#SHAREorg

SHARE
Technology · Connections · Results

Cross Platform Performance Monitoring with RMF XP

Robert Kieninger(KIENINGR@de.ibm.com)
IBM Corporation

Tuesday, March 13, 2012
Session 10892

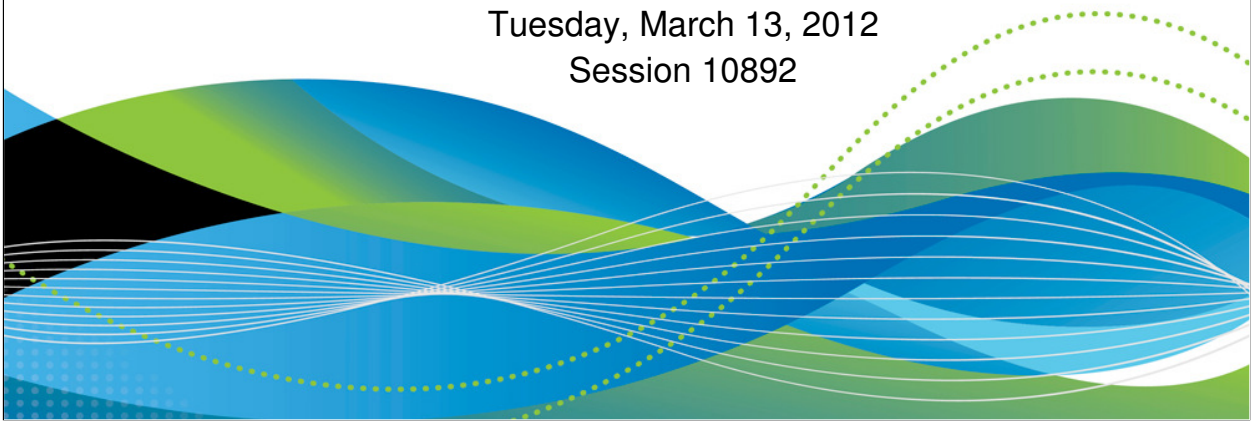

The new Component: RMF XP

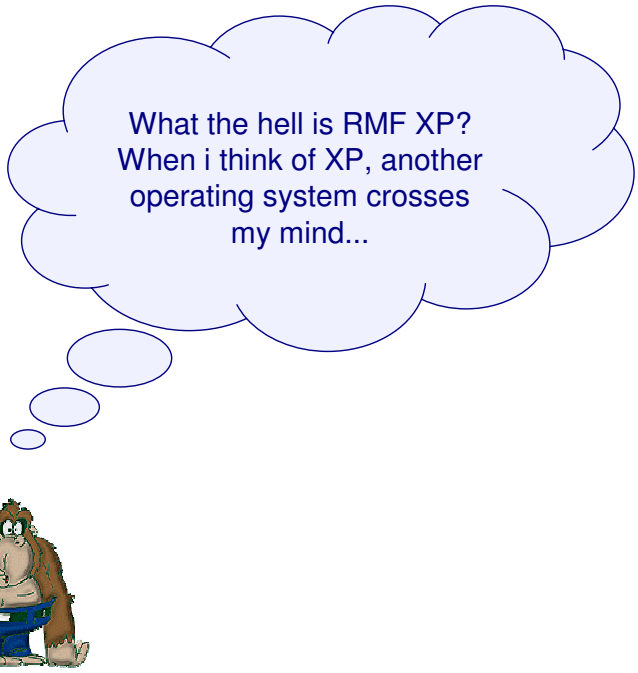

The new Component: RMF XP...

- RMF XP cannot cause Bluescreens!
- RMF XP is the solution for Cross Platform Performance Monitoring
- RMF XP supports the Operating Systems running on
  - x Blades
  - p Blades

- In addition RMF XP supports Linux on System z
  - LPAR Mode
  - VM Guest Mode

- 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 neat things together ?


We thought it is and we created the RMF XP

RMF XP – Component Overview

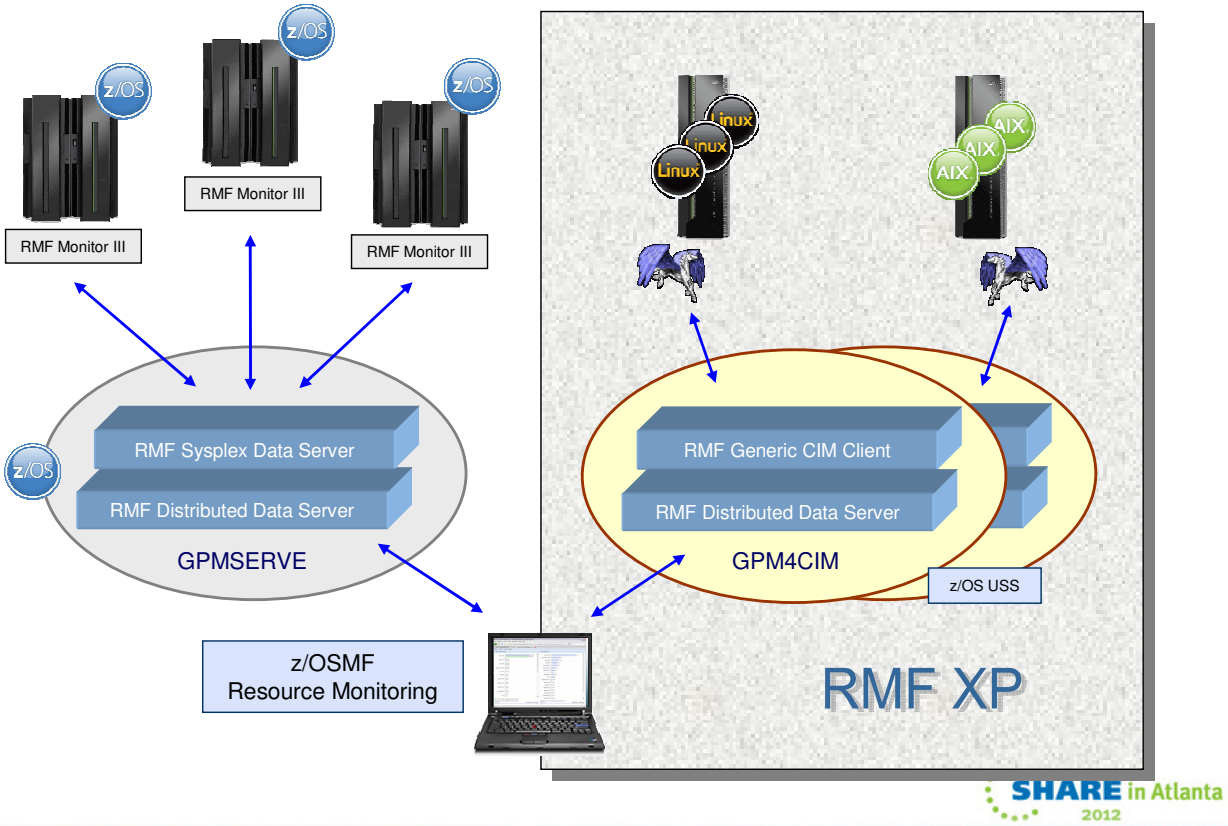

- What are the components of the new RMF XP function?
- The RMF Distributed Data Server (aka DDS) consists of two main functional entities:
  - The comunication layer to the client
  - The interface layer to the data collection backend
- For RMF XP, the existing DDS communication layer remains unchanged
- In contrast to the z/OS data collection which exploits the RMF Sysplex Data Server API's, RMF XP uses the standard CIM API's to collect the performance data from the remote Linux and AIX systems

