

SHARE
Technology • Connections • Results


RMF – The Latest and Greatest

Brad Snyder (Bradley.Snyder@us.ibm.com)
IBM Corporation

Monday, February 28, 2011



Trademarks



SHARE
Technology • Connections • Results

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.

SHARE
in Anaheim
2011

RMF Product Overview



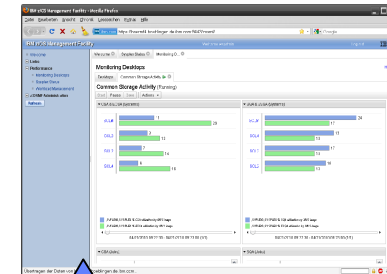
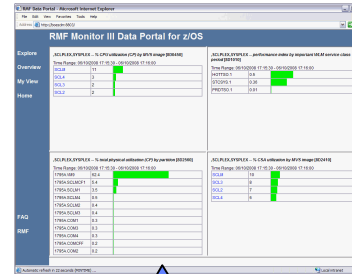
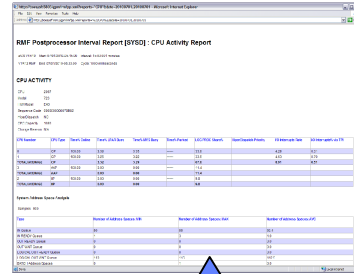

SHARE
Technology • Connections • Results

RMF Spreadsheet Reporter

Postprocessor Data Portal

Monitor III Data Portal
& RMF PM

z/OSMF Resource
Monitoring



RMF Distributed Dataserver (DDS)



RMF Postprocessor
Historical Reporting,
Analysis and Planning

RMF Monitor II and III
Real-Time Reporting,
Problem Determination

RMF Sysplex Data Server and APIs



RMF Data Gatherer

RMF Monitor I RMF Monitor II background RMF Monitor III



RMF Enhancements

- Enhanced Group Capacity Reporting
 - ▶ Examine WLM Soft-Capping Activities
 - ▶ Monitor MSU's of Capacity Groups
- Monitor II OPT Report
 - ▶ Identify active IEAOPTxx Member
 - ▶ Review Parameters at a Glance
- 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
- HTTP API to access Historical data
 - ▶ Enable RMF Distributed Data Server to request historical RMF Postprocessor data via HTTP
- Distributed Data Server Report Selection
 - ▶ Optimize GPMERVE Performance
- In-Ready Work Unit Queue Distribution
 - ▶ Identify latent Demand
 - ▶ Dispatchable Unit Granularity
- z/OSMF Resource Monitoring
 - ▶ Plugin of z/OSMF Management Facility
 - ▶ Cross-sysplex performance monitoring from a single point of control



