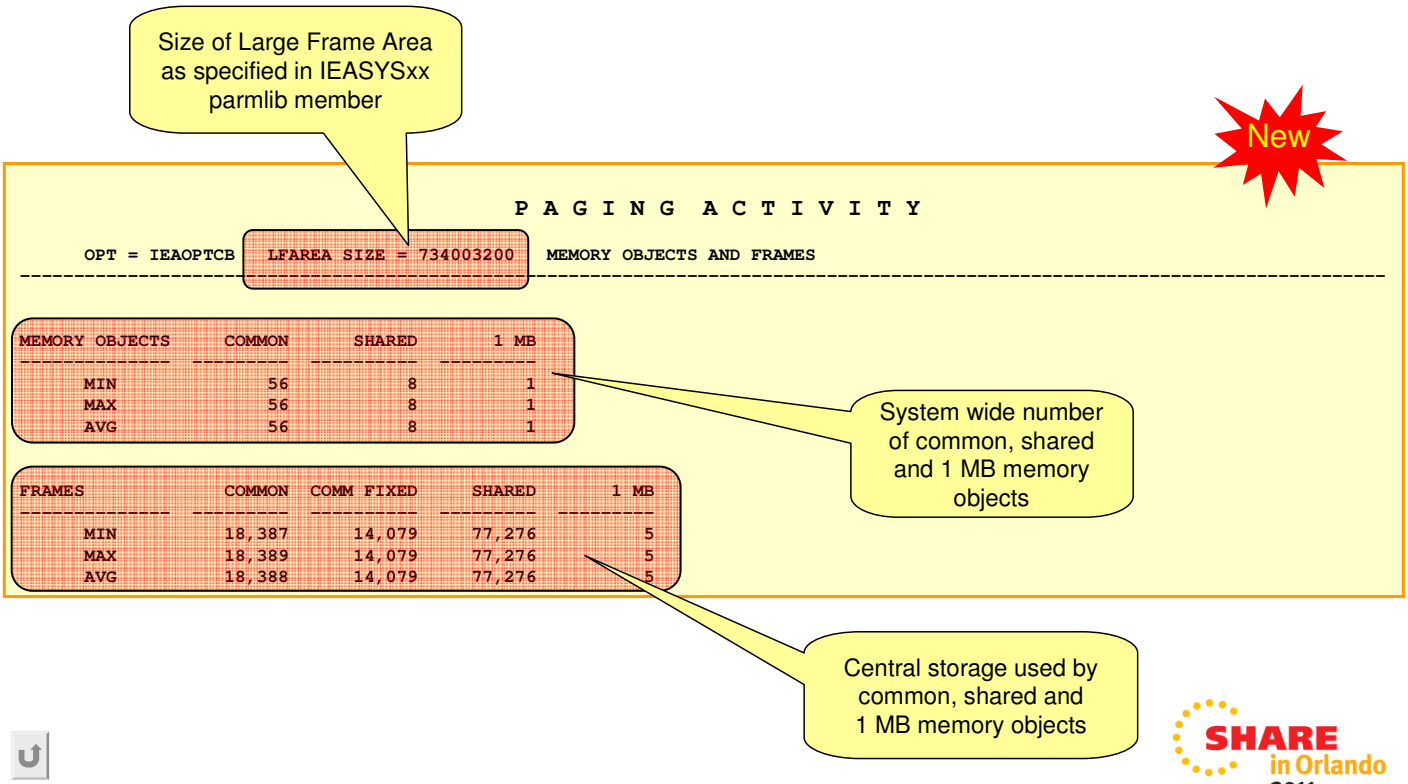


- All usage statistics for Memory Objects and Frames are moved from **FRAME AND SLOT COUNTS** section to new **MEMORY OBJECTS AND FRAMES** section:
  - MEMORY OBJECTS:
    - Number of high virtual COMMON memory objects,
    - Number of high virtual SHARED memory objects and
    - Number of LARGE memory objects (memory object backed by 1 Megabyte Page Frames )
  - FRAMES:
    - Number of high virtual COMMON memory 4K frames that are backed in central storage,
    - Number of high virtual COMMON memory 4K frames that are FIXED in central storage,
    - Number of high virtual SHARED memory 4K frames that are backed in central storage
    - Number of 1 MB frames that are backed in central storage
- The size of the large frame area (as defined by LFAREA parm in IEASYSxx parmlib member) is displayed in the section header.

# 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

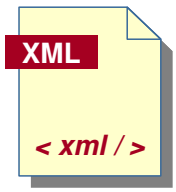


RMF Postprocessor reports can now be generated in XML Format

In the past, users of the RMF Postprocessor had to struggle with a couple of shortcomings:

- The standard text output was limited to to a width of 132 characters. Therefore some reports appeared somehow squeezed. This applies in particular to Overview Reports with a high number of OVW conditions
- The basic text format didn't exploit todays presentation capabilties like resizing windows, scrolling back and forth etc.
- Depending on the report type, the report layouts were quite different. Hence, API programs need to supply specific logic in order to extract metrics out of the reports

Starting with z/OS V1R11 RMF, the Postprocessor is able to generate a subset of the reports alternatively 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

- It is RMF's objective to introduce the XML format for the majority of Postprocessor reports over the next releases.
- With the first stage in RMF for z/OS 1.11, those reports have been selected where no similar Monitor III report is available – since Monitor III reports are already XML enabled by means of the Data Portal.
  These are the following Postprocessor single-system reports:
  - CPU Activity report
  - CRYPTO Activity report
  - FICON Director Activity report
  - OMVS Kernel Activity report,
  - ESS Disk Systems Activity report
- In addition, Postprocessor Overview reports belong also to the first stage since multiple Overview conditions have caused nasty line-feeds in the old text format
- RMF for z/OS 1.12 extends the list of Postprocessor XML reports with another single-system report and the frist sysplex report:
  - Device Activity report (single-system report)
  - Workload Activity report (sysplex report)
- With RMF for z/OS 1.13 the next two Postprocessor single-system reports are XML enabled:
  - Paging Activity report
  - Serialization Delay report (SDELAY)
- The SDELAY report is the first report that is only 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.<br>If you define this ddname, no MFRnnnnn files are created.<br>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.<br>If you define this ddname, no PPORPnnn files are created. |
| XPXSRPTS<br> 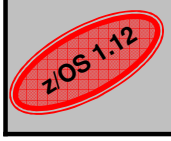 | Combined **sysplex-wide** report in XML format | One ddname for one data set to contain all sysplex reports for each interval included in the input data. | There is no dynamic allocation of this ddname, you have to define it explicitly if you want to get all reports in XML format into one data set or output class.<br>If you define this ddname, no MFRnnnnn files are created.<br>If you define this ddname and PPXSRPTS, no XML output in file XPXSRPTS is created. |

- The request for XML output from the RMF Postprocessor is triggered by means of the existence of a DD card for XPRPTS (for single-system reports) and/or XPOVWRPT (for Overview reports) in the JCL for the Postprocessor job.
- If one of these DD statements are defined in the JCL, then XML output will be generated.
- The following rules apply if requests for text and XML are given simultaneously:
  - If the user specifies a DD card for cumulative text output, i.e. PPRPTS, as well as XPRPTS, then only text output will be generated into PPRPTS and no XML is produced.
  - If single-system reports are requested on the REPORTS control statement which are not XML enabled yet and the DD card XPRPTS is defined, no output is generated for them.
- Beginning with RMF for z/OS 1.12 the new ddname XPXSRPTS was introduced to support XML formated Postprocessor sysplex report.
- For the XPXSRPTS ddname similar rules apply as for XPRTS, that is:
  - If the user specifies the PPXSRPTS ddname for cumulative text output together with XPXSRPTS, no XML output is generated.
  - If sysplex reports other than the WLMGL report are specified in the SYSRPTS control statement, no output is generated.

# Spreadsheet Reporter – XML Support

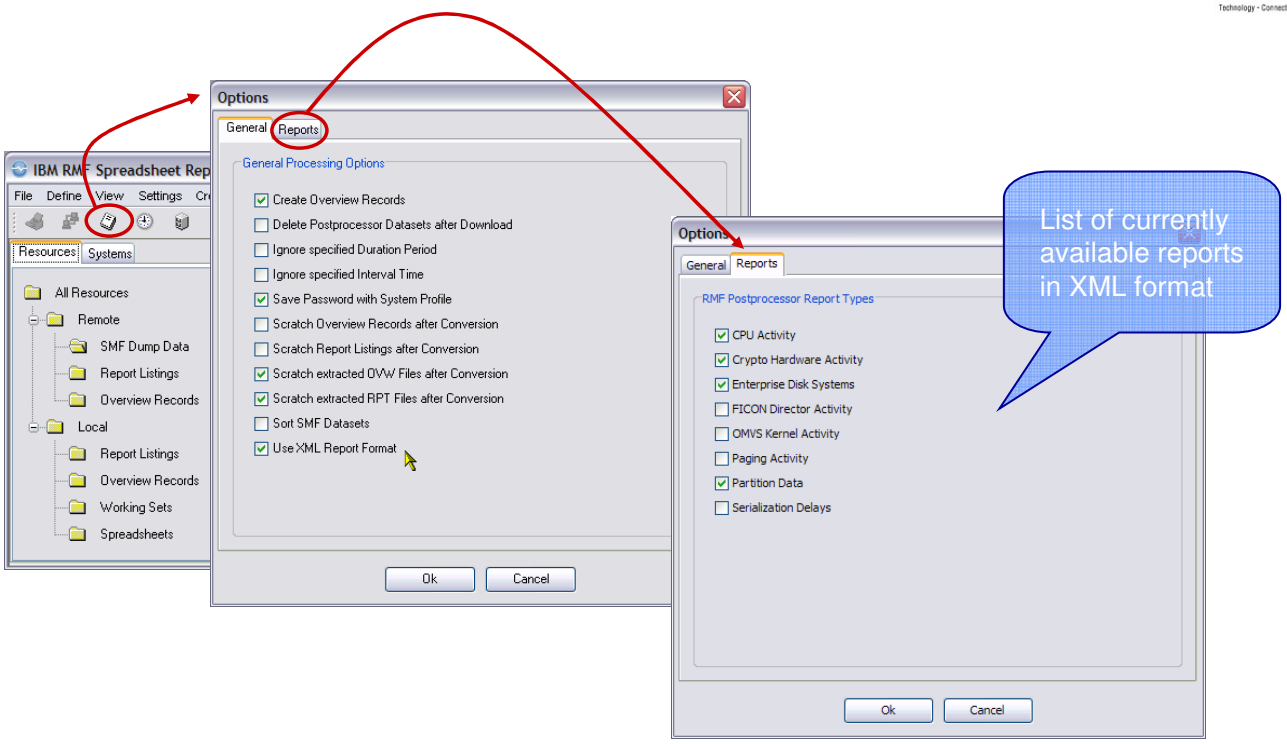


- The RMF Spreadsheet Reporter provides built-in support for the new Postprocessor XML formatted reports
- You can request the new XML format by means of the general option “Use XML Report Format“. This will cause the following results:
  - The report selection list on the Options dialogs Reports tab displays just the report types which can be generated in XML format
  - Under the cover, the generated JCL will contain the new DD names for the XML format (XPRPTS and XPOVWRPT)
  - The default filetype for local report listings changes from *.lis* to *.xml*

# Spreadsheet Reporter – XML Support

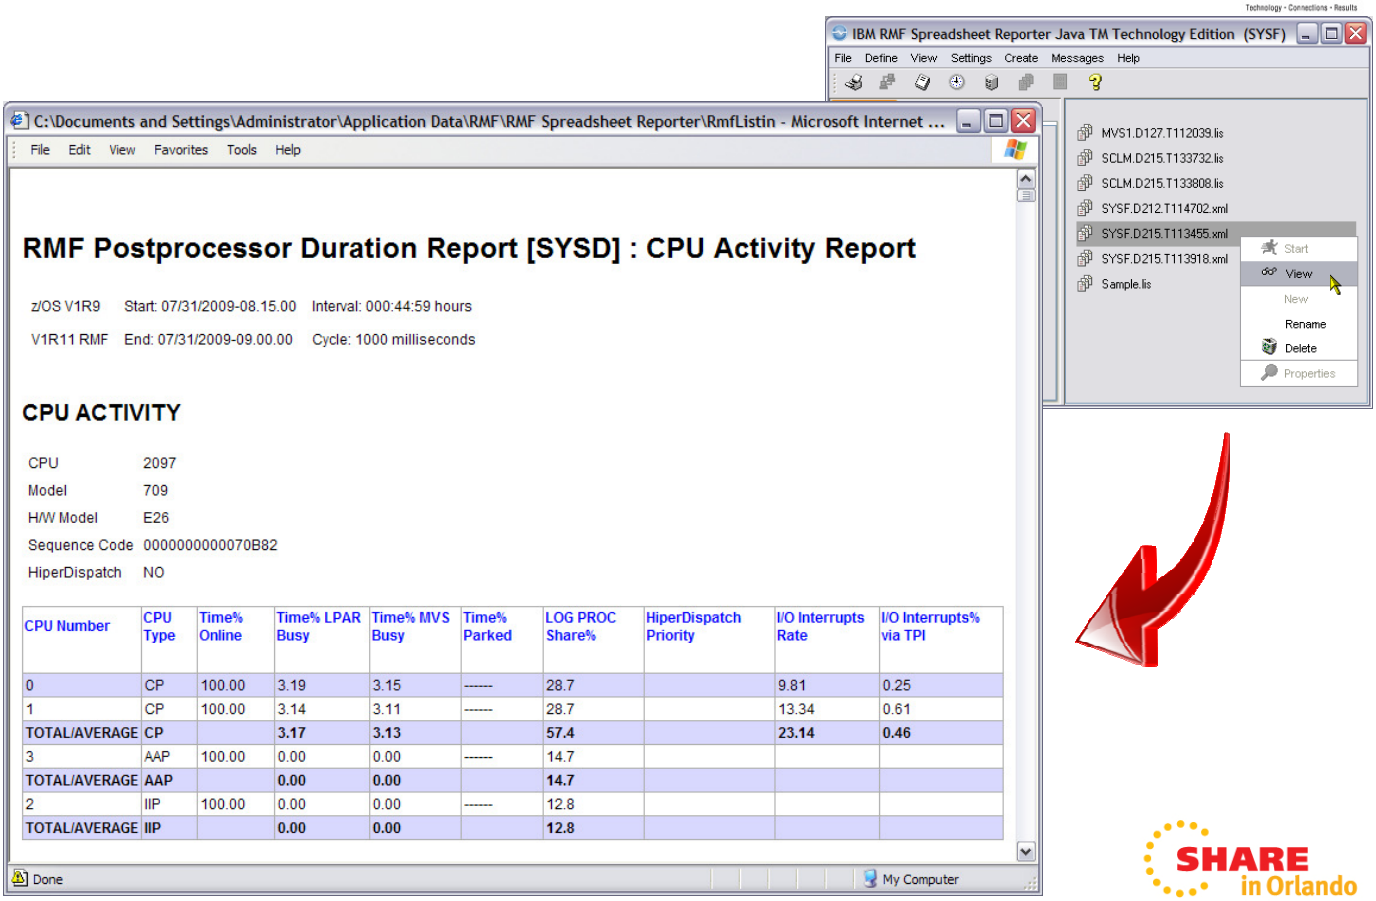


- The new XML formatted reports can be instantly displayed within the Spreadheet Reporter environment
- Once a listing is associated with the filtype *xml* the *View* action opens a browser window with the formatted report
- Report headers are displayed as name/value pairs while the report body is arranged in tabular format
- Internally the XML is visualized by means of the following stylesheets which are located in the Spreadsheet Reporters *Listing* directory:
  - ddsml-pp.xsl
  - ddsml.css