RMF XP – Linux Data Collection

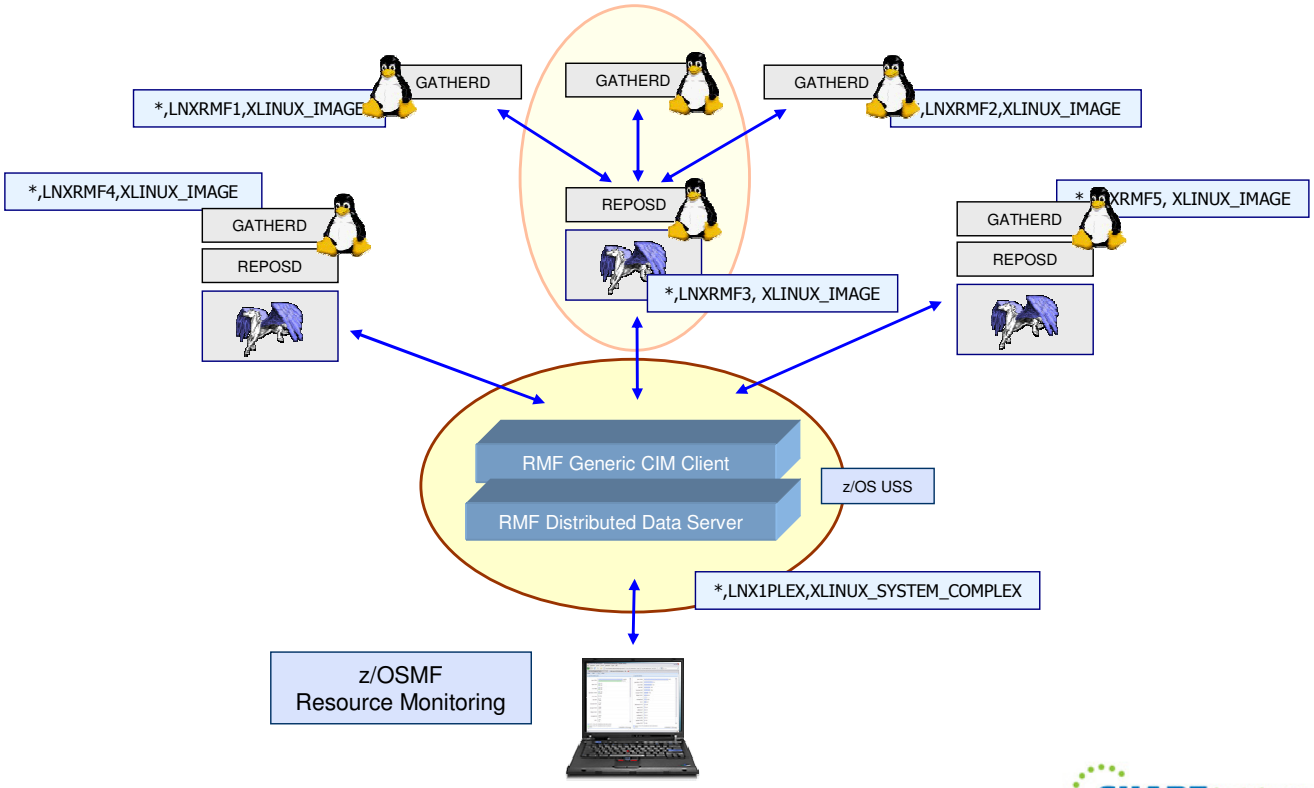

- The topology for the Linux data collection slightly differs from the AIX topology:
- On each individual enpoint, a performance data collector is needed in terms of a component called GATHERD
- The collector stores his data to a central repository which is managed by the REPOSD component.
- RMF XP can interact with this component by means of a CIM provider interface.
- Result: just one connection to a CIM server is needed to retrieve performance data from multiple Linux images

RMF XP – AIX Data Collection

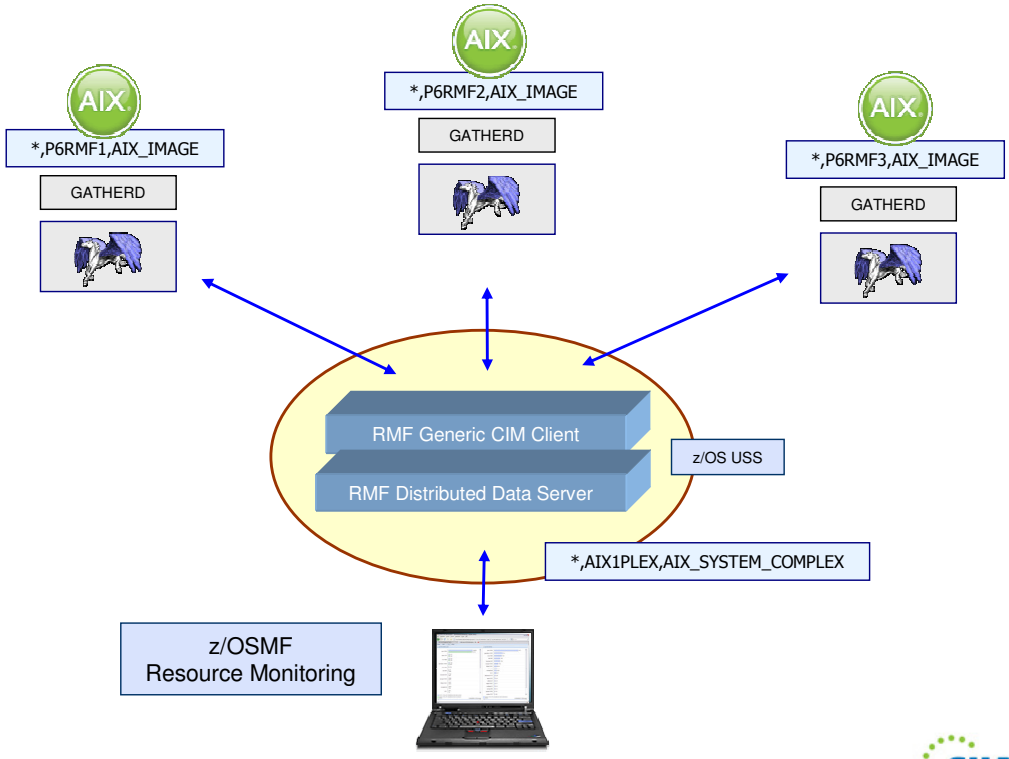

- In contrast to Linux, the concept of a repository node does not exist for the AIX operating system
- Hence, RMF XP builds up a seperate connection to each individual endpoint

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=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
```


- To start RMF XP, use the new proclib member GPM4CIM
- The gpm4cim executable runs in the Unix System Services environment and receives control from the BPXBATCH utility
- One GPM4CIM instance is needed per platform type

RMF XP – Configuration

```
//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 is almost an out-of-the-box function
- Just supply the following parameters:
  - Complex name (can be any string)
  - Image names (must be valid host names or ip addresses)
  - Interval length (optional parameter, default = 300 sec
- Even though the minimum interval length can be set to 60 seconds, it is recommended to choose longer intervals in order to prevent too much resource consumption

RMF XP – Interval Synchronization

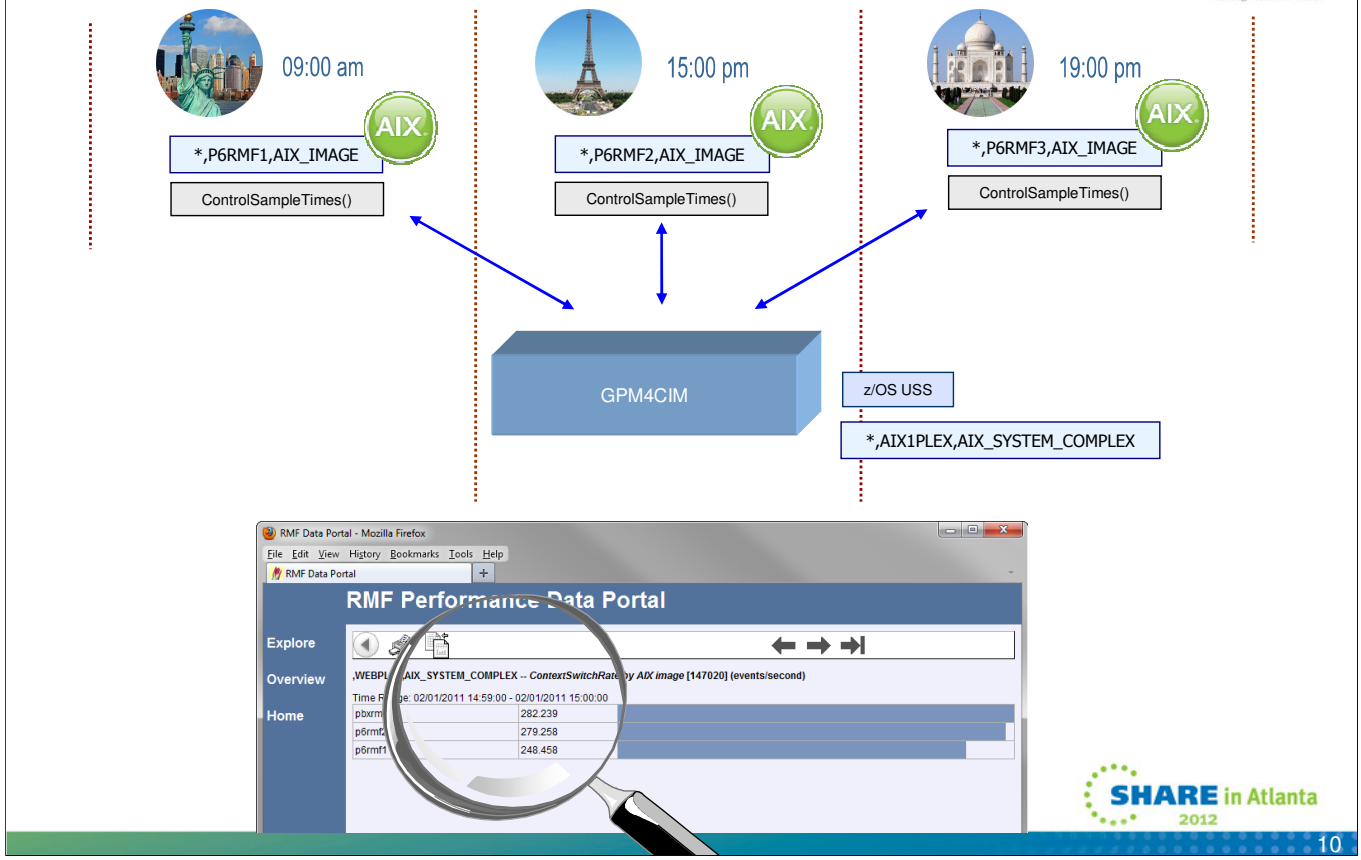

- For the AIX operating system the remote CIM API allows to
  - Start the data collection at a certain point in time
  - Set a common collection interval for all monitored endpoints
- RMF XP synchronizes the interval
  - On a one minute boundary if the interval is < 5 minutes
  - On a five minute boundary if the interval is > 5 minutes
- For the Linux operating system it is the responsibility of the system administratior to
  - Start the data collection on the monitored enpoints synchronously
  - Set the data collection interval on all endpoints accordingly to the RMF XP interrval length