SHARE

z/OS V1.11

z/OS V1.11

z/OS V1.11

z/OS V1.12


z/OS V1.11

z/OS V1.12

z/OS V1.12

SHARE
in Anaheim
2011

Group Capacity Reporting

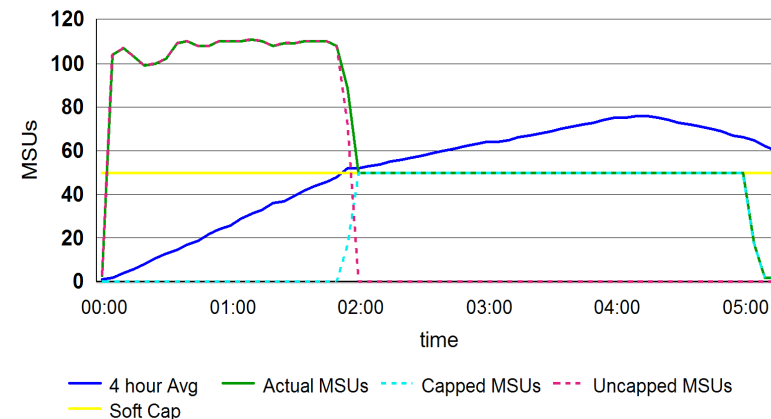


SHARE
Technology • Connections • Results

- Defined Capacity
 - Sets an upper limit for the capacity that a partition may consume
 - Specified in “million service units per hour” - MSU
 - Controlled by WLM , Enforced by PR/SM
- 4 hour rolling average: 48 buckets a 5 min
 - WLM computes MSU’s consumed by the LPAR
 - All intervals before IPL are initialized to 1 MSU/h
- Capping is enforced when 4 hour rolling average exceeds the specified limit
 - Cap removed when the 4h avg drops below
 - Usage spikes *above* the cap are possible
 - 4h rolling average may exceed cap as well!



Defined capacity limit enforcement




Group Capacity Reporting

How does WLM Softcapping work?

- ▶ WLM exploits the existing Capping mechanism in a PR/SM environment. (used by PR/SM hypervisor to limit the CPU access for a logical partition)
- ▶ When CAPPING=YES for an individual partition is active, the maximum capacity granted to this partition is derived from the weighting factors. The partition cannot “borrow” additional capacity from other partitions (even though contention is not in effect)
- ▶ The Capping mechanism depends on partition weight vs. defined capacity:


Weight vs. defined capacity	How capping is enforced
Weight equivalent = defined capacity limit	PR/SM hypervisor caps at partition weight
Weight equivalent < defined capacity limit	Noncontiguous cap pattern
Weight equivalent > defined capacity limit	Capping via “Phantom weight”

Softcapping Example: Weight Equivalent for LPAR1 = 50 MSUs



Contiguous Cap Pattern


- ▶ Weight Equivalent = MSU Limit (e.g. 50 vs. 50)



In order to cap LPAR1 down to 50 MSU's the LPAR hypervisor can apply the capping technique unchanged and contiguously

Non Contiguous Cap Pattern


- ▶ Weight Equivalent < MSU Limit (e.g. 50 vs. 60)



In case the other LPARs don't consume their weight equivalents, up to 10 additional MSU's can be granted to LPAR 1 by disrupt the capping temporarily


Phantom Weight

- ▶ Weight Equivalent > MSU Limit (e.g. 50 vs. 40)



A phantom LPAR with 10 MSU's is created. This capacity is delivered to LPAR 2 & LPAR 3 accordingly to their weights

Group Capacity Reporting...

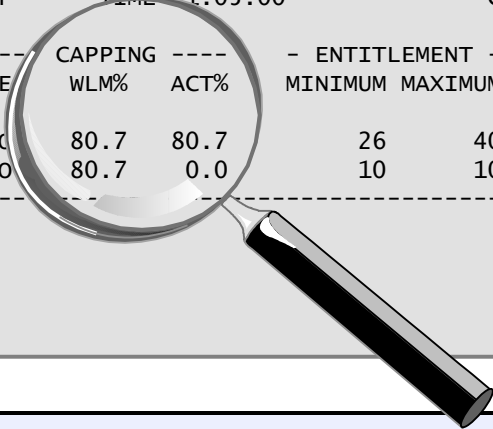


SHARE
Technology • Connections • Results

GROUP CAPACITY REPORT

z/OS V1R11 SYSTEM ID TRX1 DATE 03/19/2009 INTERVAL 05.00.000
RPT VERSION V1R11 RMF TIME 11.05.00 CYCLE 1.000 SECONDS

GROUP-CAPACITY NAME	LIMIT	PARTITION	SYSTEM	-- MSU -- DEF	ACT	WGT	---	CAPPING WLM%	----	ACT%	- ENTITLEMENT - MINIMUM	MAXIMUM
RMFGRP	40	TRX1	TRX1	0	37	400	NO	80.7	80.7		26	40
		TRX2	TRX2	10	4	200	NO	80.7	0.0		10	10
-----				TOTAL		41	600					



Field Heading	Meaning
CAPPING WLM%	Percentage of time when WLM considers to cap the partition. This applies when the partition is a member of a capacity group which is currently in capping state. That means, that the unused capacity credit for this group is exhausted and one or more members of the group will be actually capped in case they still call for additional capacity above their limits
CAPPING ACT%	Percentage of time when capping actually limited the usage of processor resources for this partition

Group Capacity Reporting...