# 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

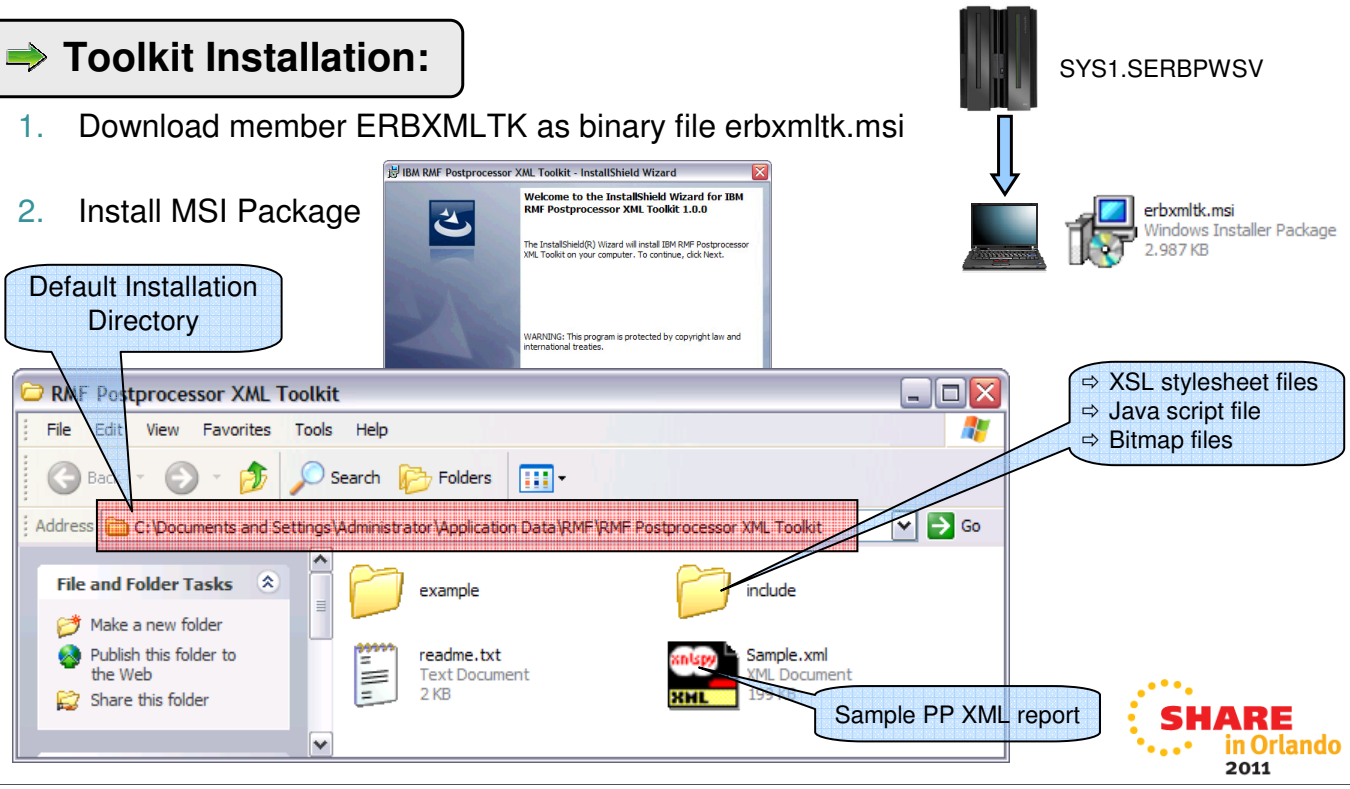


- RMF for z/OS 1.13 includes the new ***Postprocessor XML Toolkit*** which assists you in browsing Postprocessor XML reports with your internet browser. The reports are formatted by use of RMF XSL stylesheets.
- Toolkit installation:
  - The Postprocessor XML Toolkit is part of the RMF product. The application files and installation utility of the Postprocessor XML Toolkit are provided in member ERBXMLTK of the host distribution library SERBPWSV. Download this member as binary file **erbxmltk.msi.**
  - Install the MSI package using the Windows Installer, either by double-clicking on the MSI package file or by issuing the command:
    msiexec /package erbxmltk.msi [/qn]
    The Windows Installer guides you through the installation.
  - Specify the directory where to install the Postprocessor XML Toolkit. The default, for example for Windows XP, is: C:\Documents and Settings_userid_\Application Data\RMF\RMF Postprocessor XML Toolkit
  - The Postprocessor XML Toolkit is installed into program group *IBM RMF Performance Management*.
  - The installation process extracts all files necessary to format and display the XML reports (e.g. XSL stylesheet, JAVA script and Bitmap files) to the toolkit directory.
  - Furthermore, in the toolkit directory, you find hints and tips on how to exploit the produced reports in XML output format.

# XML Support – Postprocessor XML Toolkit …

## ➡ Toolkit Usage:

1. Download Postprocessor XML report into Toolkit directory

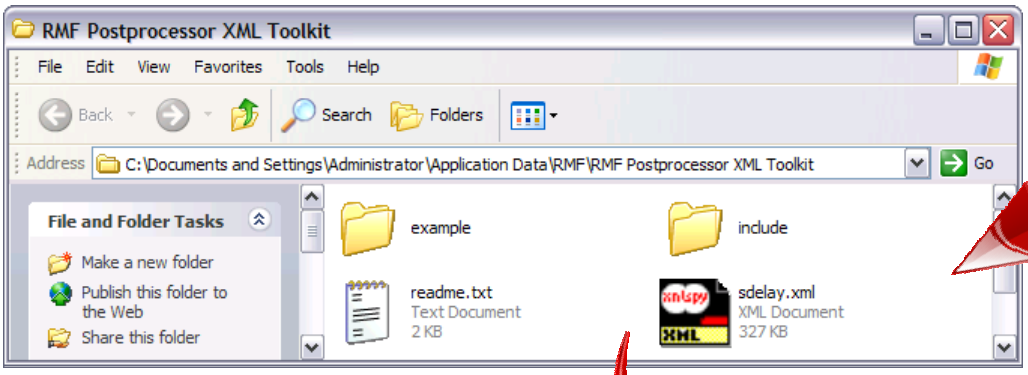


BPMU.SDELAY.LISTING

Download in ASCII format

2. Open XML report with web browser

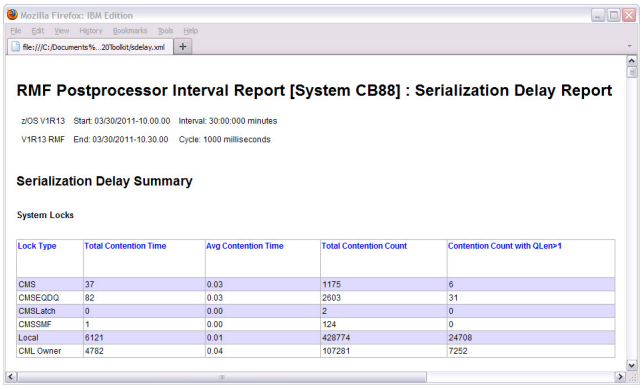


- Toolkit usage:
  - To view an XML Postprocessor report:
    1. Run RMF Postprocessor JOB and direct the XML output (DDNAMES: XPRPTS, XPOVWRPT or XPXSRPTS) to a data set.
    2. Download the XML output data set into the Postprocessor XML Toolkit directory on your workstation with file extension .xml.
       Ensure to download the data set containing the XML output of the Postprocessor reports in ASCII format to the Postprocessor XML Toolkit directory.
    3. Open the XML Postprocessor reports within the Postprocessor XML Toolkit with your internet browser
  - XML formatting details:
    - The stylesheet files that are required to format the report are available in a subdirectory of the Postprocessor XML Toolkit
    - The created Postprocessor reports in XML output format contain a link to the stylesheet in this subdirectory
    - When you open the XML Postprocessor reports within the Postprocessor XML Toolkit with your internet browser, the RMF stylesheet transforms the report into an HTML document.

- Check the example subdirectory of the Postprocessor XML Toolkit. There you find more information on how to exploit RMF Postprocessor XML reports. You also find an example for an RMF Postprocessor realtime reporting setup with sysplex wide scope.

# 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 |

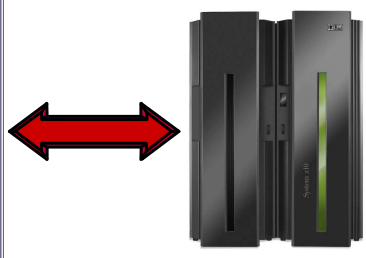

- In most installations, the access to historical data is needed for in depth performance analysis. This allows to keep track whether a critical situation has been persistent or not.
- The existing HTTP API of the RMF Distributed Data Server (DDS) already provides a sysplex-wide access of the data collected by RMF Monitor III. With RMF for z/OS 1.12 this API is extended by the capability to grant instant and easy access to RMF long-term historical data as reported by the RMF Postprocessor.
- Application programs can exploit the extended DDS HTTP API by sending standard URL requests for historical RMF Postprocessor data to the DDS.
- Since DDS returns the requested data as XML document, a web browser can act as an RMF Postprocessor Data Portal to access historical RMF data which is generated by the RMF Postprocessor in XML format.

# HTTP API to access Historical data …

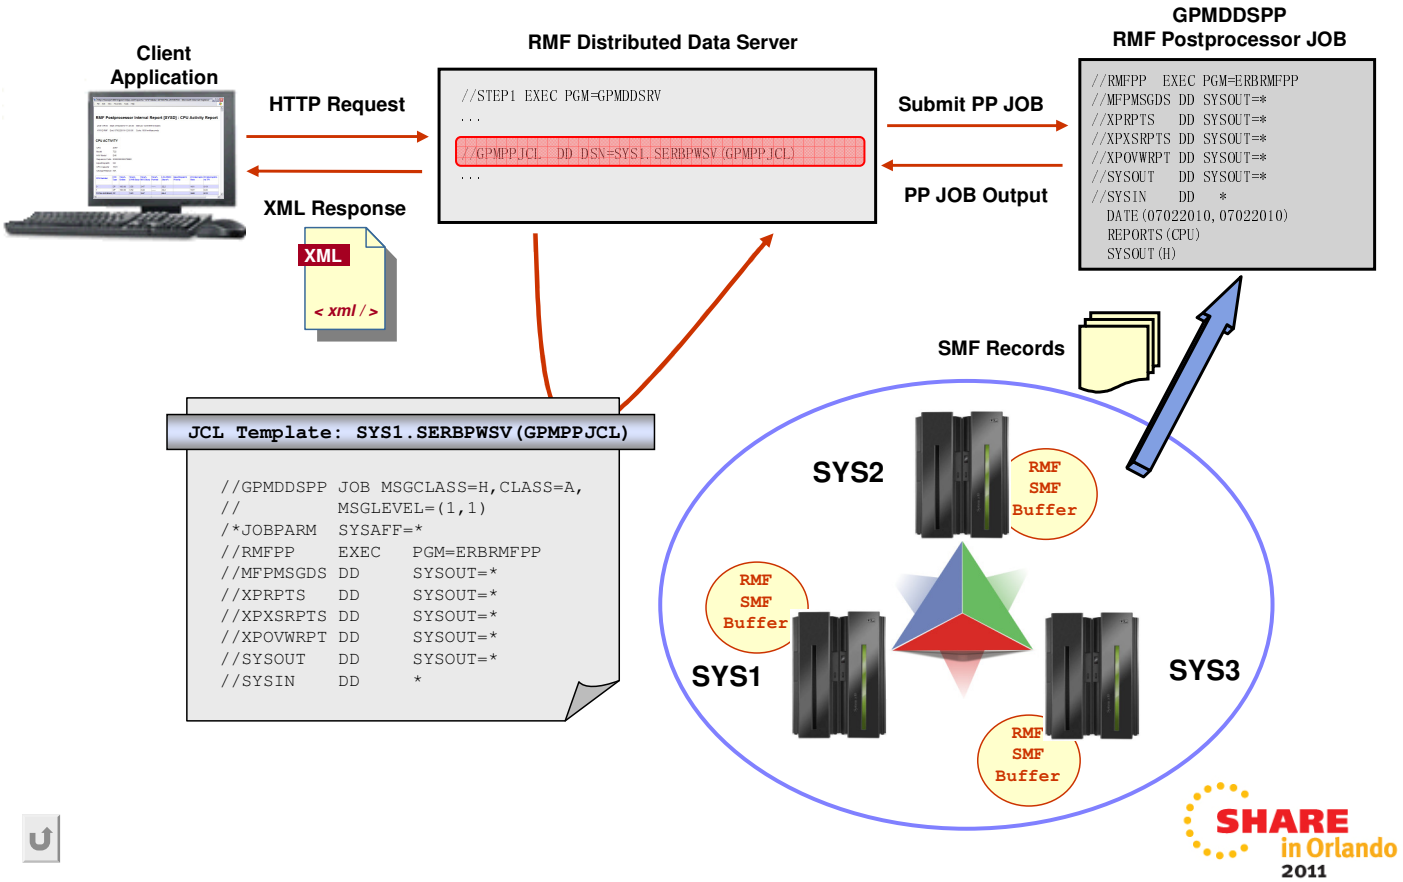


- To get access to RMF Postprocessor data provided by the DDS, the GPMSERVE started task is extended with an additional DD card:
  //GPMPPJCL DD DISP=SHR,DSN=SYS1.SERBPWSV(GPMPPJCL)
- This DD statement points to a JCL template which contains all JCL statements necessary to start a Postprocessor job from the RMF Distributed Data Server. The JCL template is stored in SYS1.SERBPWSV(GPMPPJCL).
- You must adapt or replace the GPMPPJCL member to suit your installation, ensuring that the DDS is able to run RMF Postprocessor jobs. If you do not want to request Postprocessor data with the DDS, you can omit the GPMPPJCL ddname.
- The dataset or the member specified in the GPMPPJCL DD card can be replaced with private ones containing the JCL template for running RMF Postprocessor from the DDS.
- The JCL template that is used by DDS to start the Postprocessor JOB does not contain any Postprocessor control statements. These will be dynamically created by DDS depending on the HTTP request parameters.
- For a detailed description of the DDS HTTP API please refer to the *z/OS RMF Programmer's Guide*.

- Notes:
  - The Postprocessor API functionality is only available with JES2 installed.
  - The DD card specifications for the XML output datasets must not be changed. To be able to retrieve the Postprocessor output, DDS needs these SYSOUT data sets to reside on spool.