RMF XP – Platform specific Resource Models

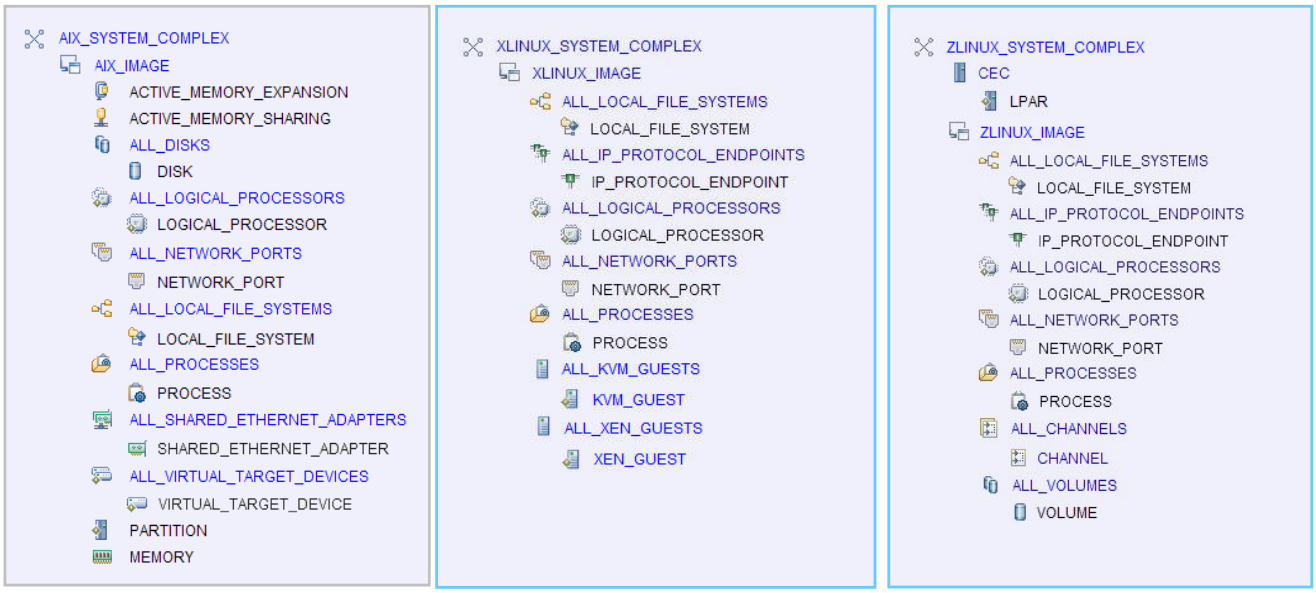

- If we could follow the pure nature of CIM, resources and metrics are common in a computing environment, regardless of the operating system.
- However, the reality looks a bit different. We have to deal with lots of platform specific extensions of the resource model.
- That‘s why we decided to supply one GPM4CIM instance per operating system type.

RMF XP – Security

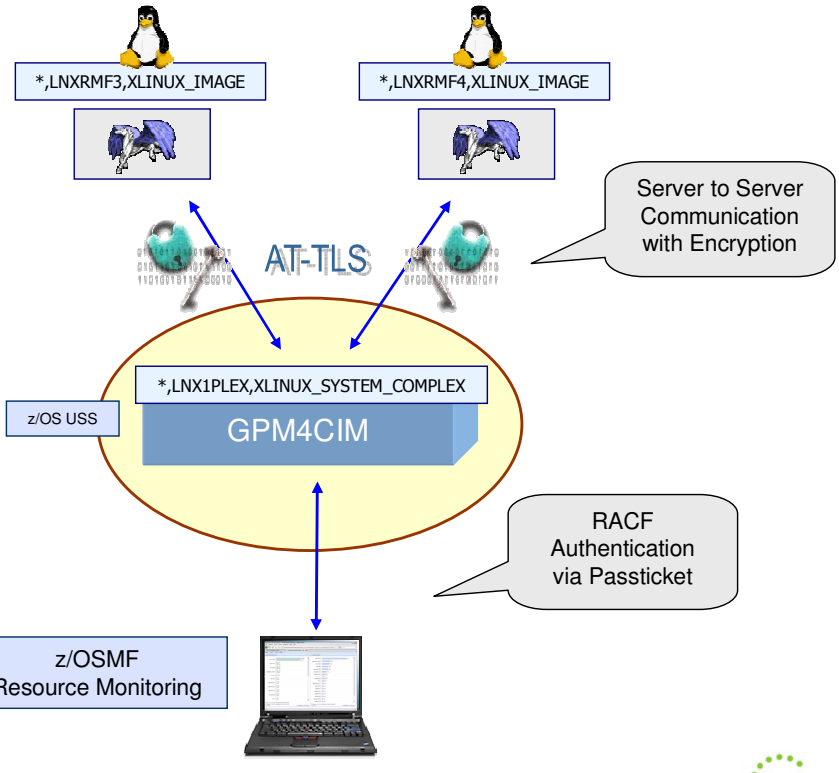

- Encryption for the communication between GPM4CIM and the endpoints can be configured via AT-TLS on the z/OS outbound side
- For the communication between the client and GPM4CIM authentication can be configured via userid and password or passtickets

RMF XP – zIIP Exploitation

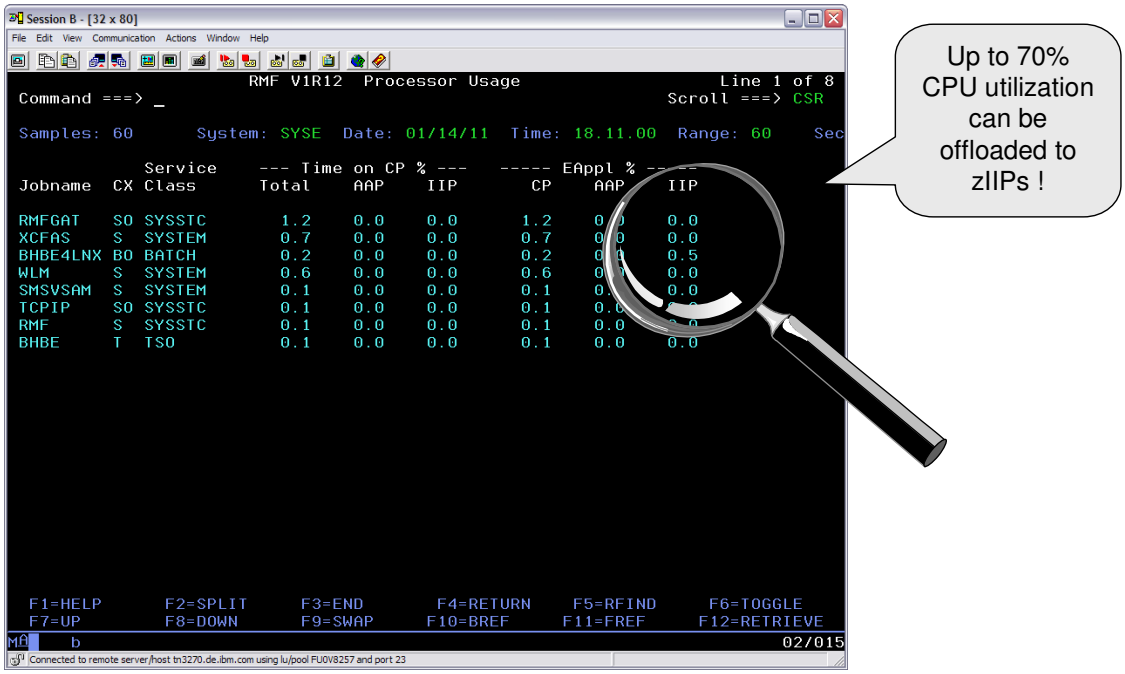

SHARE in Atlanta 2012

13

- Talking to lots of images and collect lots of performance metrics is not for free. It can become costly!
- However, RMF XP can exploit zIIP specialty engines for the CIM based data collection
- In our test environment, we have observed that approximately 70% of the CPU consumption can be offloaded to zIIP engines