PARTITION DATA REPORT

PAGE 2

z/OS V1R11 SYSTEM ID TRX1 DATE 03/19/2009 INTERVAL 05.00.000
 RPT VERSION V1R11 RMF TIME 11.05.00 CYCLE 1.000 SECONDS

MVS PARTITION NAME	TRX1	NUMBER OF PHYSICAL PROCESSORS	26	GROUP NAME	PMFGRP
IMAGE CAPACITY	40	CP	4	LIMIT	40*
NUMBER OF CONFIGURED PARTITIONS	58	AAP	1	AVAILABLE	12
WAIT COMPLETION	NO	IFL	18		
DISPATCH INTERVAL	DYNAMIC	ICF	2		
		IIP	1		




----- PARTITION DATA ----- -- LOGICAL PARTITION PROCESSOR DATA -- -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --

NAME	S	WGT	---MSU---			PROCESSOR- NUM	TYPE	---DISPATCH TIME DATA---		LOGICAL PROCESSORS		--- PHYSICAL PROCESSORS ---			
			DEF	ACT	DEF			WLM%	EFFECTIVE	TOTAL	EFFECTIVE	TOTAL	LPAR MGMT	EFFECTIVE	TOTAL
TRX1	A	400	50	4	NO	0.0	3.0	CP	00.00.10.184	00.00.10.501	1.13	1.17	0.03	0.85	0.88
H05LP45	A	10	0	1	NO	0.0	2	CP	00.00.01.866	00.00.02.218	0.31	0.37	0.03	0.16	0.18
TRX2CFA	A	100	0	1	YES	0.0	1	CP	00.00.04.139	00.00.04.210	1.38	1.40	0.01	0.34	0.35
H05LP59	A	100	0	1	NO	0.0	3	CP	00.00.02.594	00.00.03.274	0.29	0.36	0.06	0.22	0.27
H05LP60	A	10	0	8	NO	0.0	2	CP	00.00.21.928	00.00.23.760	3.65	3.96	0.15	1.83	1.98
TRX2	A	200	60	3	NO	0.0	3.0	CP	00.00.09.073	00.00.09.495	1.01	1.06	0.04	0.76	0.79
PHYSICAL										00.00.22.805			1.90		1.90
TOTAL									00.00.49.787	00.01.16.267			2.21	4.15	6.36

Field Heading	Meaning
LIMIT	Capacity limit (in MSU's) defined for the partition's capacity group. An '*' following the limit value indicates that this partition started to be a member of this capacity group less than four hours ago. Therefore, the capacity of this partition is still managed according to it's 4 hour rolling MSU average limitations and not yet according to the group capacity limit
AVAILABLE	Long-term average of CPU service units which would be allowed by the limit of the capacity group but are not used by its members. If the value is negative, WLM will start to cap one more members of this group in case they still call for additional capacity above their limits

Group Capacity Reporting...



SHARE
Technology • Connections • Results


The Monitor III CPC Capacity report is enhanced:

- ▶ New metric: *4h unused group capacity average*. This is the average available capacity for the group during the last 4 hours. If the value is negative the members have consumed capacity beyond the defined group limit. Therefore WLM sets the group into the capping state.
- ▶ New metric: *Remaining time until group capping*. This is the projected time until the usage of processor resources for one or more members of the group might be limited.
- ▶ Both new metrics are available as hidden fields in the ISPF version of the CPC report
- ▶ The threshold applied when calculating the *remaining time until capping* is changed for partitions being member of a capacity group: The MSU limit for the projection is now derived from the image capacity or group capacity definitions.