# Reporting of system serialization delays

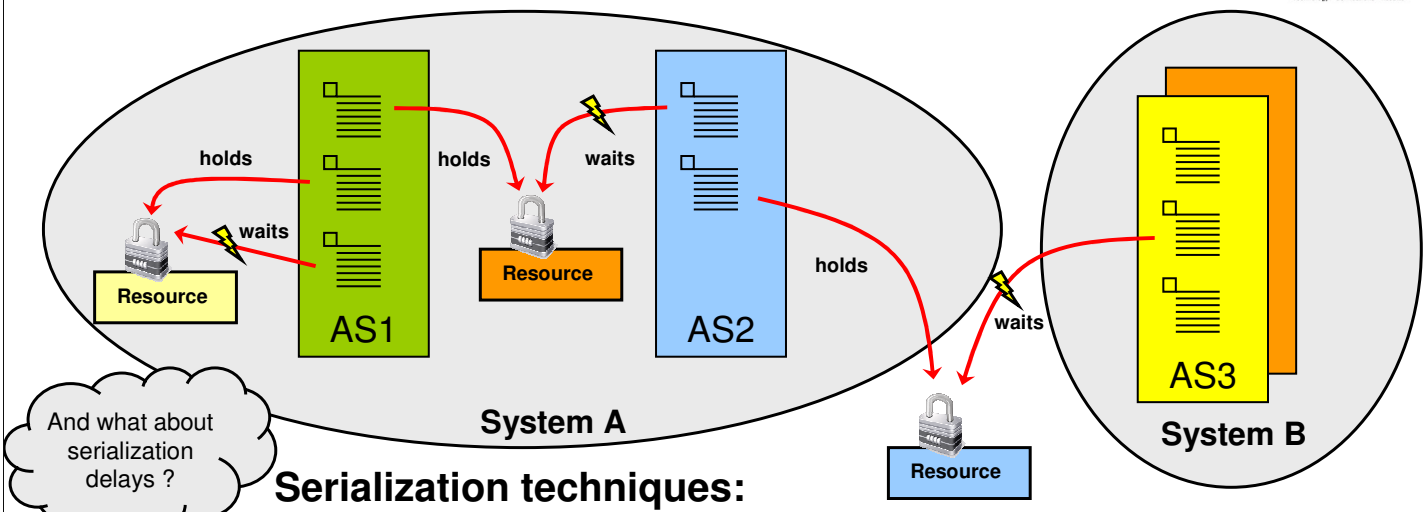


## Serialization techniques:

- 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

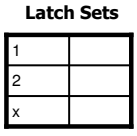


SHARE in Orlando 2011

- There are various techniques to serialize access to resources. These techniques serialize the resources between work units running in the same address space, in different address spaces on the same system and in address spaces on multiple systems. Here is a quick overview of serialization techniques:
- SETLOCK service for system locks
  - Two categories:
    - Global locks -- protect serially reusable resources related to more than one address space (CPU and CMS locks)
    - Local locks -- protect the resources assigned to a particular address space. When the local lock is held for an address space, the owner of the lock has the right to manipulate the queues and control blocks associated with that address space. (CML and LOCAL locks)
  - Lock types:
    - CPU (processor lock) -- serializes on the processor level, providing system-recognized disablement.
    - CMS (general cross memory services lock) -- serializes on more than one address space where this serialization is not provided by one or more of the other global locks.
    - CML (cross memory local lock) -- serializes resources in an address space other than the home address space.
    - Local storage lock (LOCAL) -- serializes functions and storage used by the local supervisor within an address space. There is one lock for each address space.
- GRS provides two sets of system serialization services:
  - GRS ENQ service allows to serialize an abstract resource within the scope of a JOB STEP, SYSTEM or multi-system complex (GRS Complex). The GRS complex is usually equal to the sysplex but it does not have to be.
  - The GRS latch services provide a high speed serialization service for authorized callers. Latch services know nothing of the intended scope. Scoping is completely controlled by the user. It uses user provided storage to manage a lock/latch table that is indexed by a user defined lock/latch number. GRS latch is also widely used. Very big users are USS, Logger, RRS, MVS, etc… Compared to GRS ENQ,the GRS latch services have a significant shorter instruction path so that it's a fast technique to serialize resources within a single address space or across several address spaces in a single MVS system with a minimal overhead.
- With the increased number of systems and workload, resource contention can become more and more a factor that impacts the overalll performance. However, it may be difficult to detect and debug performance problems due to resource contention since system dumps or traditional performance reports may not be adequate tools to identify the address space that is causing a contention.

# 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

- Beginning with z/OS 1.13, RMF collects the following serialization delay data:
  - System suspend lock contention information
  - GRS enqueue and latch performance statistics
- The data is reported in the new RMF PP Serialization Delay report (SDELAY) as a system-wide summary and on address space level.

# 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

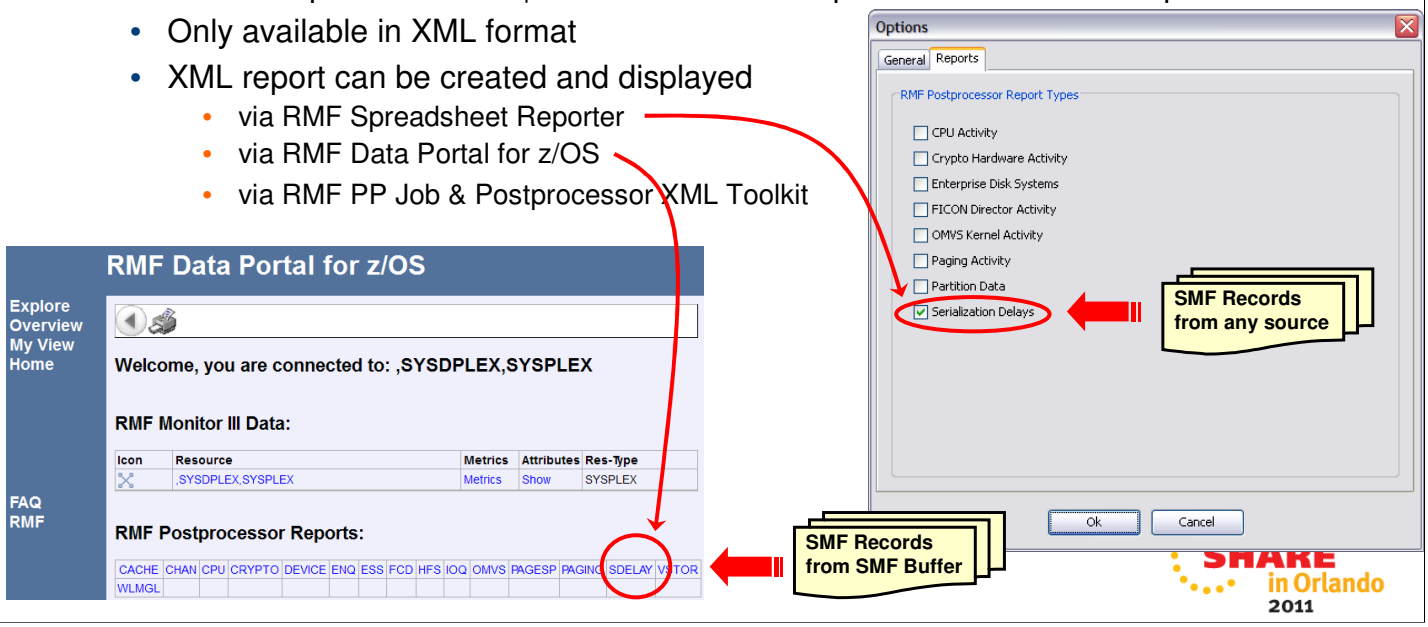


- The serialization delay data is collected by RMF Monitor III and stored in the new SMF 72 subtype 5.
- There is no RMF data gathering option. Data gathering is controlled by the SMF 72-5 setting in the active SMF parmlib member ERBSMFxx. Use TYPE/NOTYPE in the SMFPRMxx parmlib member to enable/disable the data gathering.
- Specify suboption SDELAY in the RMF Postprocessor REPORTS control statement to create the new ***Postprocessor Serialization Delay*** report.
- The Serialization Delay report is the first Postprocessor report that is only available in XML format.
- XML report can be created and displayed
  - via RMF Data Portal (SMF data from SMF buffer only)
  - via Spreadsheet Reporter (any SMF source)
- There is another option to display the SDELAY report in a web browser:
  - Download the XML output created by the RMF Postprocssor JOB to the toolkit directory provided with the RMF Postprocessor XML Toolkit
  - Open the report in the toolkit directory with your web browser.

# Serialization delays: System Summary

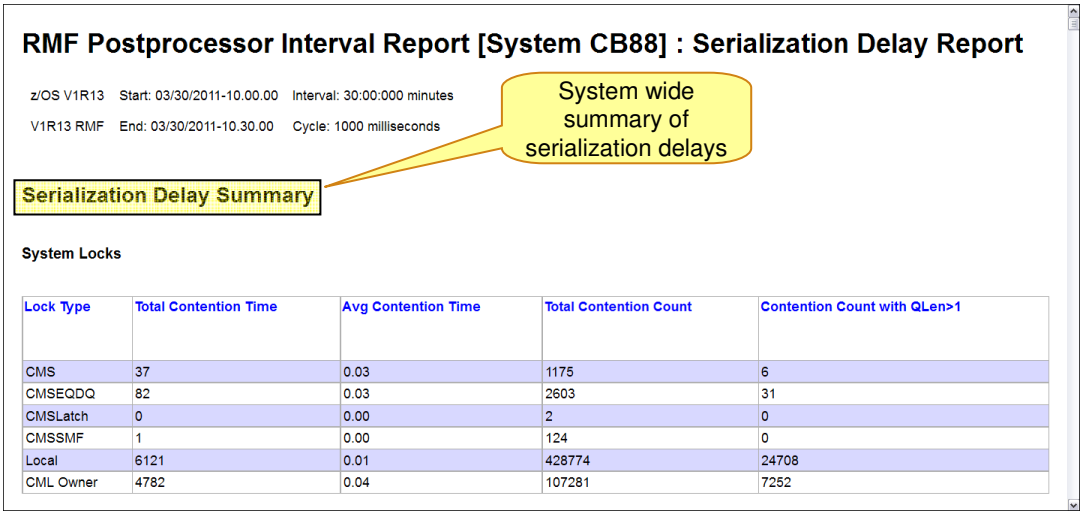

**RMF Postprocessor Interval Report [System CB88] : Serialization Delay Report**

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.

- The ***Serialization Delay report*** consists of two sections:
  - The ***Serialization Delay Summary*** section
  - The ***Serialization Delay Details*** section

- The ***Serialization Delay Summary*** section contains system-wide summary data for all address spaces and is divided into three subsections:
  - The ***System Locks*** subsection displays summary data for system suspend locks.
  - The ***GRS Latch*** subsection displays summary data about GRS latches.
  - The ***GRS Enqueue*** subsection displays summary data about GRS enqueue requests.

- The screenshot shows the ***System Locks*** subsection

- **Lock Type:**

| | |
|---|---|
| **CMS** | CMS lock |
| **CMSEQDQ** | CMS Enqueue/Dequeue lock |
| **CMSLatch** | CMS Latch lock |
| **CMSSMF** | CMS SMF lock |
| **Local** | LOCAL lock |
| **CML Owner** | CML lock owner |

- Total Contention Time — The total amount of time in milliseconds that a unit of work was suspended by a lock of the indicated type.
- Avg Contention Time — The average amount of time in milliseconds that a unit of work was suspended by a lock of the indicated type.
- Total Contention Count — The total number of times that a unit of work was suspended by a lock of the indicated type.
- Contention Count with QLen>1 — The total number of times that a unit of work was suspended by a lock of the indicated type when there was already at least one other unit of work suspended for the lock (that is, queue length > 1).

# Serialization delays: System Summary

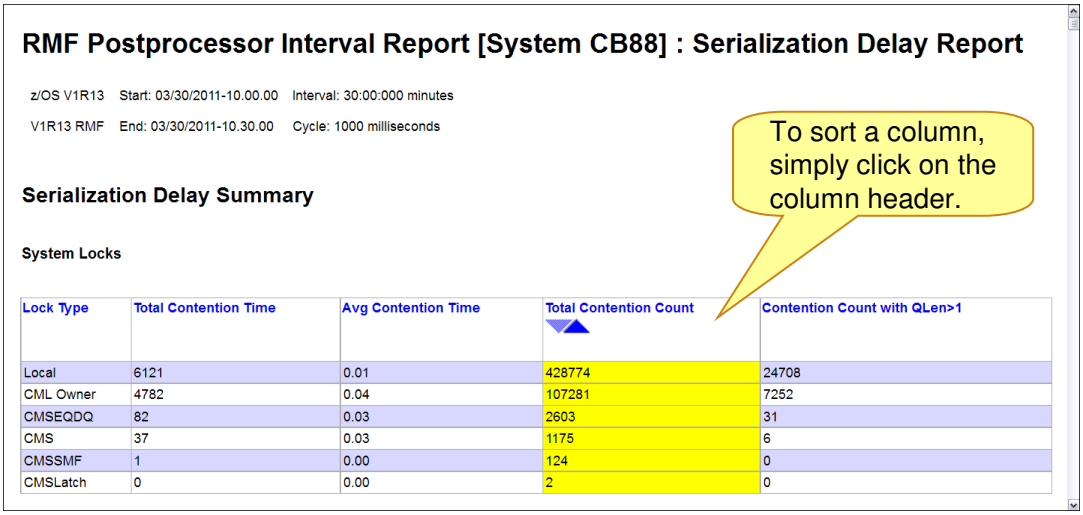

**RMF Postprocessor Interval Report [System CB88] : Serialization Delay Report**

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

- The XSL stylesheet and JAVA script files provided within the RMF Spreadsheet reporter or RMF Postprocessor XML Toolkit allows to format and display the Serialization Delay report with a web browser.
- Each column of the report can be used as sort criteria to sort the report section:
  - 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 …

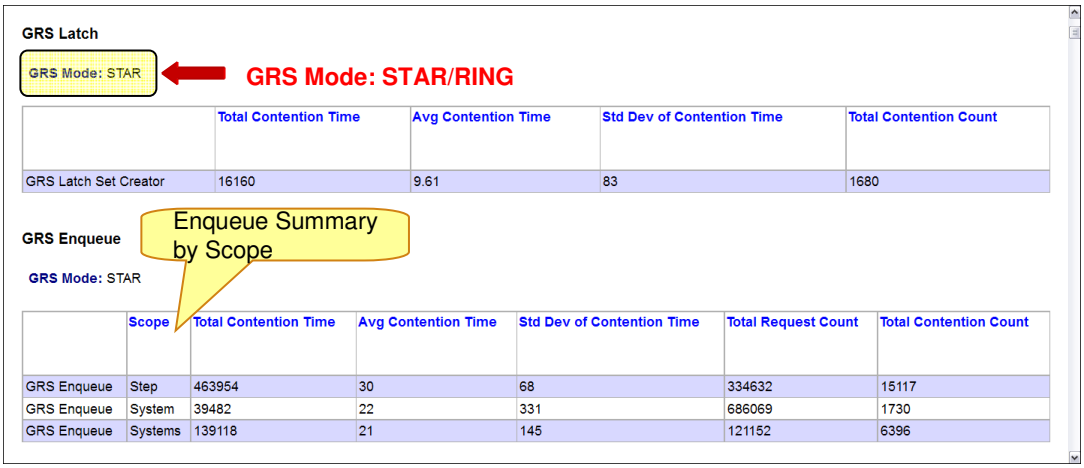

GRS Latch

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

***GRS Latch*** section: Contains summary data about GRS latches for all address spaces:

| | |
|---|---|
| • GRS Mode | The operation mode of GRS: NONE, RING or STAR |
| • Total Contention Time | The total amount of time in milliseconds that latch obtain requests were suspended. |
| • Avg Contention Time | The average amount of time in milliseconds that latch obtain requests were suspended. |
| • Std Dev of Contention Time | The standard deviation of the total contention time in milliseconds. |
| • Total Contention Count | The total number of suspended latch obtain requests. |

***GRS Enqueue*** section: Contains summary data about GRS enqueue requests for all address spaces

| | |
|---|---|
| • GRS Mode | The operation mode of GRS: NONE, RING or STAR |
| • Scope | The scope of an GRS enqueue request: STEP, SYSTEM or SYSTEMS<br>One line is displayed for requests of a certain scope. |
| • Total Contention Time | The total amount of time in milliseconds that the GRS ENQ requests with the specified ***Scope*** were suspended. |
| • Avg Contention Time | The average amount of time in milliseconds that the GRS ENQ requests with the specified ***Scope*** were suspended. |
| • Std Dev of Contention Time | The standard deviation of the ***Total Contention Time*** in milliseconds. |
| • Total Request Count | The total number of GRS ENQ requests with the specified ***Scope***. |
| • Total Contention Count | The total number of GRS ENQ requests with the specified ***Scope*** that were suspended. |

# Serialization delays: CMS Lock Details

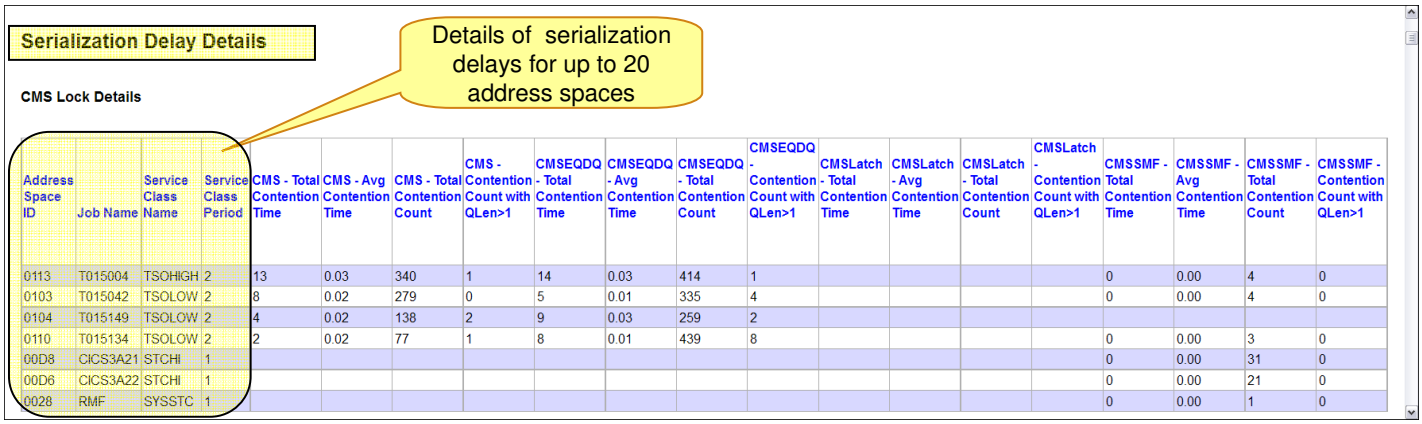

**Serialization Delay Details**

**CMS Lock Details**

Details of serialization delays for up to 20 address spaces

| 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.