RMF XP – Performance Considerations

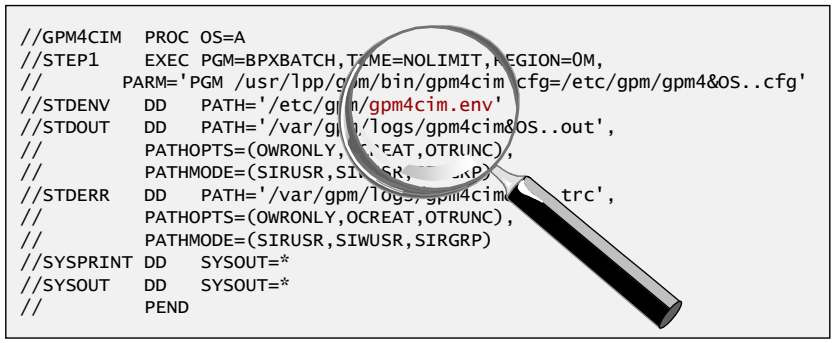

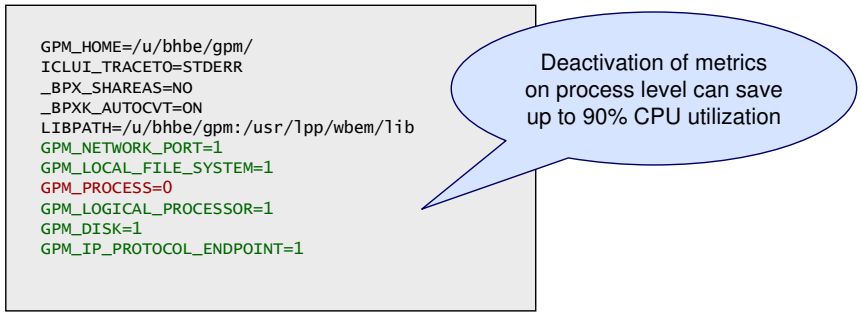

Exclude individual metric categories from the data collection

- Most variables within the GPM environment file are common for all platforms and should not be changed
- However, individual metric categories can be included in – or excluded from – the data collection by means of environment varables
- This feature is only applicable for metric catgories whith multiple instances per system
- Specifically for metrics on process level, the deactivation can save up to 90% CPU utilization
- The remote deactivation of metrics categories depends on the CIM server implementation resp. the ControlMetricsByClass() method
- This method is currently only valid together with the AIX CIM server
- Anyway, for the Linux platforms it is possible to exclude individual metric categories as well. But this cannot achieved by means of a remote API
- It is the reponsibilty of the administrator to deactivate specific provider modules on the endpoints manually by means of the *unload* command