The following new metrics will be provided by the RMF Distributed Dataserver:

Resource	New metric
Sysplex	<ul style="list-style-type: none">▶ available capacity (MSU/h) for group by partition▶ remaining time until group capping in seconds by partition
LPAR	<ul style="list-style-type: none">▶ available capacity (MSU/h) for group▶ remaining time until group capping in seconds

Monitor II OPT Report



SHARE
Technology • Connections • Results



The screenshot shows the IBM WLM Work Queue Viewer web page in Microsoft Internet Explorer. The browser address bar shows <http://www-03.ibm.com/servers/eserver/zseries/zos/wlm/tools/wlmque.html>. The page content includes a navigation menu, a search bar, and a main section titled "WLM Work Queue and OPT Parameter Viewing Tools". This section describes the WLMQUE and WLMOPT tools, their installation instructions, and related links. A red arrow points from the "Download" section of the web page to the terminal window on the right.


```
RMF Monitor II Library List and OPT Settings Selection Menu
Selection =====

Enter selection number or command on selection line.

1 Link list      LNKLSTxx - Link Library list      (LLI)
2 LPA list       LPALSTxx - LPA Library List       (LLI LPA)
3 APF list       IEAAPFxx - Authorized Program List (LLI APF)

4 OPT            IEAOPTxx - OPT Settings          (OPT)
```

Monitor II OPT Report



SHARE
Technology • Connections • Results

```

RMF - OPT Settings                               Line 1 of 29
Command ==>                                     Scroll ==> PAGE


                                CPU=  1/  1 UIC= 65K PR=  0          System= TRX2 Total

OPT: 00                                         Time: N/A
-- Parameter -- -- Default -- -- value -- Unit ----- Description -----
ABNORMALTERM          Yes          Yes Y/N  Abnormal terminations in routing
BLWLINTHD              20           20 sec  Time blocked work waits for help
BLWLTRPCT              5            5 0/00  CPU cap. to promote blocked work
CCCAWMT               12000        12000 usec  Alternate wait management time
ZAAPAWMT              12000        12000 usec  AWM time value for ZAAPS
ZIIPAWMT              12000        12000 usec  AWM time value for ZIIPS
CNTCLIST              No           No Y/N   Clist commands count individually
CPENABLE              10,30|0,0    10,30 %   Threshold for TPI (low,high)
DVIO                  Yes          Yes Y/N   Directed VIO is active
ERV                   500         500/CB  SU      Enqueue residency CPU Service/DP
HIPERDISPATCH        No           No/No Y/N  Hiperdispatch is desired/active
IFAHONORPRIORITY      Yes          Yes Y/N   Allows CPs to help ZAAPS
IIPHONORPRIORITY      Yes          Yes Y/N   Allows CPs to help ZIIPS
INITIMP               0            0/FE #   INITIMP value/DP for initiators
IRA405I               70,50,50     70,50,50 %   Fixed storage of <16M,16M-2G,tot
MAXPROMOTETIME        6            6 *10s  Holder allowed to run promoted
MCCAFCTH              1933,3867    400,800 #   Threshold for storage (low,ok)
MCCFXEPR              92           92 %     Fixed storage threshold < 16 MB
MCCFXTPR              80           80 %     Fixed online storage threshold
PROJECTCPU            No           Yes Y/N   CPU projection for ZAAPS, ZIIPS
    
```

Current Member
& Activation Time

Explanation
of
Parameters


Postprocessor XML Formatted Reports



SHARE
Technology • Connections • Results


Rationale:

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



RMF Postprocessor reports can now be generated in XML Format

Postprocessor XML Formatted Reports



SHARE
Technology • Connections • Results

- 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


- With z/OS 1.12, the first sysplex report is available in XML format

- ⇒ Workload Activity report

z/OS 1.12

- Overview reports are also available in XML format
- Summary and Exception reports are not available in XML format
- Interval reports based on data collected during a Monitor II background session are not available in XML format