- The ***Serialization Delay Details*** section provides the following information in four subsections:
  - The ***CMS Lock Details*** subsection contains detail data about CMS/CMSEQDQ/CMSLatch/CMSSMF locks per address space.
  - The ***CML and LOCAL Lock Details*** subsection contains detail data about CML and LOCAL locks per address space.
  - The ***GRS Latch Details*** subsection contains detail data about GRS latches.
  - The ***GRS Enqueue Details*** subsection contains detail data about GRS enqueue requests.

**Note:** For each lock type, a maximum of the top twenty address spaces with the longest contention times are reported.

***CMS Lock Details*** section: Detail data about CMS/CMSEQDQ/CMSLatch/CMSSMF locks per address space

| | |
|---|---|
| • Address Space ID | The hexadecimal address space identifier (ASID) of the job for which lock data was collected. |
| • Jobname | The name of the job. |
| • Service Class Name | The name of the service class that the job has been running in. |
| • Service Class Period | The service class period that the job has been running in. |
| • Total Contention Time | The total amount of time in milliseconds that a unit of work of the indicated address space was suspended on the respective lock type. |
| • Avg Contention Time | The average amount of time in milliseconds that a unit of work of the indicated address space was suspended on the respective lock type. |
| • Total Contention Count | The number of times that a unit of work of the indicated address space was suspended on the respective lock type. |
| • Contention Count with QLen>1 | The number of times that a unit of work of the indicated address space was suspended on the respective lock type when there was already at least one other unit of work suspended for the 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:

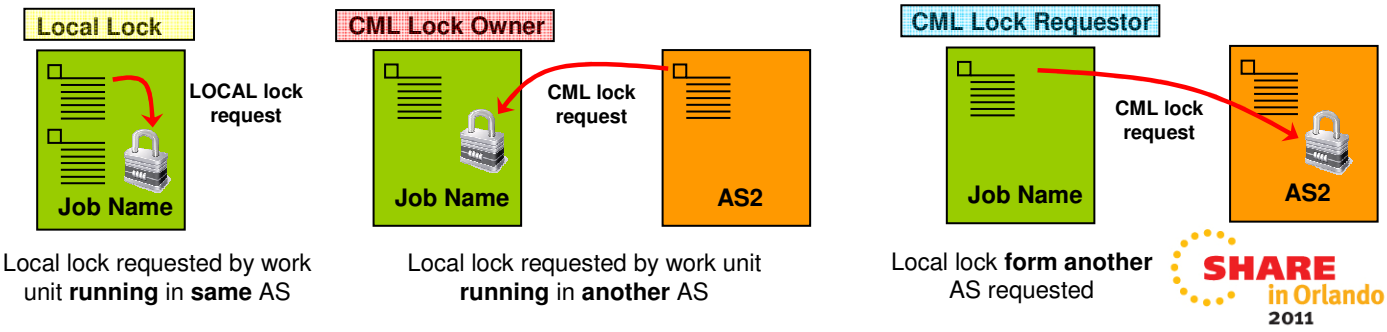


Local lock requested by work unit **running** in **same** AS

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

Local lock **form another** AS requested

SHARE in Orlando 2011

***CML and LOCAL Lock Details*** section: Delays on address space level for local locks

**CML Lock Owner :**

- Total Contention Time — The total amount of time in milliseconds that a unit of work from another address space was suspended when requesting the local lock of the indicated address space.
- Avg Contention Time — The average amount of time in milliseconds that a unit of work from another address space was suspended when requesting the local lock of the indicated address space.
- Total Contention Count — The number of times that a unit of work from another address space was suspended when requesting the local lock of the indicated address space.
- Contention Count with QLen>1 — The number of times that a unit of work from another address space was suspended when requesting the local lock of the indicated address space and there was already at least one other unit of work waiting for this lock.

**Local Lock:**

- Total Contention Time — The total amount of time in milliseconds that a unit of work of the indicated address space was suspended on a local lock.
- Avg Contention Time — The average amount of time in milliseconds that a unit of work of the indicated address space was suspended on a local lock.
- Total Contention Count — The number of times that a unit of work of the indicated address space was suspended on a local lock.
- Contention Count with QLen>1 — The number of times that a unit of work of the indicated address space was suspended on a local lock when there was already at least one other unit of work suspended.

**CML Lock Requestor:**

- Total Contention Time — The total amount of time in milliseconds that a unit of work of the indicated address space was suspended when requesting the local lock of another address space.
- Avg Contention Time — The average amount of time in milliseconds that a unit of work of the indicated address space was suspended when requesting the local lock of another address space.
- Total Contention Count — The number of times that a unit of work from this address space was suspended when requesting the local lock of another address space.
- Contention Count with QLen>1 — The number of times that a unit of work from this address space was suspended when requesting the local lock of another address space and there was already at least one other unit of work waiting for that lock.

# 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


***GRS Latch Details*** section: Detail data about GRS latches on address space level

| | |
|---|---|
| ***Latch Set Creator:*** | Statistics for latch obtain requests ***against latch sets created by this address space*** |
| ***Latch Set Requestor:*** | Statistics for latch obtain requests ***issued from this address space***: |
| • Total Contention Time | The amount of contention time in milliseconds that was caused by latch obtain requests. |
| • Avg Contention Time | The average amount of contention time in milliseconds that was caused by latch obtain requests. |
| • Std Dev of Contention Time | The standard deviation of the total contention time. |
| • Total Contention Count | The number of times a latch obtain request was suspended. |

# 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

***GRS Enqueue Details*** section: Detail data about GRS enqueue requestson on address space level

Grouped by enqueue scopes: STEP, SYSTEM and SYSTEMS

| | |
|---|---|
| • Total Contention Time | The total amount of contention time in milliseconds that was caused by GRS ENQ requests of the indicated scope for this address space. |
| • Avg Contention Time | The average amount of contention time in milliseconds that was caused by GRS ENQ requests of the indicated scope for this address space. |
| • Std Dev of Contention Time | The standard deviation of the total contention time in milliseconds for GRS ENQ requests of the indicated scope for this address space. |
| • Request Count | The total number of GRS ENQ requests of the indicated scope for this address space. |
| • Contention Count | The total number of GRS ENQ requests of the indicated scope that were suspended for this address space. |

# 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

- In the past, RMF has reported the number of address spaces waiting for or using a CPU (see In-ready statistics within the System Address Space Analysis of the Postprocessor CPU Activity report). However, with more than one task per address space, this does not really reflect the demand for CPU. Performance analysts need to know how many work units (WEBs) are running or waiting for a CPU.
- For this reason the CPU Activity report has been extended and new Overview Conditions have been implemented.

# Work Unit Queue Distribution: Monitor I CPU Report

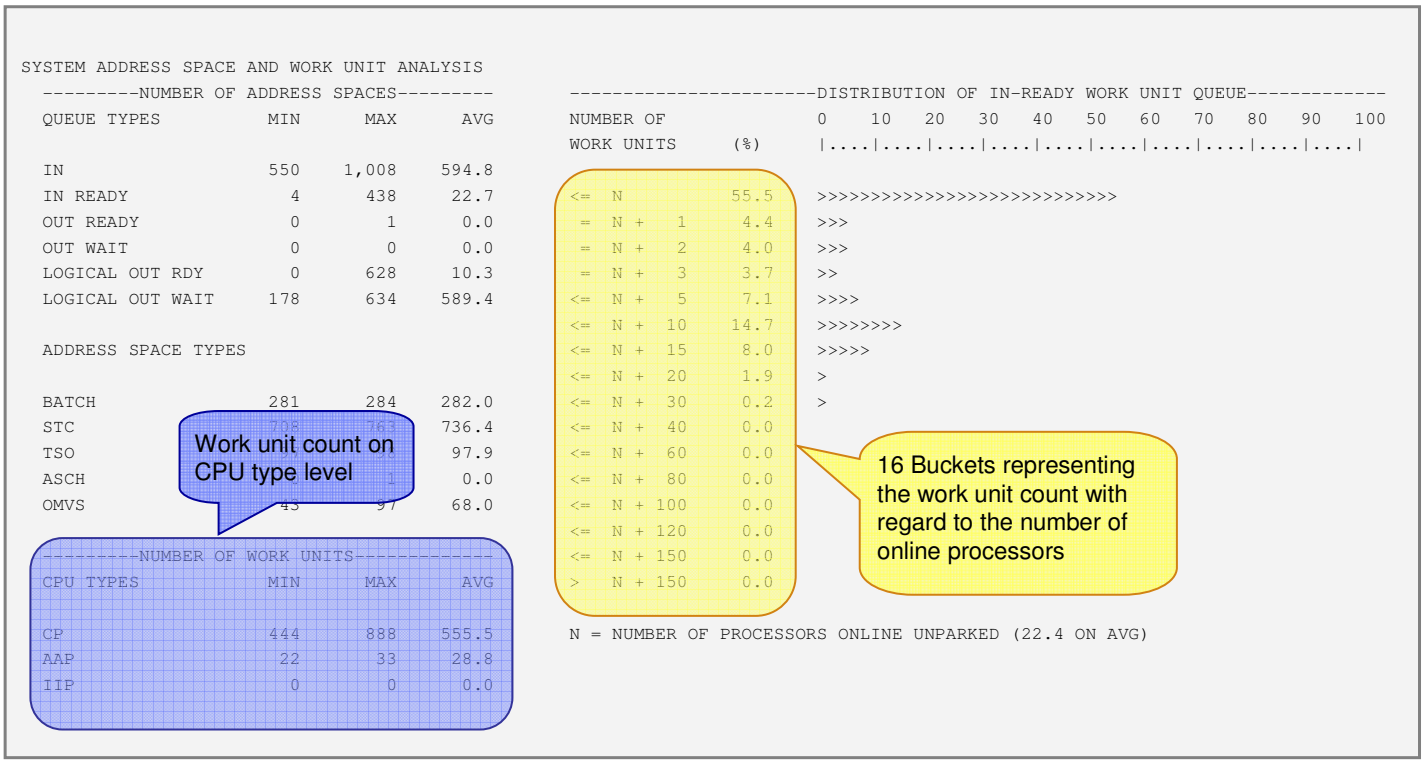


The SYSTEM ADDRESS SPACE section of the Monitor I CPU Activity Report has been changed as follows:

- The In-Ready Queue Distribution is now based on the real number of dispatchable work units
  - We display 16 buckets of ascending queue lengths with regard to the number of actual online processors when the sample was taken
  - This allows to detect latent demand for the processor resource and to distinguish uncritical spikes from severe shortages
  - The number of work units within the buckets includes standard CPs as well as specialty engines
- The new NUMBER OF WORK UNITS statistics allow to determine the CPU demand on processor type level
  - The MIN, MAX and AVG queue lenghts are shown for CPs, AAPs and IIPs

# Work Unit Queue Distribution Example

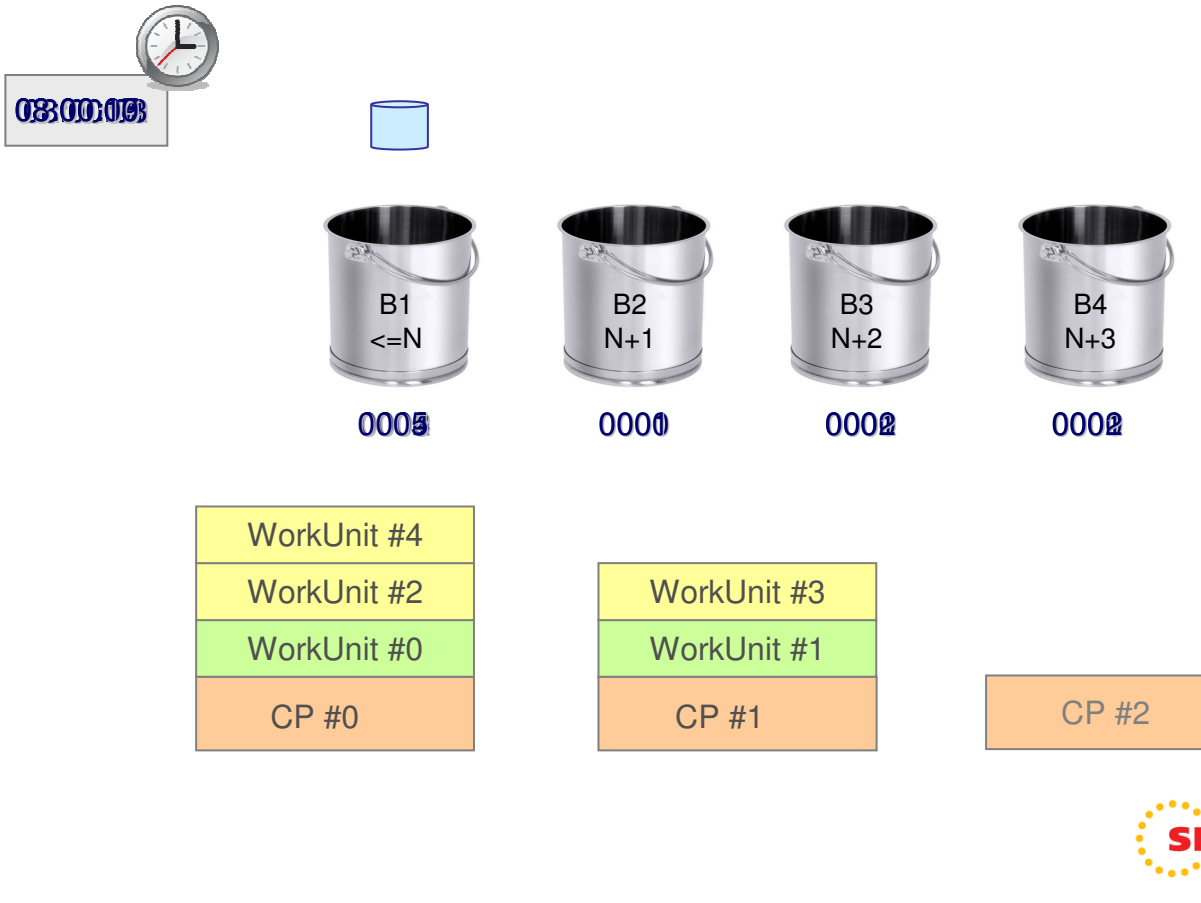


- This example points out the concept of buckets with regard to the number of online processors. For simplification reasons, there are just four buckets (instead of 16) and five work units as maximum
- As long as there are less or equal work units than processors, all these samples contribute to bucket number one
- In case there are more work units than processors the corresponding bucket is increased by one for each sample
- When an additional processor becomes available – proposed that the number of work units remains unchanged – the left-adjacent bucket will be increased

# Work Unit Queue Distribution: Spreadsheet Reporter

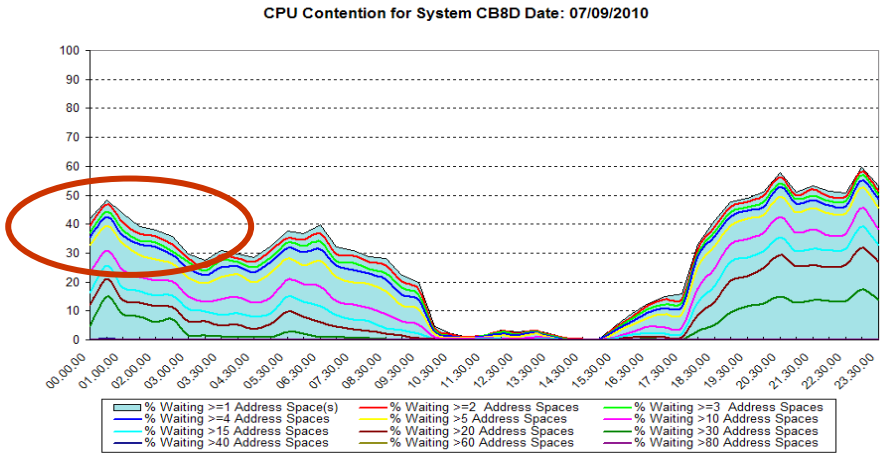


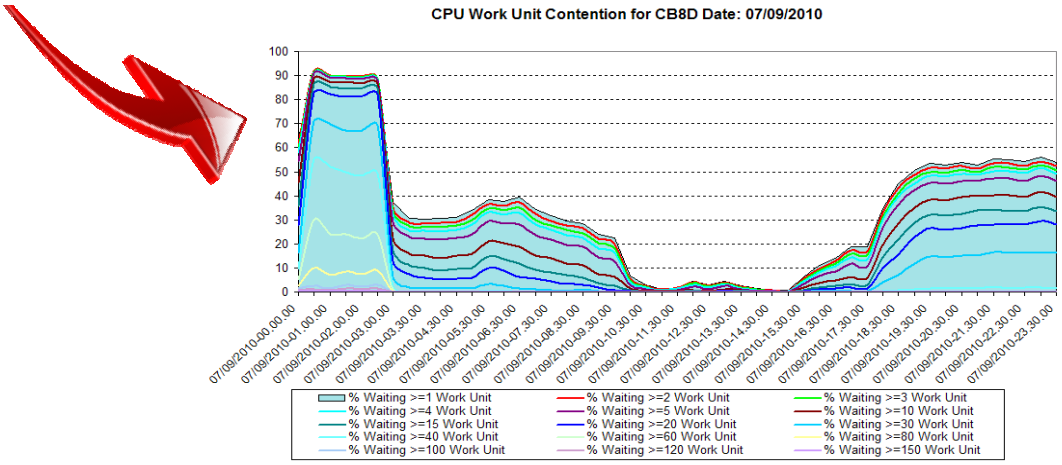


- The new ***OneWebCont*** sheet in the System Overview Spreadsheet of the RMF Spreadsheet Reporter allows to analyze cpu contention on work unit granularity.