RMF XP – Resource Tree

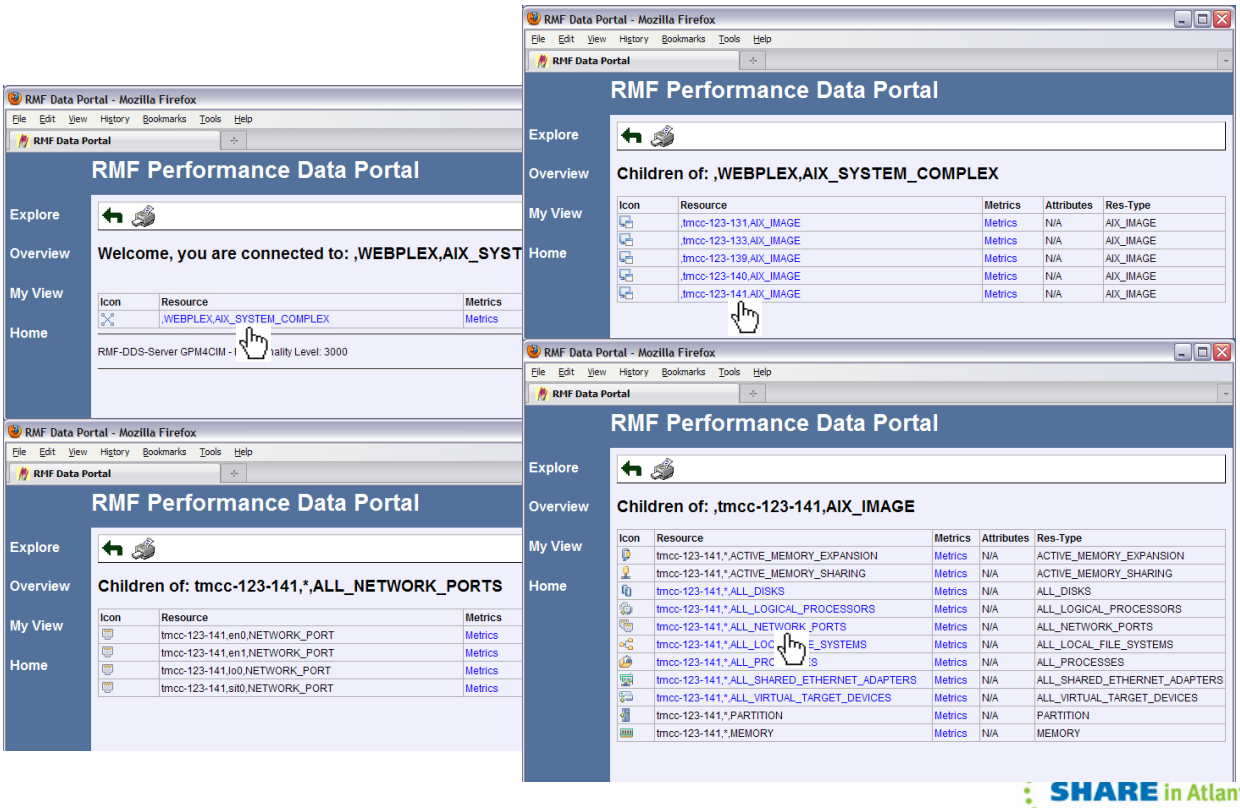

RMF XP – Metrics

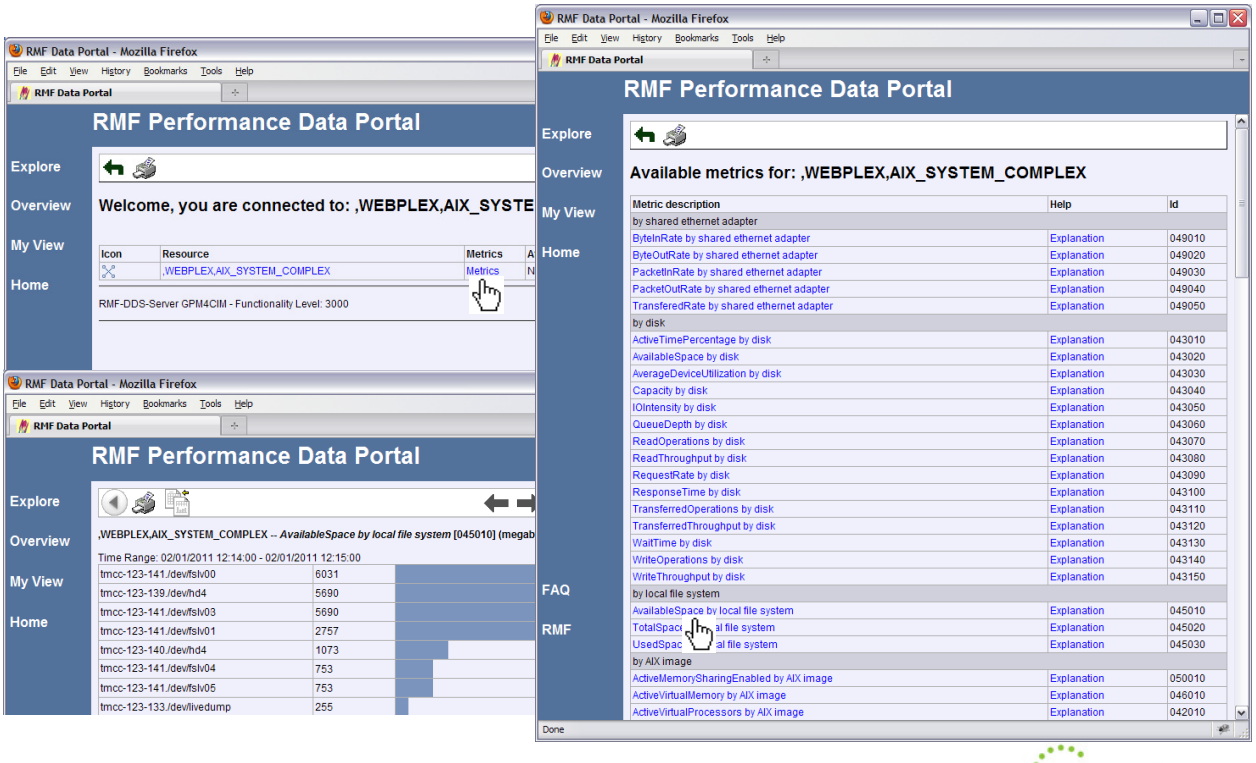

RMF XP – Metric Values

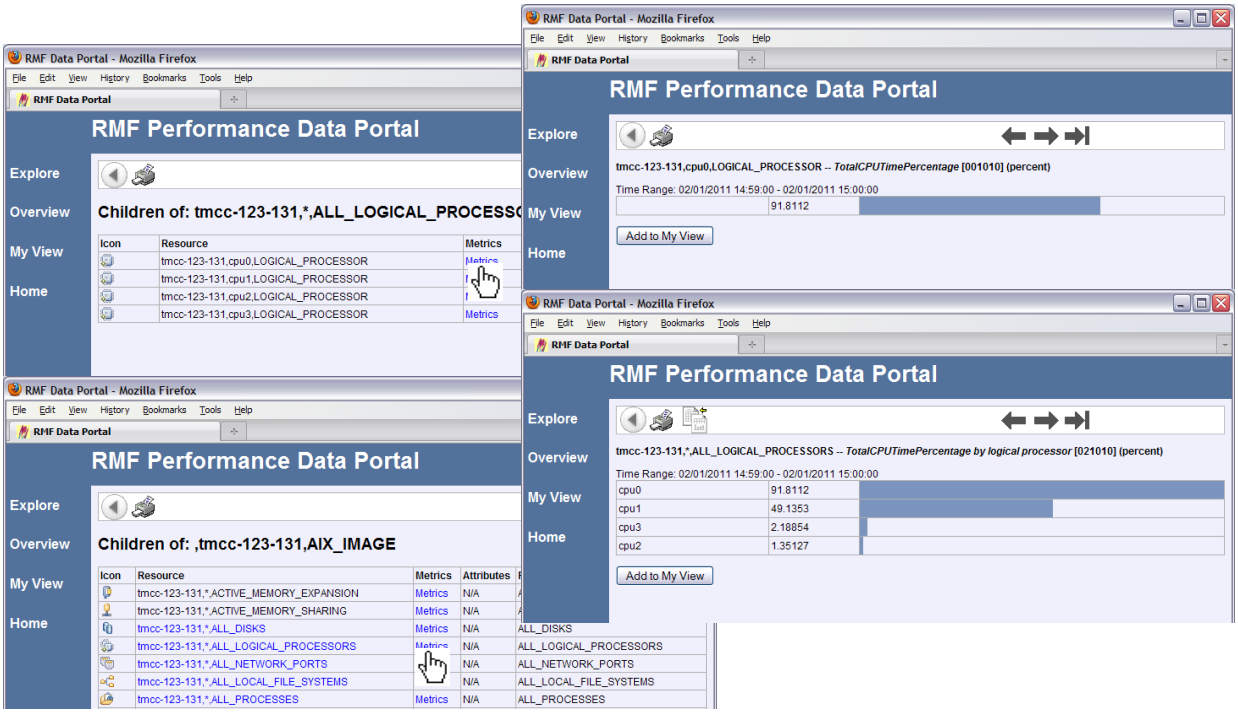

RMF XP – Metric Scope

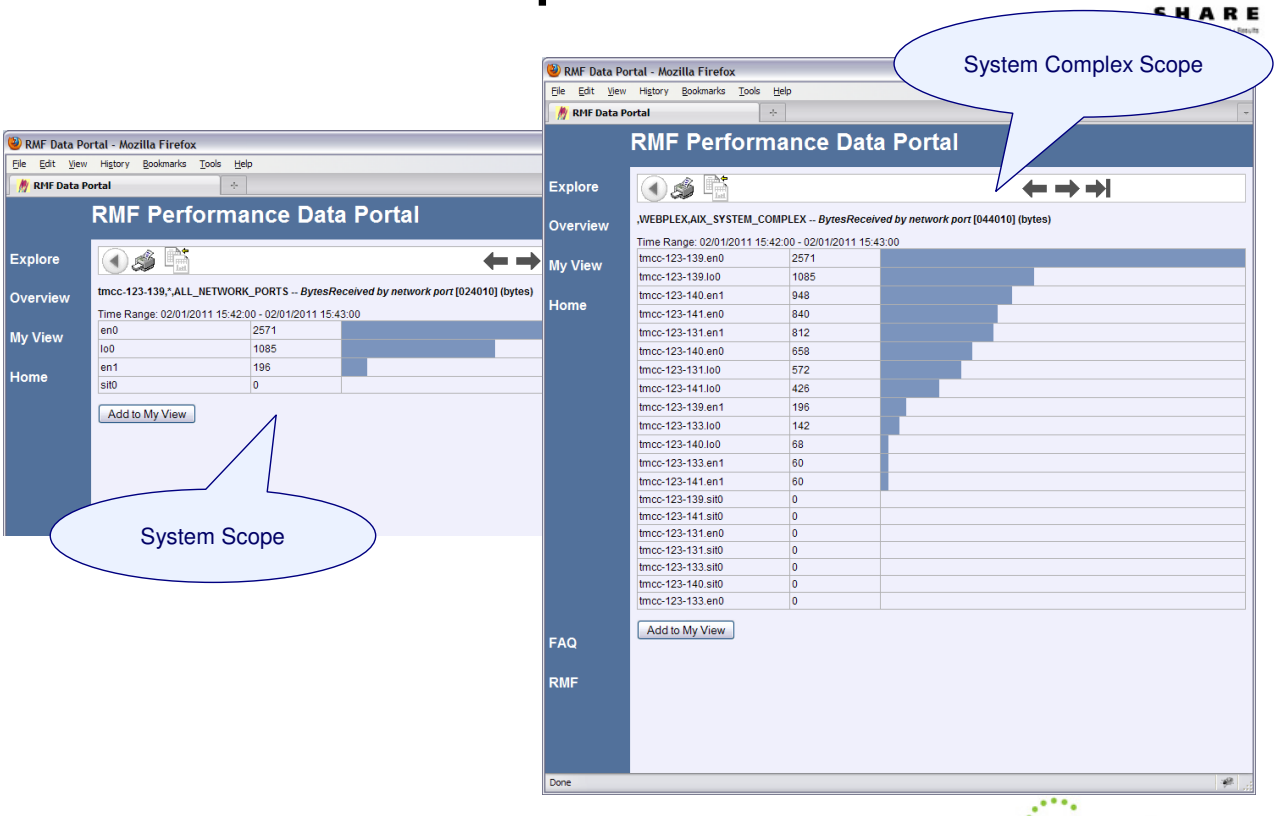

RMF XP – z/OSMF Integration

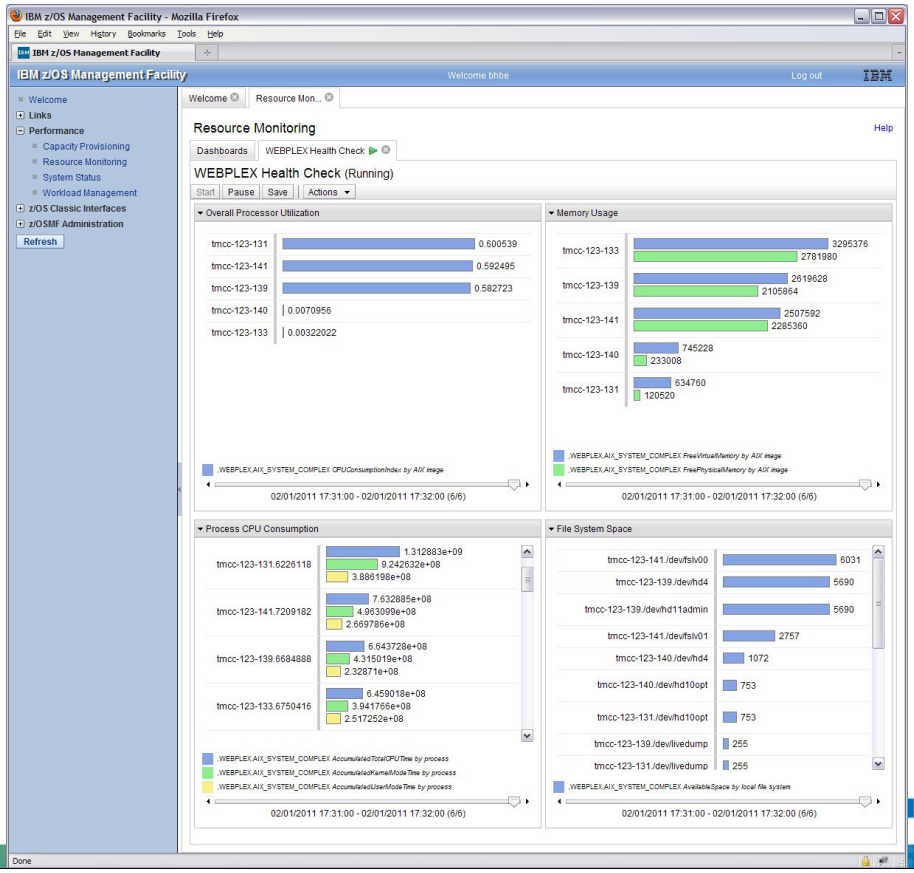

- The new RMF XP capabilities are completely applicable in the z/OSMF Resource Monitoring plugin
- Hence, you are able to monitor seamlessly all the systems within your enterprise, independent of the operating system type
- In a short time frame you can take advantage of the various capabilities of z/OSMF Resource Monitoring:
  - One workstation as single point of control
  - Dashboards with state-of-the-art graphical views
  - Maximum flexibility with user defined dashboards
  - Powerful data reduction and filtering