Postprocessor XML Formatted Reports




SHARE
Technology • Connections • Results

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

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

z/OS 1.12

Spreadsheet Reporter – XML Support



SHARE
Technology • Connections • Results

The image shows the IBM RMF Spreadsheet Reporter application interface. The main window has a menu bar (File, Define, View, Settings, Cr) and a toolbar with icons for file operations. The 'Resources' pane on the left shows a tree view of systems, including 'Remote' and 'Local' folders with sub-folders like 'SMF Dump Data', 'Report Listings', 'Overview Records', 'Working Sets', and 'Spreadsheets'. The 'Options' dialog box is open, with the 'Reports' tab selected. The 'Reports' tab contains a list of 'RMF Postprocessor Report Types' with checkboxes for each: CPU Activity, Crypto Hardware Activity, Enterprise Disk Systems, FICON Director Activity, OMVS Kernel Activity, and Partition Data. A blue speech bubble points to this list with the text 'List of currently available reports in XML format'. A red arrow points from the 'Reports' tab in the Options dialog to the 'Reports' tab in the main application window.

Options

General Reports

General Processing Options

- Create Overview Records
- Delete Postprocessor Datasets after Download
- Ignore specified Duration Period
- Ignore specified Interval Time
- Save Password with System Profile
- Scratch Overview Records after Conversion
- Scratch Report Listings after Conversion
- Scratch extracted QVW Files after Conversion
- Scratch extracted RPT Files after Conversion
- Sort SMF Datasets
- Use XML Report Format

Options


General Reports

RMF Postprocessor Report Types

- CPU Activity
- Crypto Hardware Activity
- Enterprise Disk Systems
- FICON Director Activity
- OMVS Kernel Activity
- Partition Data

List of currently available reports in XML format

Spreadsheet Reporter – XML Support



SHARE

IBM RMF Spreadsheet Reporter Java™ Technology Edition (SYSF)

File Define View Settings Create Messages Help

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

File Edit View Favorites Tools Help

RMF Postprocessor Duration Report [SYSD] : CPU Activity Report

z/OS V1R9 Start: 07/31/2009-08.15.00 Interval: 000:44:59 hours
 V1R11 RMF End: 07/31/2009-09.00.00 Cycle: 1000 milliseconds

CPU ACTIVITY


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

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


Done My Computer

MVS1.D127.T112039.lis
 SCLM.D215.T133732.lis
 SCLM.D215.T133808.lis
 SYSF.D212.T114702.xml
 SYSF.D215.T113455.xml
 SYSF.D215.T113918.xml
 Sample.lis

- Start
- View**
- New
- Rename
- Delete
- Properties



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

RMF Postprocessor Interval Report [SYSD] : CPU Activity Report

z/OS V1R10 Start: 07/02/2010-11.45.00 Interval: 14:59:999 minutes
V1R12 RMF End: 07/02/2010-12.00.00 Cycle: 1000 milliseconds


CPU ACTIVITY

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

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



HTTP API to access Historical data ...



SHARE

Technology • Connections • Results

GPDDSP

RMF Postprocessor JOB

```
//RMFPP EXEC PGM=ERBRMFPP
//MFPMSGDS DD SYSOUT=*

//XPRPTS DD SYSOUT=*

//XPXSRPTS DD SYSOUT=*
//XPOVWRPT DD SYSOUT=*

//SYSOUT DD SYSOUT=*
//SYSIN DD *
```

RMF Distributed Data Server

```
//STEP1 EXEC PGM=GPDDSRV
...
//GPMPJCL DD
DSN=SYS1.SERBPWSV(GPMPJCL)
...
```


Client Application



HTTP Request



XML Response



Submit PP JOB



PP JOB Output



DATE (07022010) 07022010
REPORTS (C...