# IBM z/OSMF Management Facility

*z/OS application, browser access*

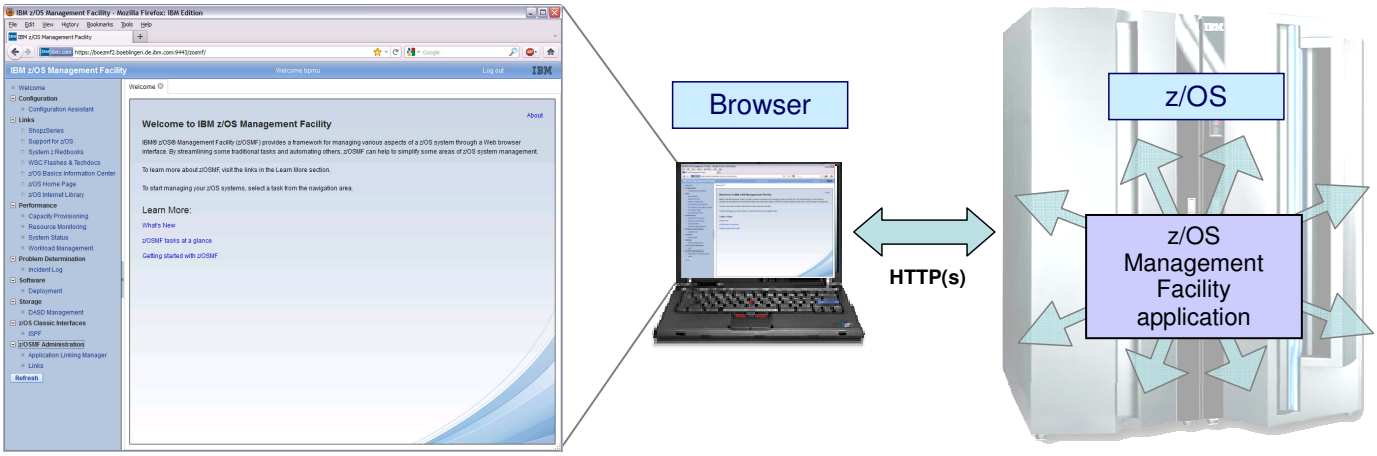


- *z/OS Management Facility is a Web 2.0 application on z/OS*
  - Manages z/OS from z/OS
  - Browser communicates with z/OSMF via secure connection, anywhere, anytime

- IBM z/OS Management Facility (z/OSMF), a new zero priced product, will simplify, optimize and modernize the z/OS **system programmer** experience starting in V1R11 with problem data management and TCP/IP Policy based configuration. With z/OS V1R12 the Workload Management and the Resource Monitoring Application have been added
- z/OSMF will deliver solutions in a task oriented, Web browser based user interface with integrated user assistance. And z/OSMF will make the day to day operations and administration of the mainframe z/OS systems easier to manage for both new and experienced system programmers. The focus is to help improve system programmer productivity, and make the functions easier to understand and use
- This chart explains the structure for z/OSMF and how it fits into the z/OS environment.
- z/OSMF is on the right hand side and it manages z/OS from z/OS itself. It is not an external application, nor does it have an external client. z/OSMF is an application on z/OS with direct access to z/OS data and information, and it has a browser interface from the workstation. z/OSMF contains the GUIs and the application code. Everything is installed on the z/OS server and there are no client side install requirements.
- In the middle of the screen is a workstation with a browser and it communicates with z/OSMF via HTTP(s). z/OSMF is a Web 2.0 based solution. It incorporates a browser interface that communicates with the z/OS system. The browser can be anywhere... in the data center ... or around the world. You just need a secure connection.
- And on the left is a screen capture of the z/OSMF welcome page once you log into z/OSMF.

# IBM z/OSMF Management Facility …

*Applications / R13 plugins*

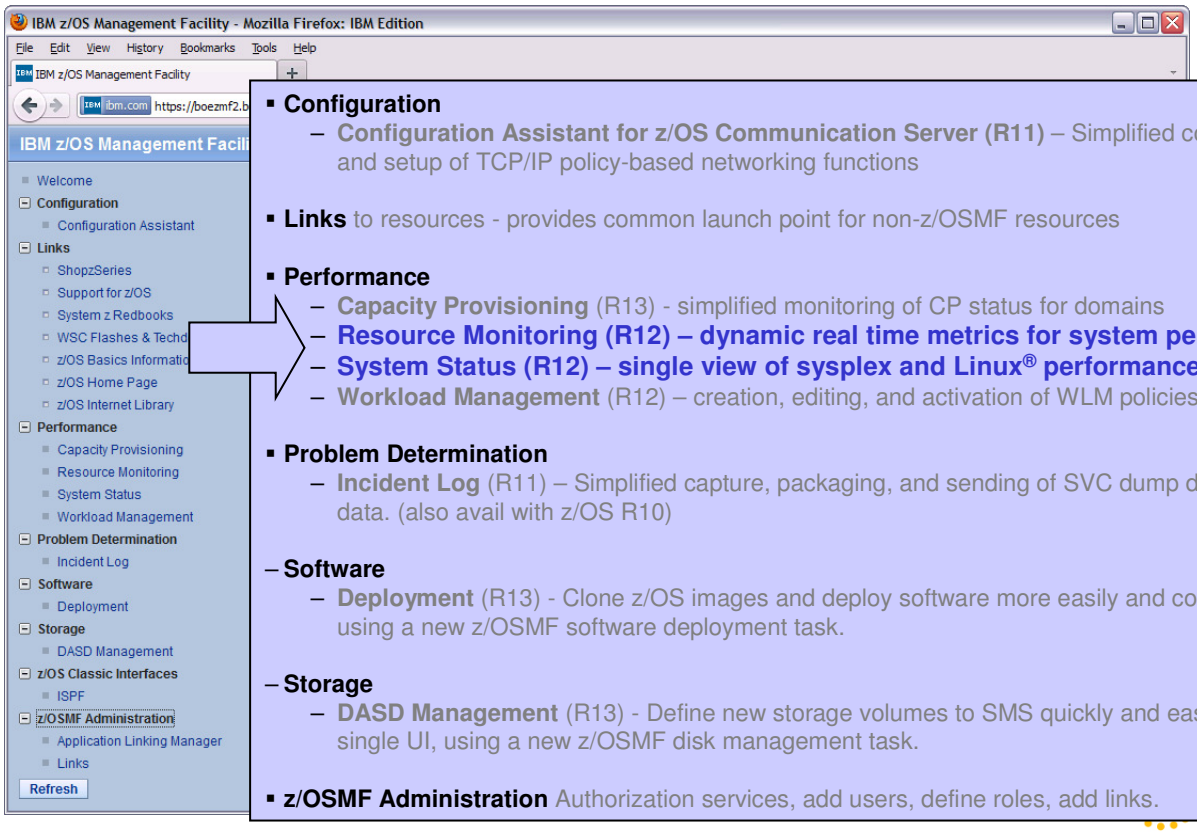


- Once you've logged in , this screen shows you the full scope of what z/OSMF provides in this first release. And really the first user that logs in as a z/OSMF administrator. So when z/OSMF is first set up, always the first ID is that of an administrator – that is a requirement for setup. And the reason for this it allows the first person to get in and to add and enable others.
- z/OSMF offers the following system management tasks:

- **Configuration** category: Simplified configuration and setup of TCP/IP policy-based networking functions
- **Links** category:
  - This is like 'My favorites' - Provides common launch point for accessing resources beyond the IBM z/OS Management Facility.
  - Some links are pre-defined in the product.
  - The administrators can define additional links to share commonly used resources for their installation.
  - With z/OSMF R12 you can add 'links' and launch points to anywhere in the left hand navbar – and not just the "Links" category. You can customize /OSMF.
- **Performance** category
  - **Capacity Provisioning** (R13) **–** simplified monitoring o z/OS Capacity Provisioning status (monitoring only, not management) – as a reminder z/OS Capacity Provisioning is part of the base of z/OS and can automate System z On/Off Capacity on Demand.
  - The ***z/OSMF Resource Monitoring*** plugin allows cross-sysplex performance monitoring from a single point of control. From the z/OSMF task tree, you can select the following subtasks:
    - The ***System Status task*** provides an enterprise-wide health check of all z/OS sysplexes.
    - For further analysis, the ***Resource Monitoring task*** can graphically display RMF Monitor III metrics as well as Linux metrics by means of customizable views.
  - **Workload Manager Policy Editor** -- Facilitate the creation and editing of WLM service definitions, installation of WLM service definitions, and activation of WLM service policies
- Under the **Problem Determination** category on the navigation bar, you will find the **Incident Log task.**
  - This will help all system programmers with problem data management tasks, providing experienced teams with procedural advantages through an incident log summary and detail views of z/OS dump incidents. The Incident Log provides a consolidated list of SVC Dump related problems, along with details and diagnostic data captured with each incident. It also facilitates sending the data for further diagnostics.
  - Task is updated to include support for adding comments, the ability to attach user-defined data, and use of encrypted parallel FTP for sending data to IBM.
- **Software** category
  - **Deployment** (R13) - Clone z/OS images and deploy software more easily and consistently, using a new z/OSMF software deployment task.
- **Storage** category
  - **DASD Management** (R13) - Define new storage volumes to SMS quickly and easily using a single UI, using a new z/OSMF disk management task.
- **z/OSMF Administration** category
  - Provides Authorization services for the administrator
  - Allows Dynamic addition of links to non-z/OSMF resources

# IBM z/OSMF Management Facility

## *Resource Monitoring – System Status Task*

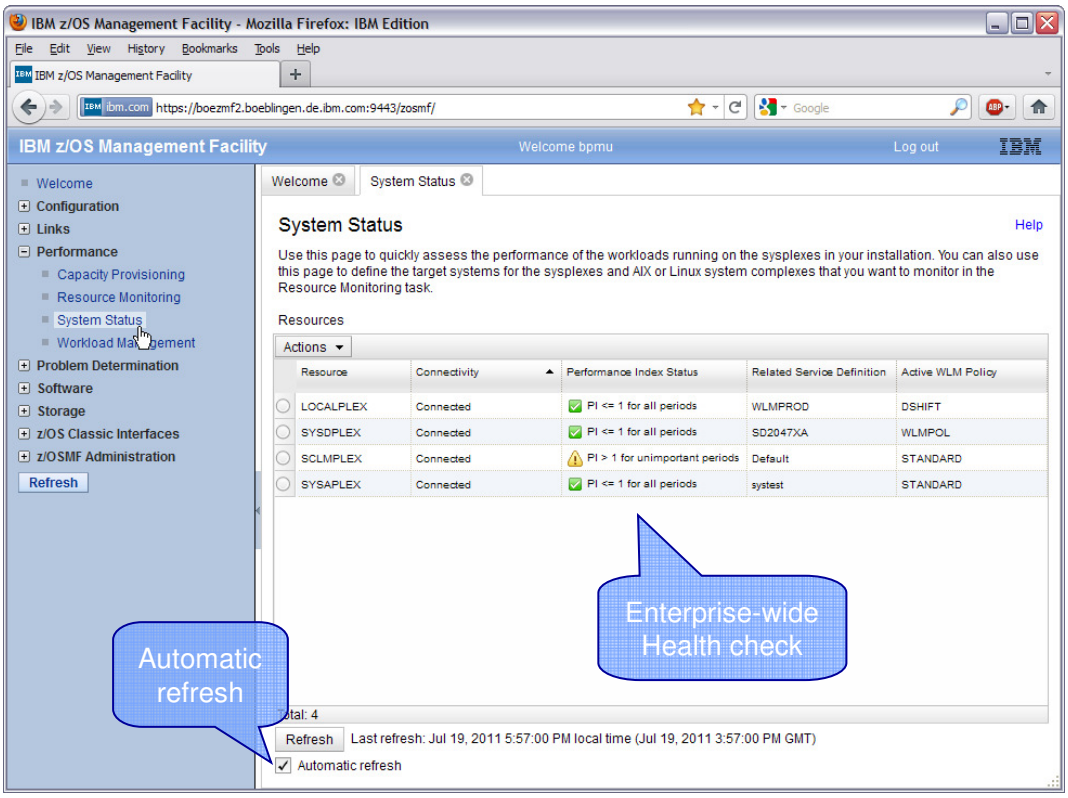


- The starting point for the monitoring is the *System Status* task. This task provides a comfortable way to assess the health status of all systems in your installation at a glance.
- The table contains the list of z/OS sysplexes and Linux images. You can add, modify and delete items in the table using the *Actions* menu.
- When you start the task for the first time, one default entry is provided as *LOCALPLEX* that points to the DDS in the sysplex in which z/OSMF is running. If you have a running DDS in this sysplex the *System Status* task detects it automatically without requiring an explicit declaration of its host name or IP address. Also, the *System Status* task can automatically detect the DDS movement within the sysplex during the z/OSMF session.
- In addition to the *LOCALPLEX* you can add all target sysplexes you want to monitor to the table. Now let's explore the column output as shown in the chart:
  - The PI-Status column gives you a red-yellow-green indicator for the sysplex health. This indicator is based on the WLM service class period goals and actuals.
  - If all service class periods on the system are meeting their goals (that is, have a *performance index* (PI) of less than or equal to 1), the PI Status is green.
  - If service classes with importance of 3, 4, and 5 with the PI greater than 1 are detected, the indicator is yellow.
  - If at least one important service class (that is, importance of 1 or 2) misses the WLM goal, the indicator is red
- In this case it becomes essential to figure out the reasons by drilling down into the details on the respective sysplex. This can be done within the *Resource Monitoring* task.