RMF XP – z/OSMF System Status Task

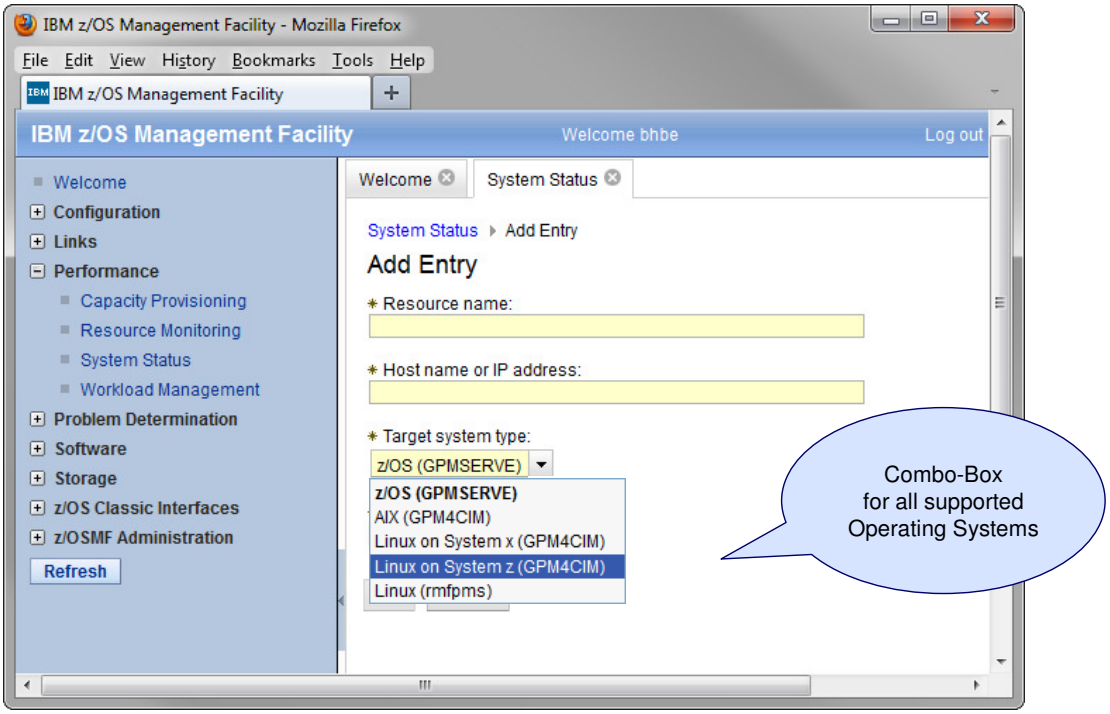

- The z/OSMF *System Status* task let's you easily define the new Operating System types
- From the *Target system type* combo-box just select the Operating System that your GPM4CIM instance is actually serving
- Then enter a valid hostname or IP address and a free selectable resource name, which denotes the top-level resource of the resource tree
- Now your definition is instantly visible in the *System Status* task and ready for use in the *Resource Monitoring* task without limitations

RMF XP – Resource Monitoring Task

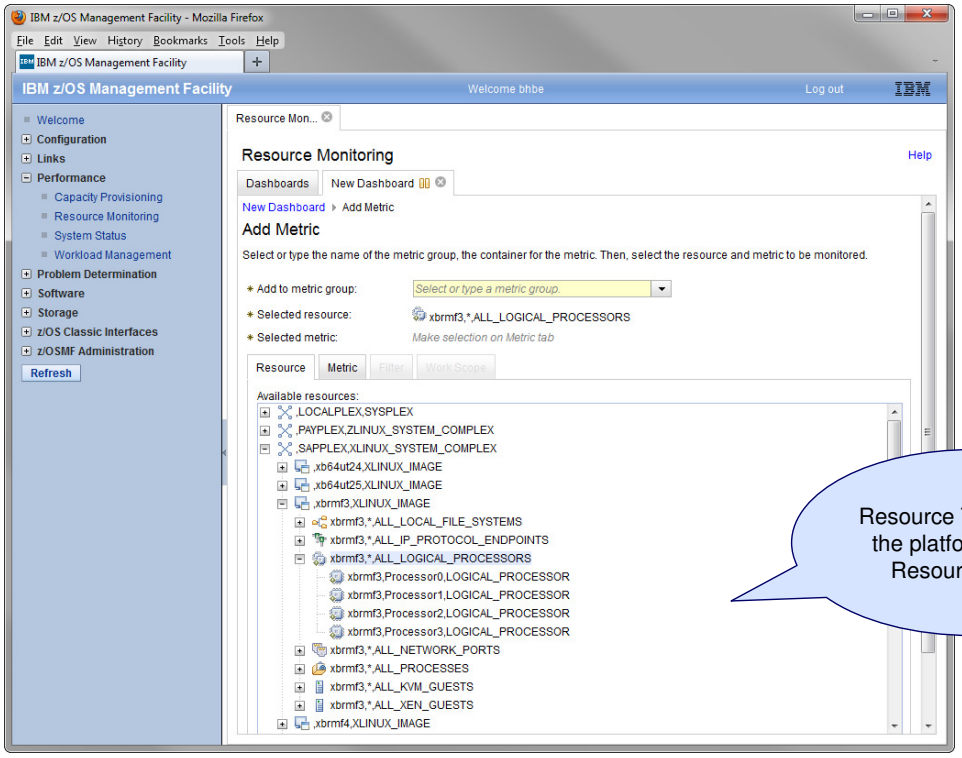

- Once you define a new metric in the z/OSMF Resource Monitoring Task, you start with the selection of the associated resource
- The *Available resources* view shows you all defined top-level resources, regardless of the operating system type
- Now you are able to expand the resource tree and the individual resources will become visible accordingly to the platform specific resource model
- All subsequent working steps – metric selection, filter definition etc. – are identical for all operating system types

RMF XP & z/OSMF – Single Point of Control

SHARE in Atlanta 2012

22

- One workstation can serve as singlepoint of control when RMF XP is used together with the z/OSMF Resource Monitoring facility
- The user can customize his own Dashboards for each platform and run all the Dashboards in parallel
- Hence, each platform can be monitored in the same browser session by simply switching between tabs

RMF XP & z/OSMF – Single Point of Control

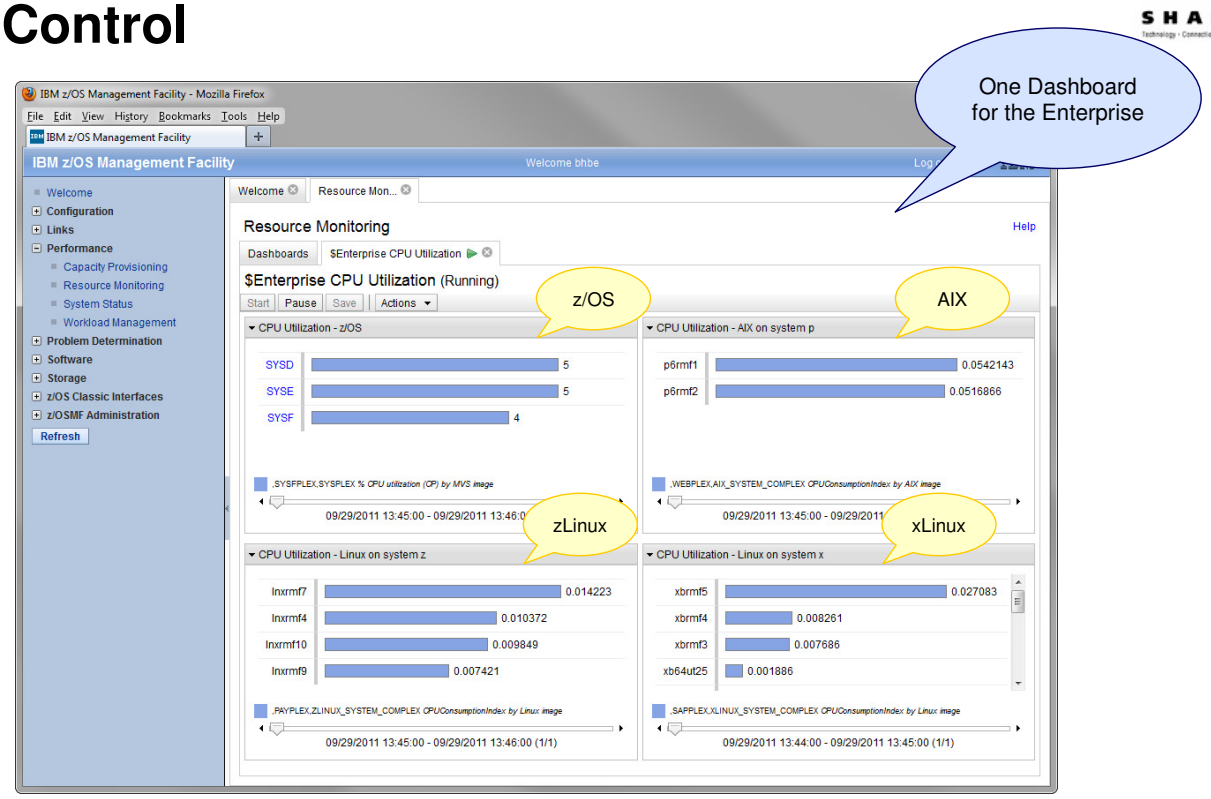

- One workstation can serve as singlepoint of control when RMF XP is used together with the z/OSMF Resource Monitoring facility
- The user can and define a Metric Group for each platform within the same Dashboards
- Hence, the selected key metrics can be monitored for all platforms at a glance