# IBM z/OSMF Management Facility …

## *Resource Monitoring – Monitoring Dasboards*

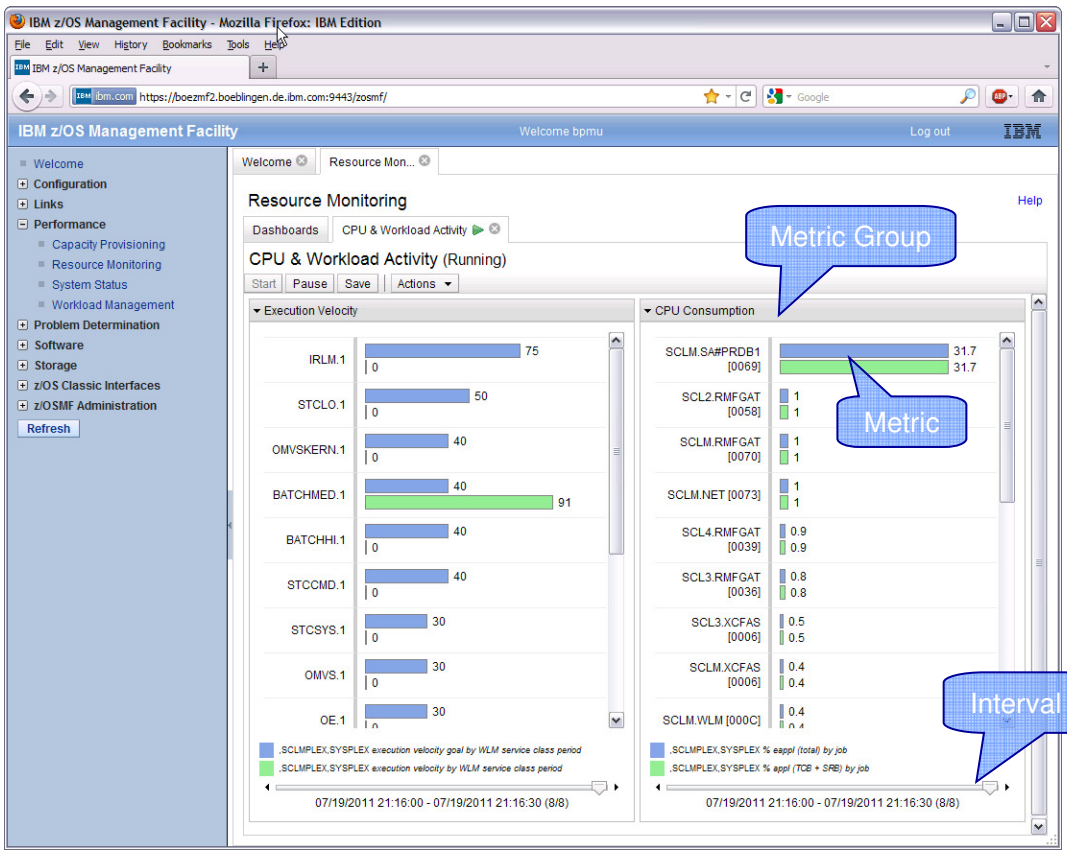


- The ***Resource Monitoring*** task allows you to monitor the performance of the z/OS sysplexes, AIX system complexes (System p), Linux system complexes (System z and System x) in your environment. The performance data is displayed in so called ***Monitoring Dashboards*.**
- A *Monitoring Dashboard* is a customizable view containing different performance metrics that you can group and arrange flexibly.
- You can create and save your own dashboards or open and modify the predefined dashboards that are with the task in z/OSMF.
- When you start a Monitoring Dashboard it begins retrieving the online data from the DDS periodically.
- By default, the most current snapshot is displayed
- It's possible to browse through the data collected since the start of the current session.
- A dashboard contains multiple metric groups; their size and column count are adjustable. You add one or multiple metrics to each group.
- Now let's create a new *Monitoring Dashboard*. Suppose, you want to inspect the common storage usage and identify jobs with the highest storage consumption.
  - To do this, you can create a new dashboard using the drop down menu on the Dashboards tab. You can now add a metric you want to monitor to the dashboard by using the "Add Metric" dialog.
  - The first thing you need to select in this dialog is a resource you're interested in. The resources are presented in a tree structure.
  - Starting with the sysplex as top level resource, you can drill down into the contained resources, for example, the MVS images or coupling facilities.
  - Once you select a resource (for example,  storage on a system), you can select a performance metric that you want to monitor.
  - The metrics of the selected resource are also presented in a tree. For a better navigation, you can use the quick-filter feature.
- *Resource Monitoring* provides sophisticated filtering capabilities:
  - Suppose you select the metric *% CSA utilization by job,* which will list all the jobs active on the system and their CSA consumption in %.
  - Because the number of jobs in the system can be pretty high, you can specify filters and work scopes: that is, display only jobs that match a certain name pattern, belong to a specific WLM service class, or meet a threshold value.

# IBM z/OSMF Management Facility …

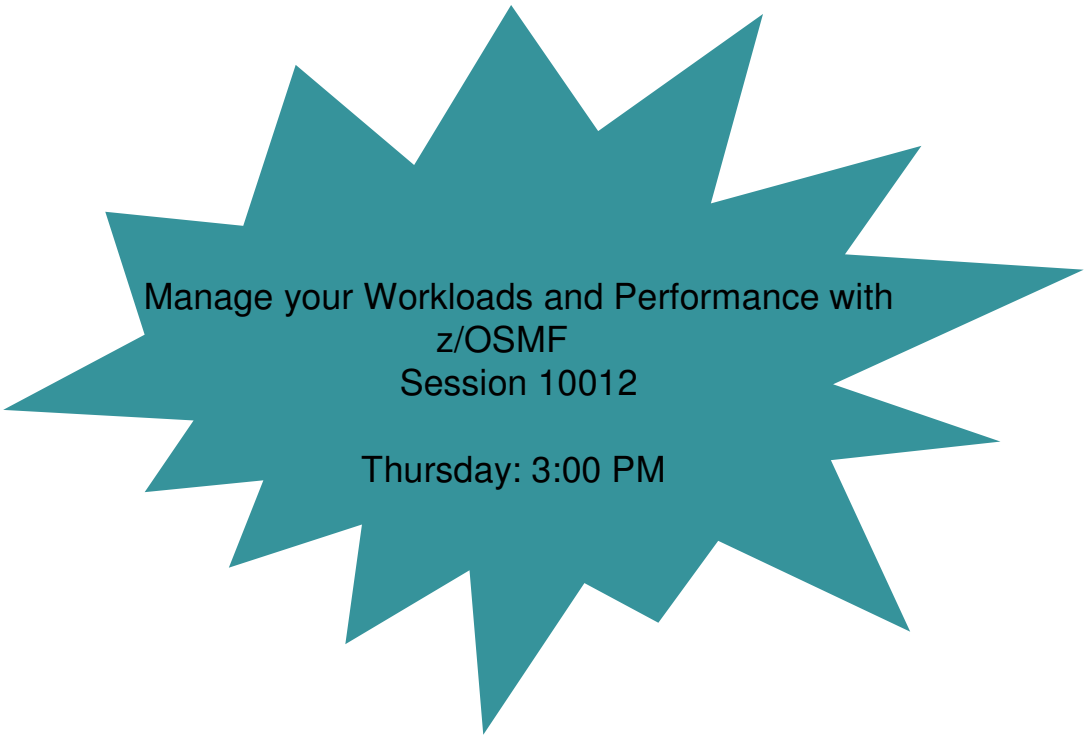

# RMF XP – What is it ?

- RMF XP is the solution for
  Cross Platform Performance Monitoring

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

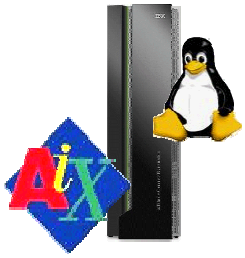


✓ Available with z/OS V1R13 RMF and
z/OS V1R12 RMF (APAR OA36030)

- RMF XP is the new solution to monitor the performance of heterogeneous environments. RMF XP supports the operating systems running on the IBM zEnterprise Bladecenter Extension:
  - AIX on System p
  - Linux on System x
- In addition, Linux on System z is supported as well

# 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

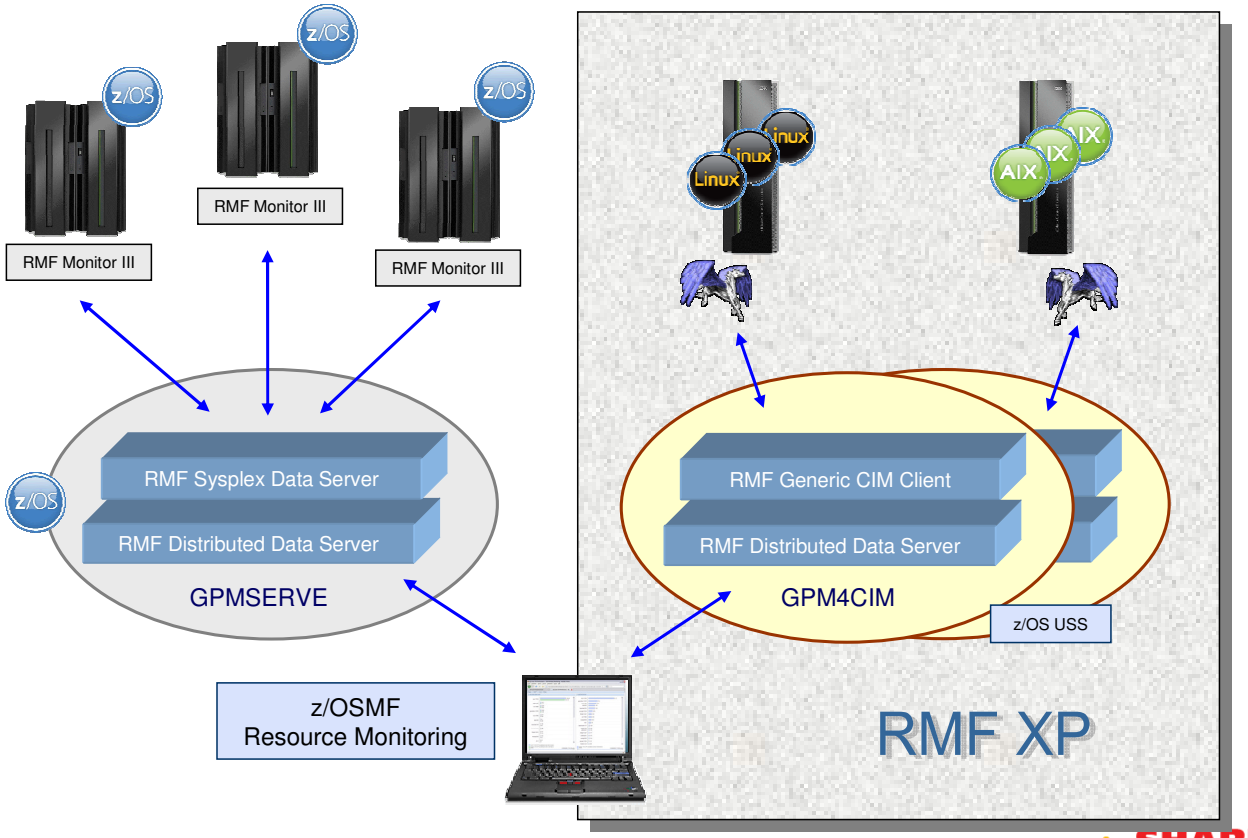

# RMF XP – Linux Topology

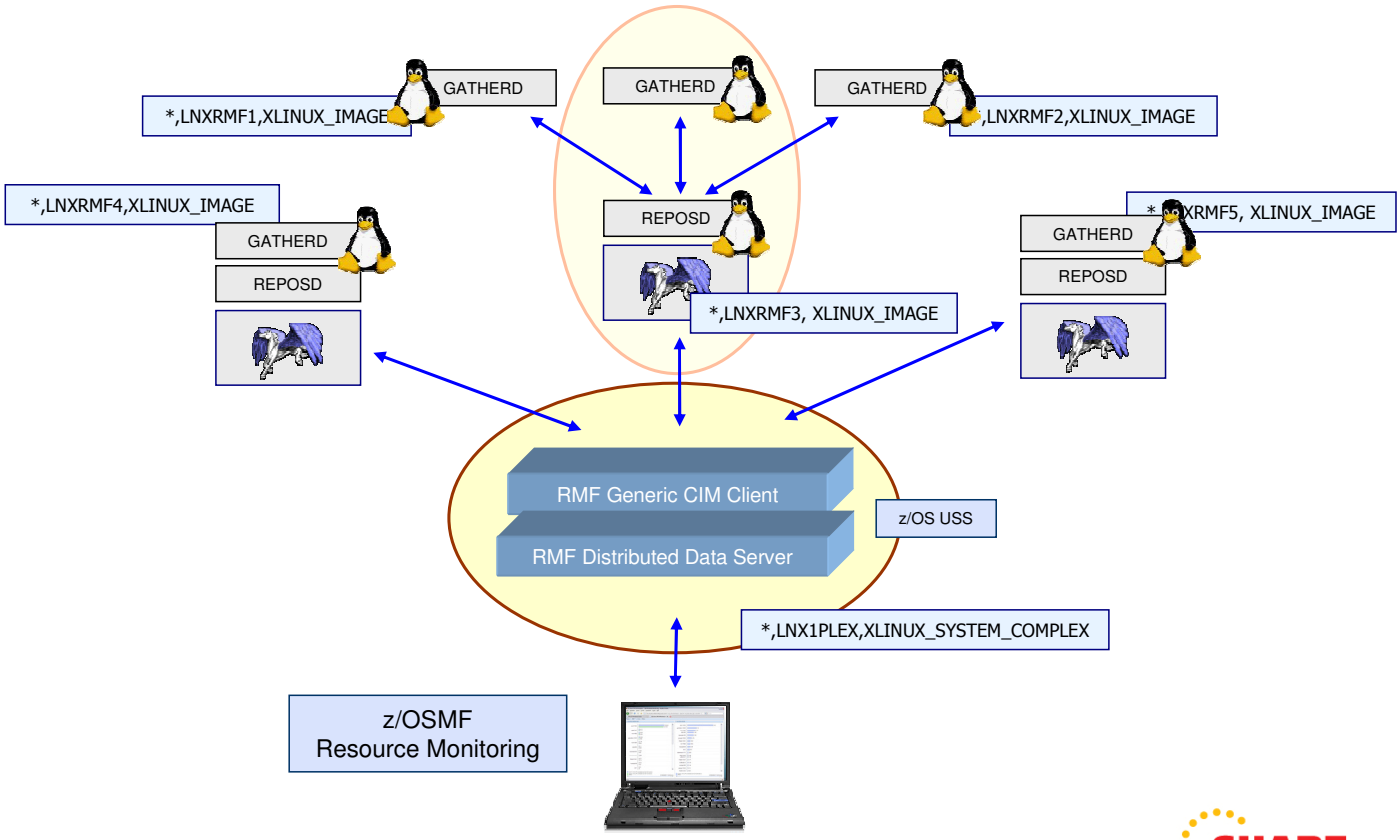

# RMF XP – AIX Topology

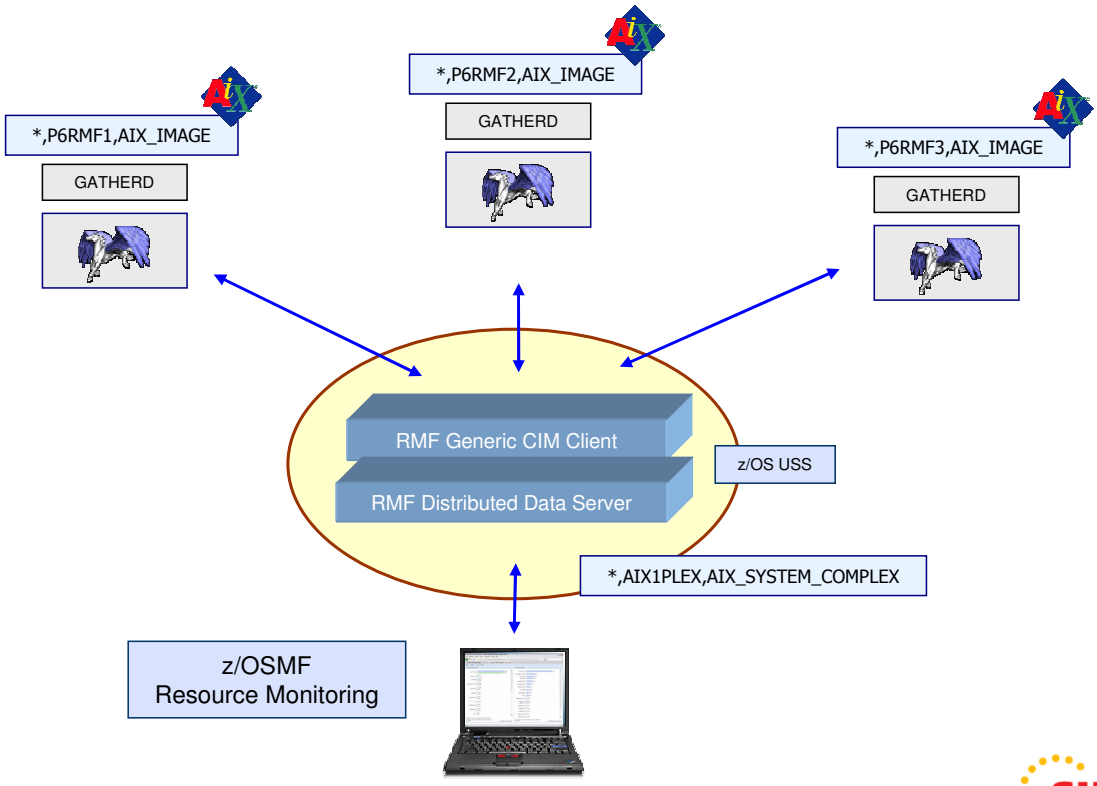

# RMF XP - Invocation

- 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/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
```

```
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 – Interval Synchronization

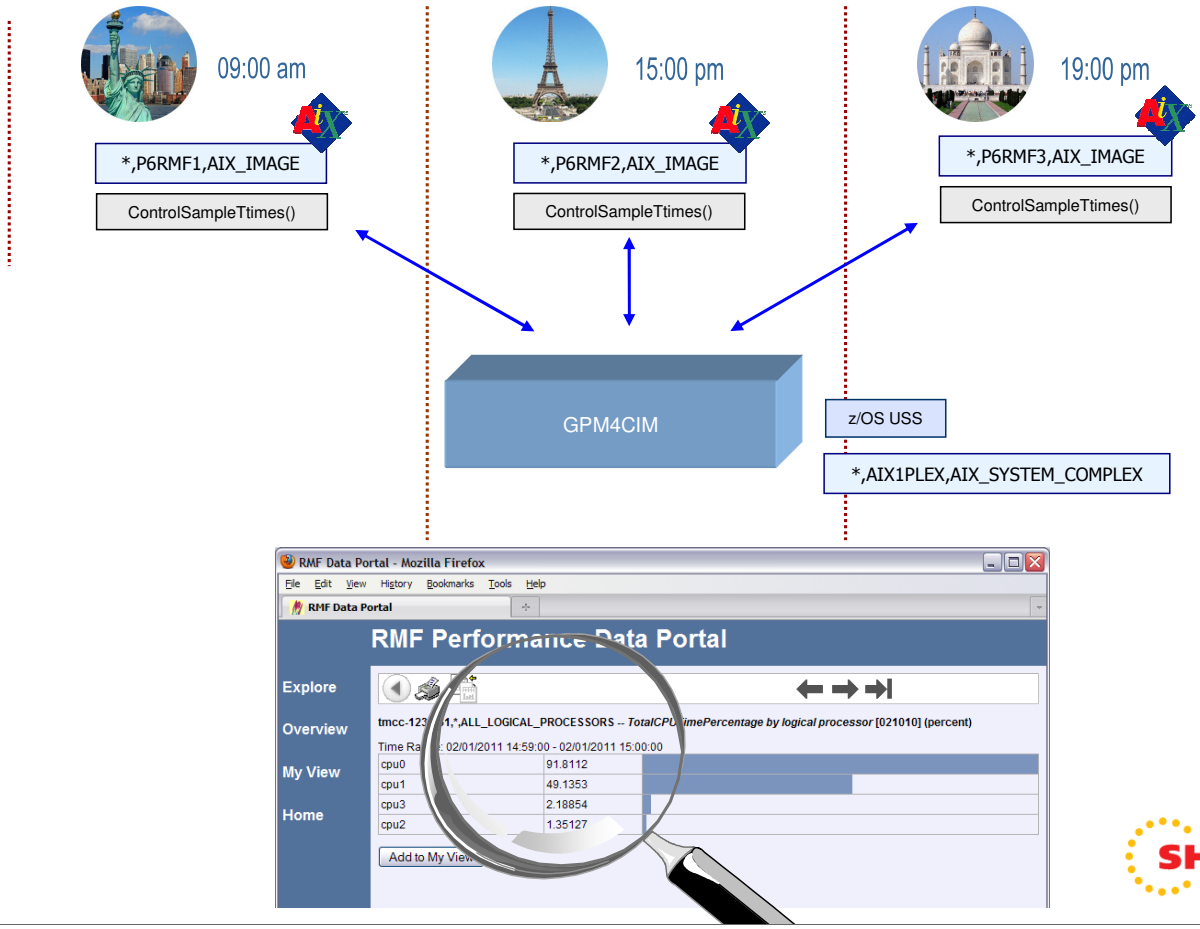


- The method call ControlSampleTimes allows to
  - start the data collection at a certain point in time
  - set a common collection interval for all monitored endpoints
- RMF XP syncronizes the interval
  - on a one minute boundary if the interval is < 5 minutes
  - on a five minute boundary if the interval is > 5 minutes