RMF XP – The Metric Promotion Concept

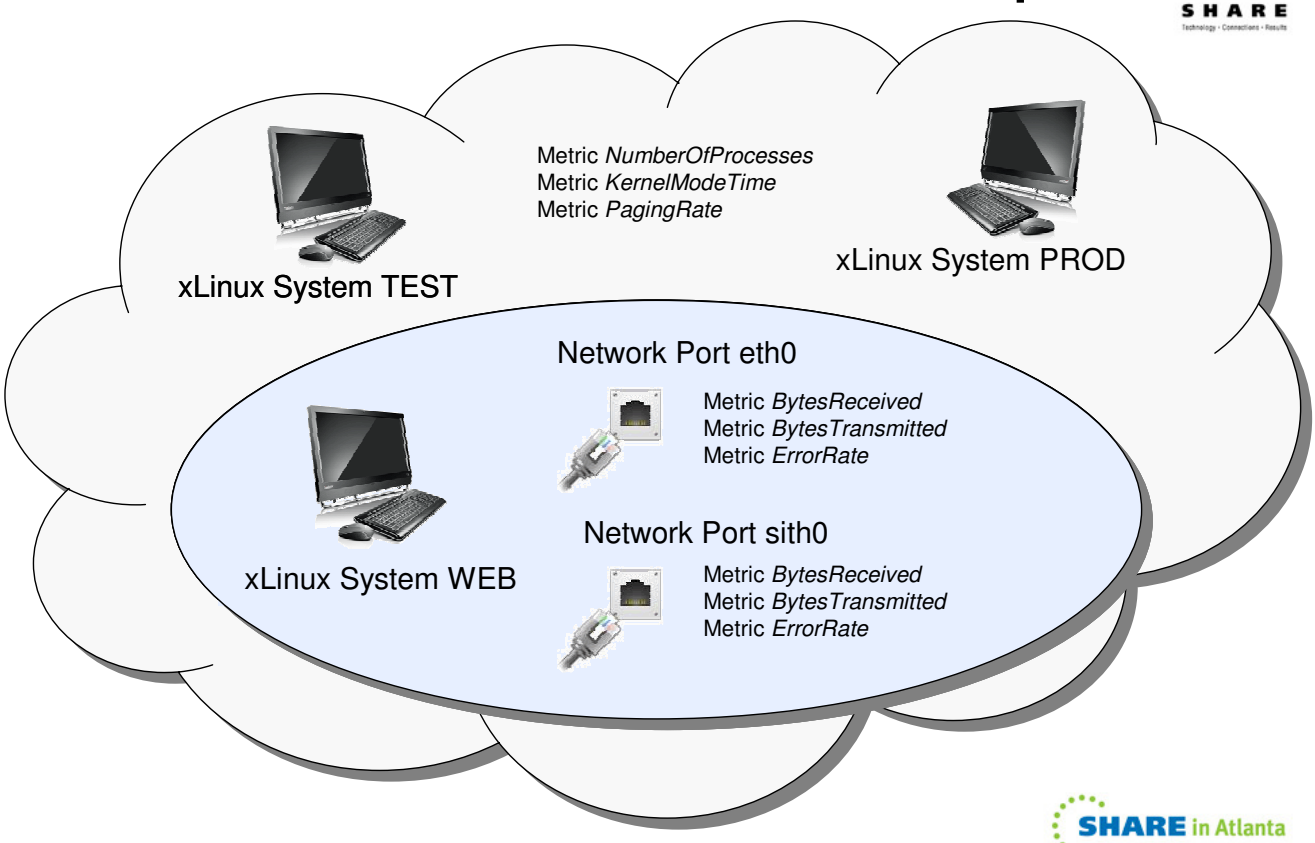

- The performance of a computing system can be evaluated by different kinds of metrics:
  - Metrics which are applicable to the entire system (e.g. KernelModeTime) or a unique resource within a system (e.g. PagingRate for the Memory resource)
  - Metrics which are applicable to resources where multiple instances can exist within a system (e.g. BytesReceived for Network Port resources)
- Hence, exactly one measurement value can arrive for the first kind of metrics while multiple measurement values can arrive for the other ones
- RMF XP can collect all metric values by means of the CIM API in terms of name-value pairs
- In the above example, exactly 9 measurement values can be retrieved for each system

RMF XP – The Metric Promotion Concept ...

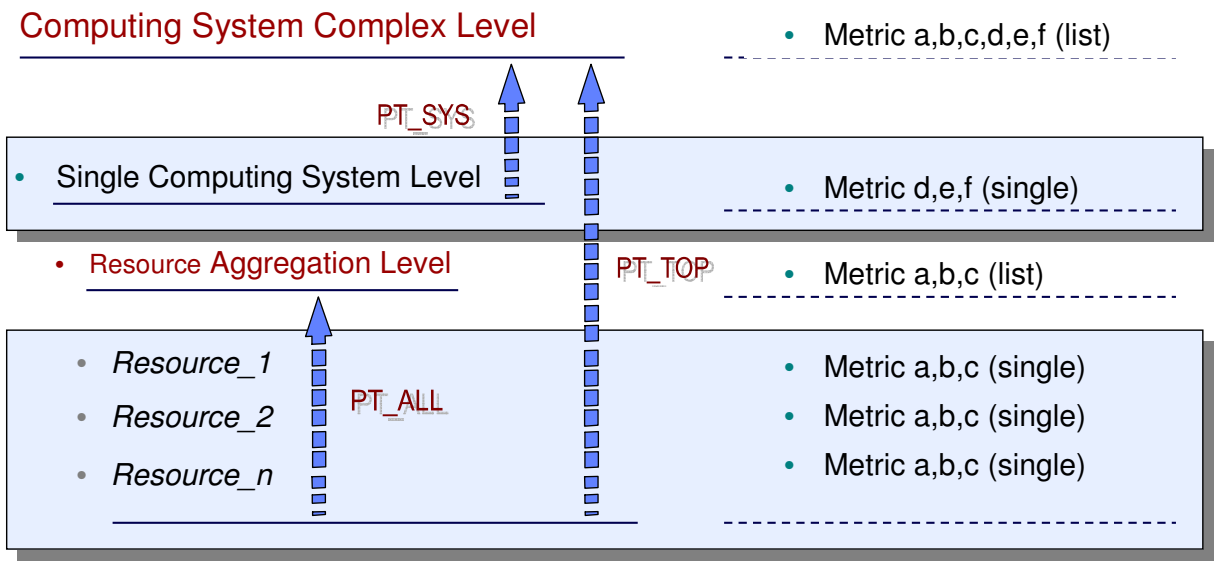

- PT_SYS: Promotion Type System
- PT_ALL: Promotion Type All
- PT_TOP: Promotion Type Top

- The starting point for the promotion are the basic metrics supported by the CIM API
- RMF XP introduces new (virtual) resources and assigns the existing metrics to to those resources in terms of *list-valued* metrics
- Following a well defined scheme, the new resources are created with on two levels:
  - Resource Aggregation Level (multiple new resources)
  - System Complex Level (exactly one new resource)
- We distinguish the following promotion types:
  - PT_ALL
    - A new resource of the type ALL is created (e.g. ALL_NETWORK_PORTS)
    - The original metrics of type single are assigned to the ALL resource in terms of list valued metrics (e.g. BytesReceived by Network Port)
  - PT_TOP
    - Similar than PT_ALL, but the promotion is performed beyond the boundaries of a single system to the system complex level
    - In order to identify a resource within a system complex, the instance name of a resource needs to be prefixed with it's system name
  - PT_SYS
    - A metric which exists only once within a system is promoted to the system complex level
    - The original metrics of type single are assigned to the system complex in terms of list valued metrics (e.g. PagingRate by System)

Metric Promotion – “Making Of“

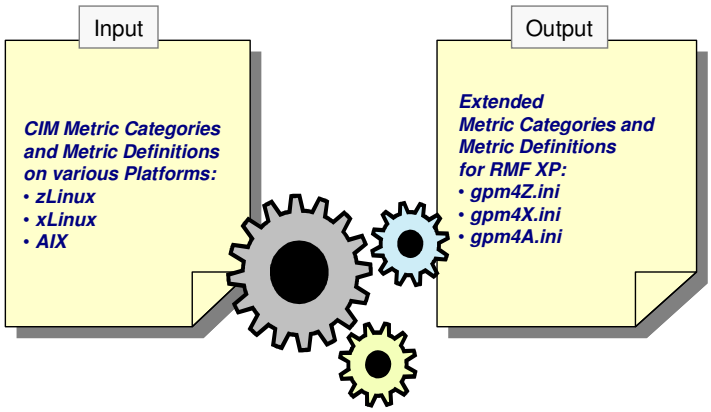

- The remote CIM API allows to retrieve the supported metric categories and metric definitions for the AIX and Linux operating system
- The transformation to the extended metric categories and metric definitions is performed automatically and the results are stored to platform specific initialization file
- Finally the initialization files are processed during the GPM4CIM startup and the contained definitions are visible for the RMF XP users