# RMF XP - Resource Models

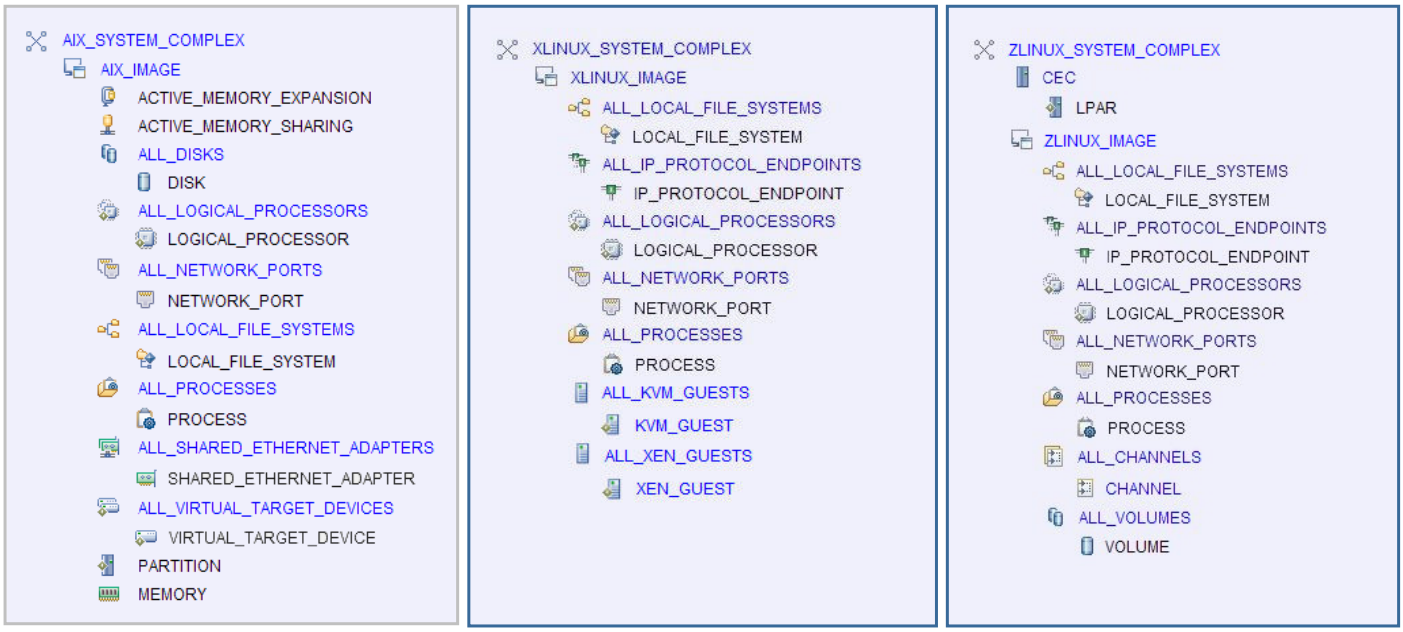

# RMF XP – Security

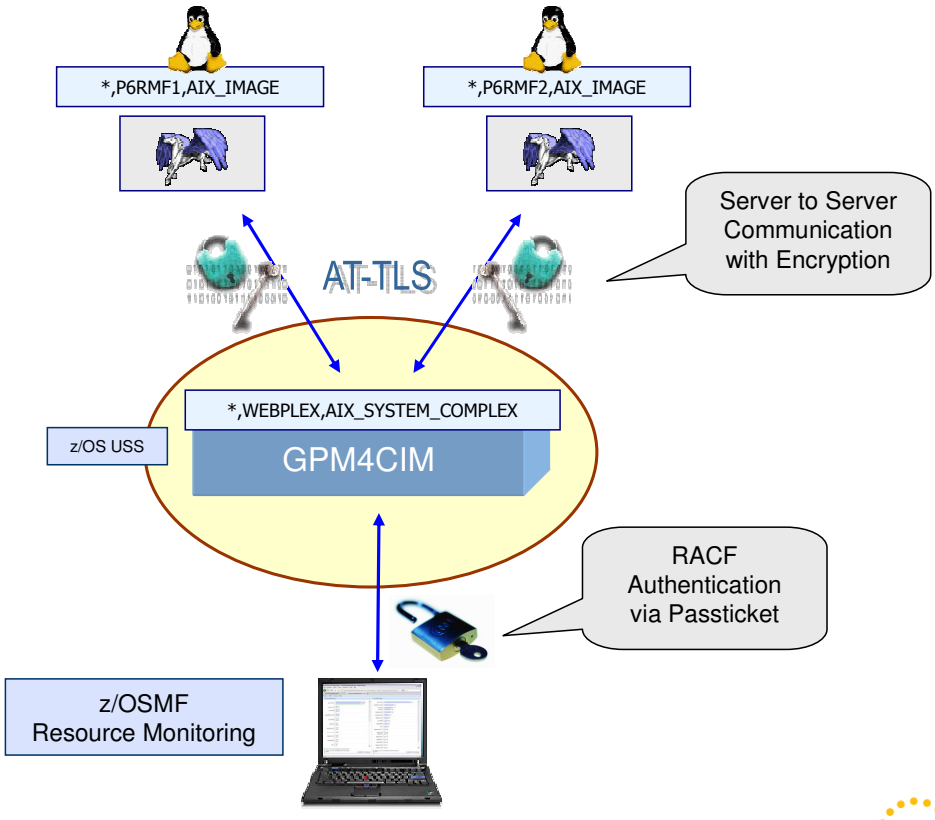

# RMF XP - zIIP Exploitation

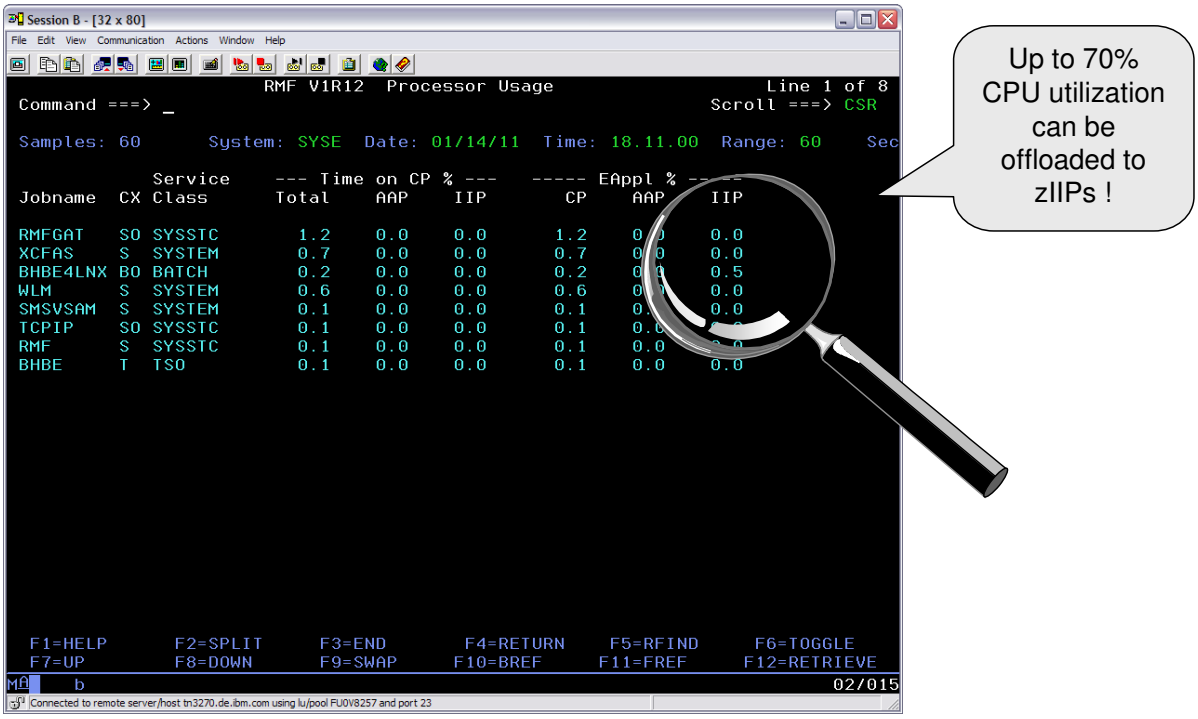

# RMF XP – Resource Tree

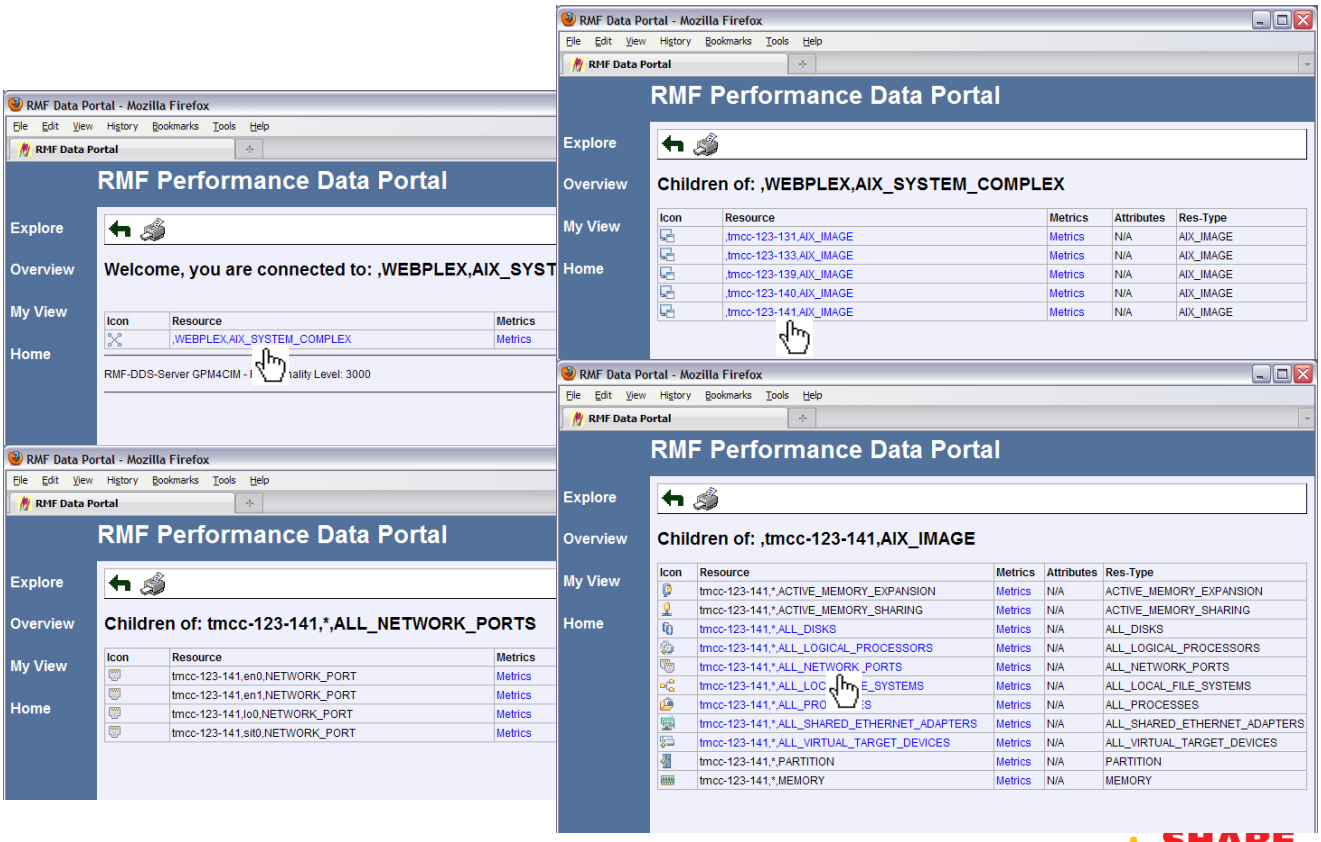

# RMF XP – Metrics

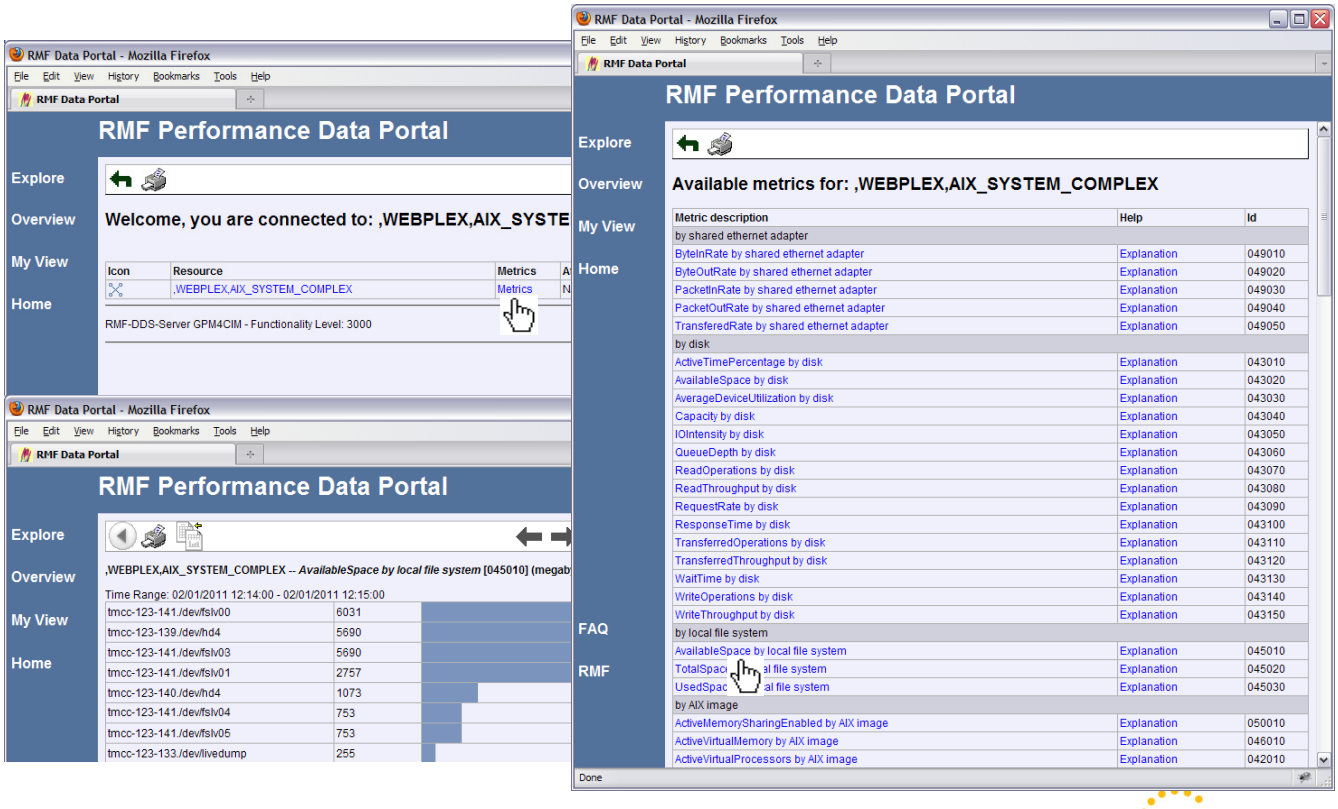

# RMF XP – Metric Values

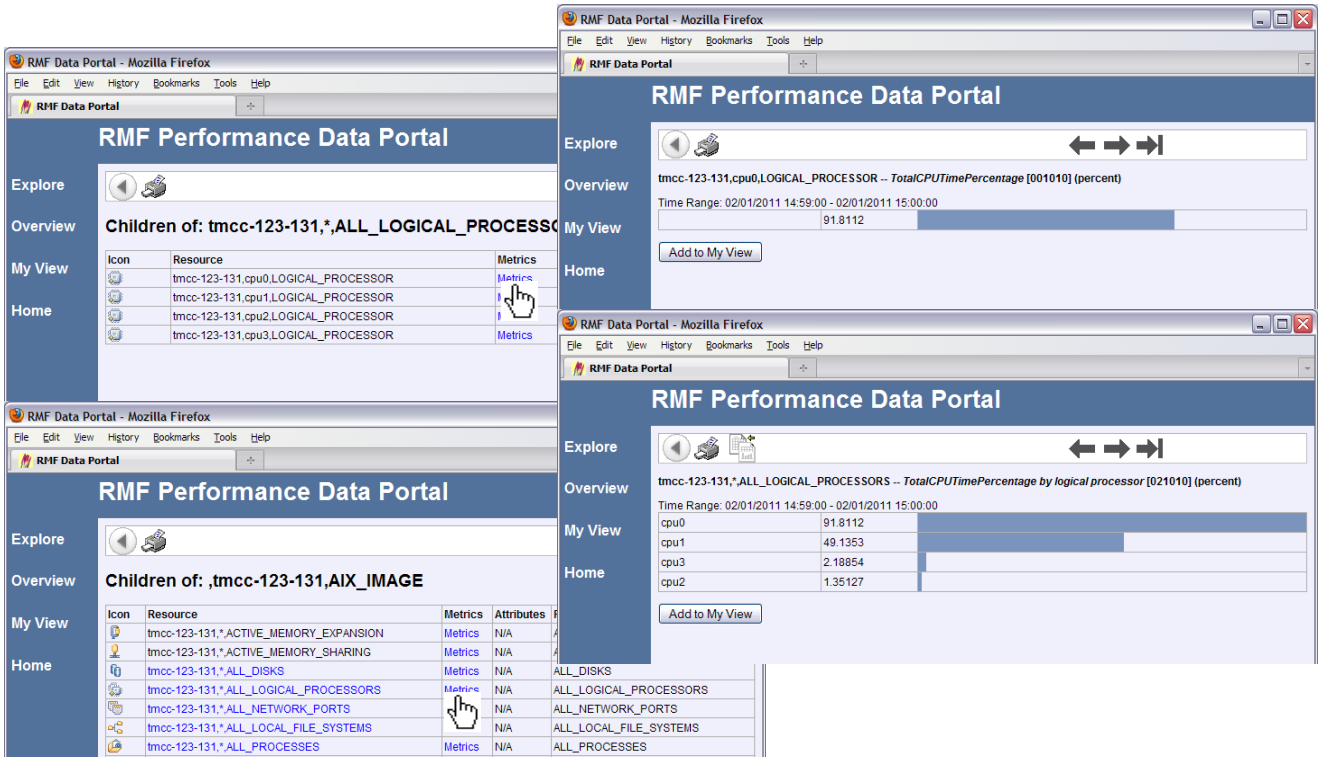

# RMF XP – Metric Scope

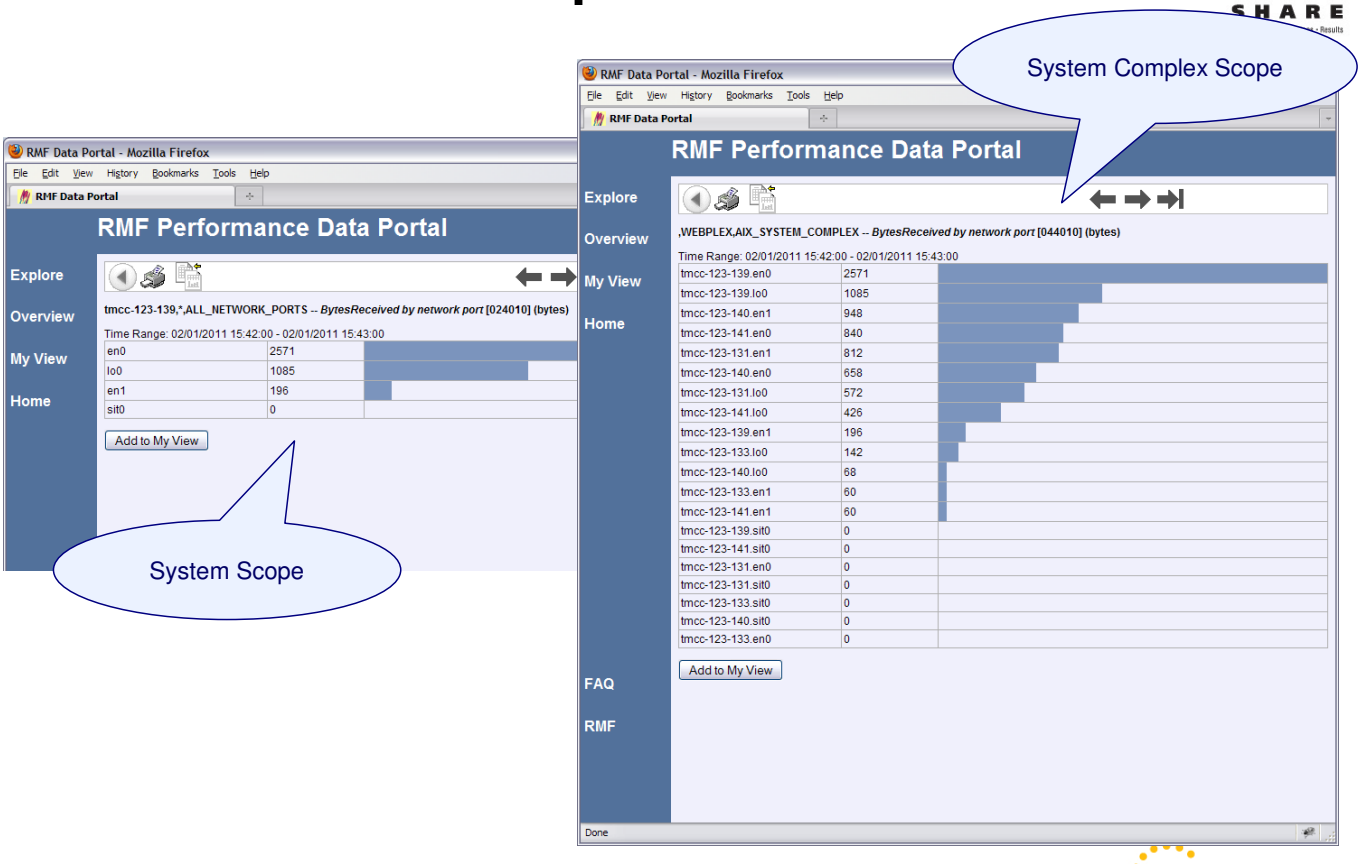

# RMF XP – z/OSMF Integration

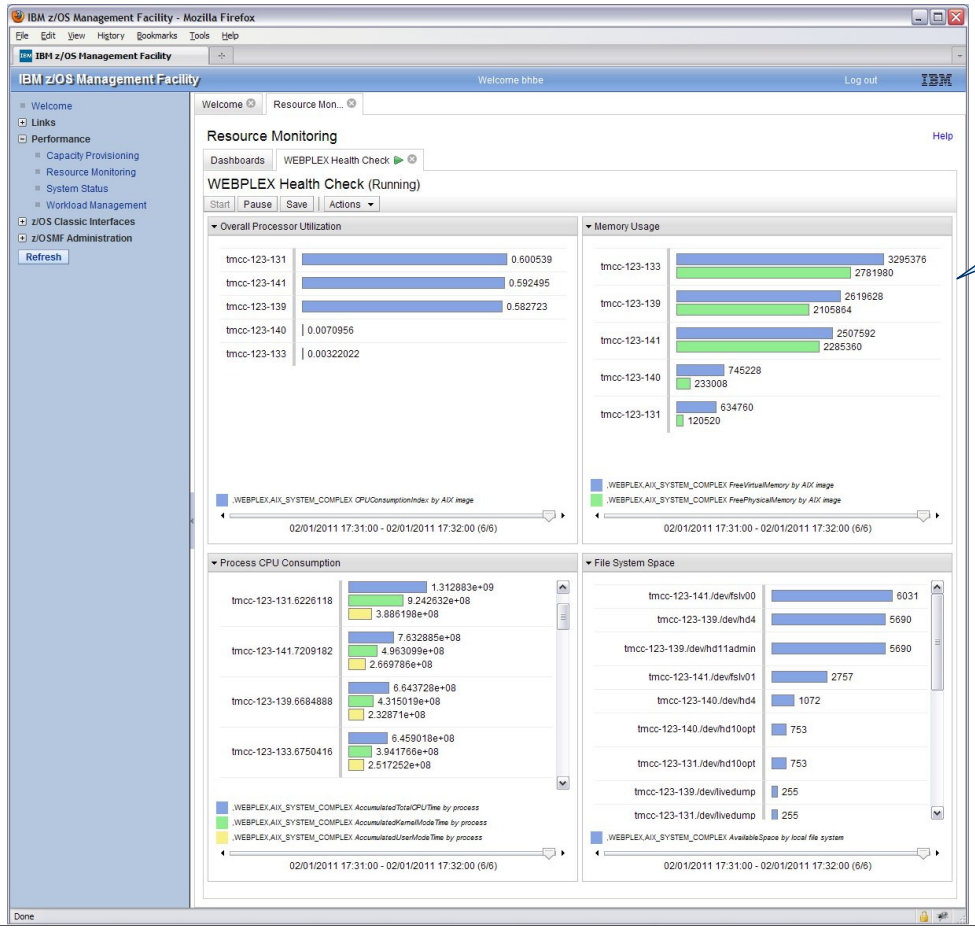


Health Check:

- ✓ 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)

- Alltogether, RMF XP provides a seamless, cross-plattform performance monitoring solution
- The new function strengthens the z/OS platform as management platform to administer and control distributed environments
- RMF XP works out-of-the box, in a couple of minutes everything is up-and-running
- The end-user has the choice between two graphical, state-of-the-art workstation frontends
- Approximately 70% of the CPU consumption can be offloaded to zIIP engines. This helps to reduce costs significantly.

# 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).

SHARE in Orlando 2011

# 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

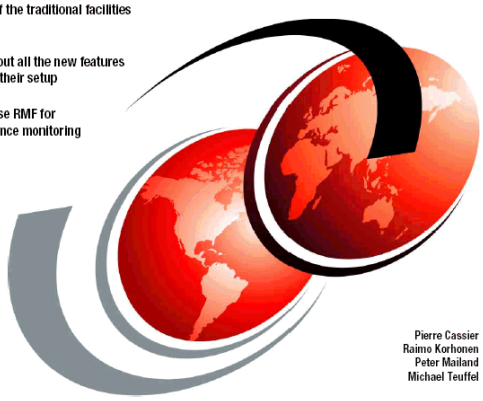

Pierre Cassier
Raimo Korhonen
Peter Mailand
Michael Teuffel

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:<br>Resource Monitoring Plugin | z/OS 1.12 zOSMF |
| RMF XP | z/OS 1.13 RMF<br>z/OS 1.12 RMF (APAR OA36030) |