“Making Of“ – The Starting Point

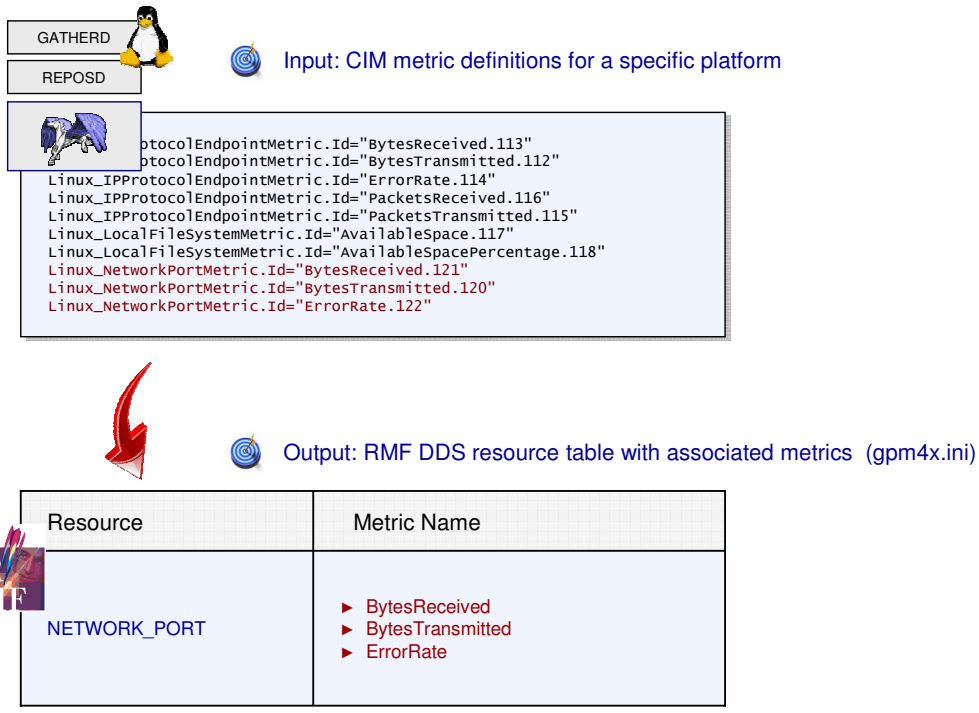

- The first step of the transformation is the pure one-to-one mapping of the CIM definitions
- No additional resources or metrics are created at this time

“Making Of“ – Broaden the Metric Scope (1)

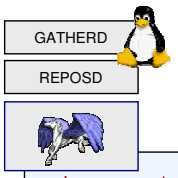

Input: CIM metric definitions for a specific platform

```
Linux_NetworkPortMetric.Id="BytesReceived.121"
Linux_NetworkPortMetric.Id="BytesTransmitted.120"
Linux_NetworkPortMetric.Id="ErrorRate.122"
```

Output: RMF DDS resource table with associated metrics (gpm4x.ini)

Resource	Metric Name	
LINUX_SYSTEM_COMPLEX	► BytesReceived by network port ► BytesTransmitted by network port ► ErrorRate by network port	☑ promotion type: TOP
ALL_NETWORK_PORTS	► BytesReceived by network port ► BytesTransmitted by network port ► ErrorRate by network port	☑ promotion type: ALL
NETWORK_PORT	► BytesReceived ► BytesTransmitted ► ErrorRate	☑ metric base

- In the second step, the promotion types PT_ALL und PT_TOP are performed:
  - PT_ALL metrics of type *single* will be available on parent level in terms of *list valued* metrics
  - PT_TOP metrics of type *single* will be available on system complex level in terms of *list valued* metrics

“Making Of“ – Broaden the Metric Scope (2)

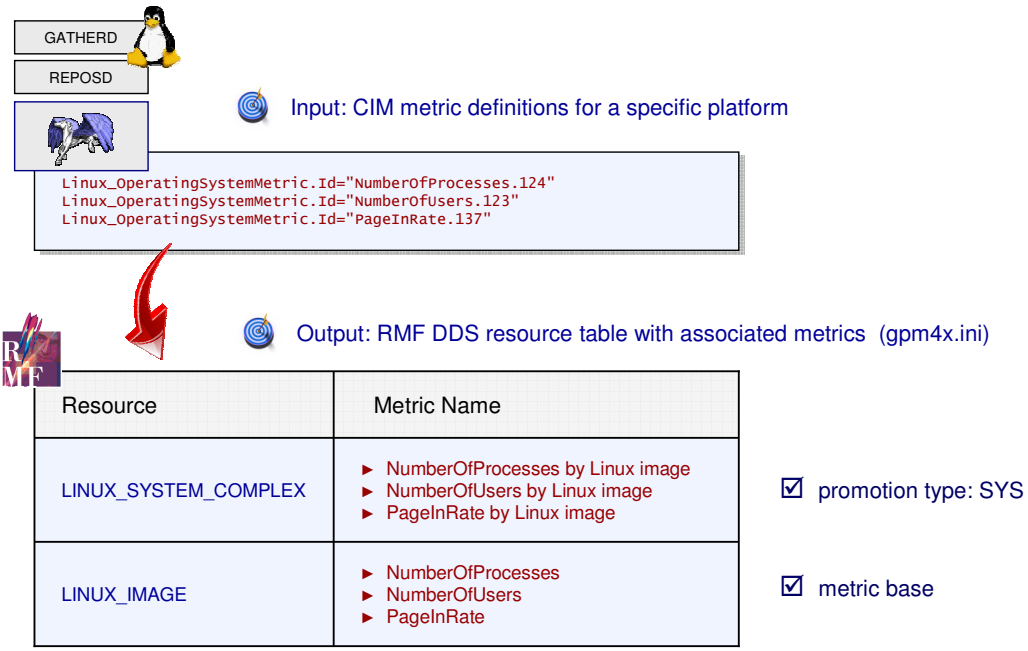

- For CIM metrics with just one instance within a system, the promotion type PT_SYS is applied: metrics of type *single* will be available on system complex level in terms of *list valued* metrics

“Making Of“ – Result (xLinux example)

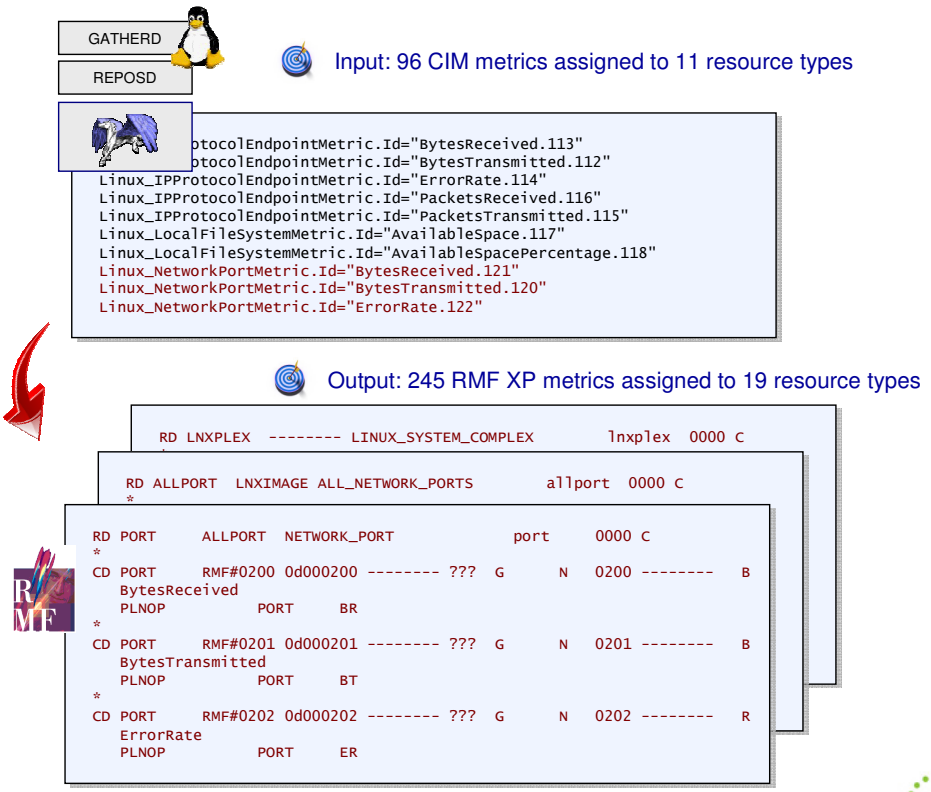

- The promotion concept creates additional resource types with additional metric definitions
- As an example for the xLinux platform, the promotion has created 19 RMF XP resource types (from 96 CIM resource types) and 245 metric definitions (from 96 CIM metric definitions)
- Alltogether, the huge benefit of the promotion is the comprehensive monitoring beyond the boundaries of one single system

RMF XP – Summary

- Seamless performance monitoring solution for z/OS and distributed platforms
- Promotion concept allows monitoring beyond the boundaries of a single system
- z/OS as monitoring platform for distributed environments
- Easy to setup, almost no customization needed
- Two graphical frontends
  - Instant access via web browser
  - z/OSMF with advanced capabilities
- zIIP exploitation helps to reduce costs
- Available with z/OS V1R13 RMF and z/OS V1R12 RMF (APAR OA36030)

RMF XP –

HOT TOPICS z/OS Hot Topics
Learn about key z/OS topics and features

SHARE
Technology · Connections · Results

Now performing!

RMF XP

BY KARIN GENTHER, HARALD BENDER, AND JÜRGEN BAUMANN

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

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

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

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

Performance data at a glance!

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

No rehearsal

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

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

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

SHARE in Atlanta 2012

Information & Tools

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

- Product information, newsletters, presentations, ...
- Downloads
  - Spreadsheet Reporter
  - RMF PM Java Edition
  - Postprocessor XML Toolkit

RMF email address: rmf@de.ibm.com

Users Guide:
New RMF XP
Chapter

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/#erb99l13.scr