SMF Records



JCL Template: SYS1.SERBPWSV(GPMPJCL)

```
//GPDDSP JOB MSGCLASS=H,CLASS=A,
//          MSGLEVEL=(1,1)
/*JOBPARM SYSAFF=*

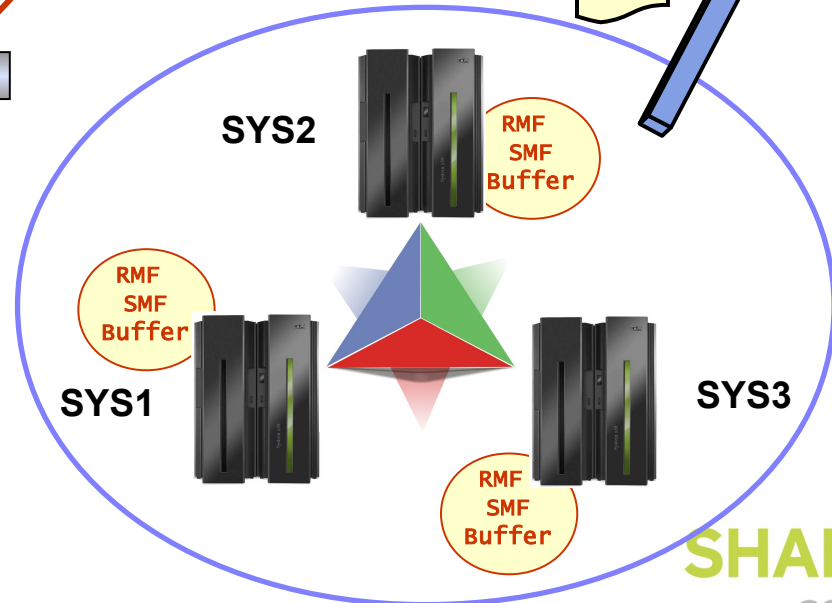
//RMFPP EXEC PGM=ERBRMFPP
//MFPMSGDS DD SYSOUT=*

//XPRPTS DD SYSOUT=*

//XPXSRPTS DD SYSOUT=*
//XPOVWRPT DD SYSOUT=*


//SYSOUT DD SYSOUT=*

//SYSIN DD *
```



SHARE
in Anaheim
2011

Distributed Data Server Report Selection



S H A R E



- The DDS supports the following kinds of metric types:

⇒ Full reports

⇒ Single valued metrics

⇒ List valued metrics

RMF Monitor III Data Portal for z/OS

Full RMF Reports:

CHANNEL	CPC	DELAY	DEV	DEVR	DSND	ENCLAVE	IOQ	LOCKSP	LOCKSU	OPD
PROC	PROCU	STOR	STORC	STORCR	STORF	STORM	STORS	SYSINFO	ZFSACT	ZFSSUM

Available metrics for: ,SYSF.MVS_IMAGE

Metric description	Help	Id
% delay	Explanation	8D0160
% idle	Explanation	8D03E0
% unknown	Explanation	8D0470
% using	Explanation	8D04A0
% workflow	Explanation	8D0550
# active users	Explanation	8D0620
# users	Explanation	8D0D50
execution velocity	Explanation	8D0EF0
transaction ended rate	Explanation	8D1200
by enclave		
% delay by enclave	Explanation	8D28A0
% idle by enclave	Explanation	8D2A90
% using by enclave	Explanation	8D2B10
by job		
% delay by job	Explanation	8D0200
% idle by job	Explanation	8D03F0
% unknown by job	Explanation	8D0480
% using by job	Explanation	8D04E0
% workflow by job	Explanation	8D0560

- The one and only metric source for the DDS are Monitor III report data tables
- By default, the DDS exploits a defined set of Monitor III sysplex-wide reports and single system reports


RMF Monitor III Data Portal for z/OS

Full RMF Reports:

CACHDET	CACHSUM	CFACT	CFOVER	CFSYS	SPACED	SPACEG	SYSSUM	XCFGROUP	XCFOWW	XCFPATH
---------	---------	-------	--------	-------	--------	--------	--------	----------	--------	---------

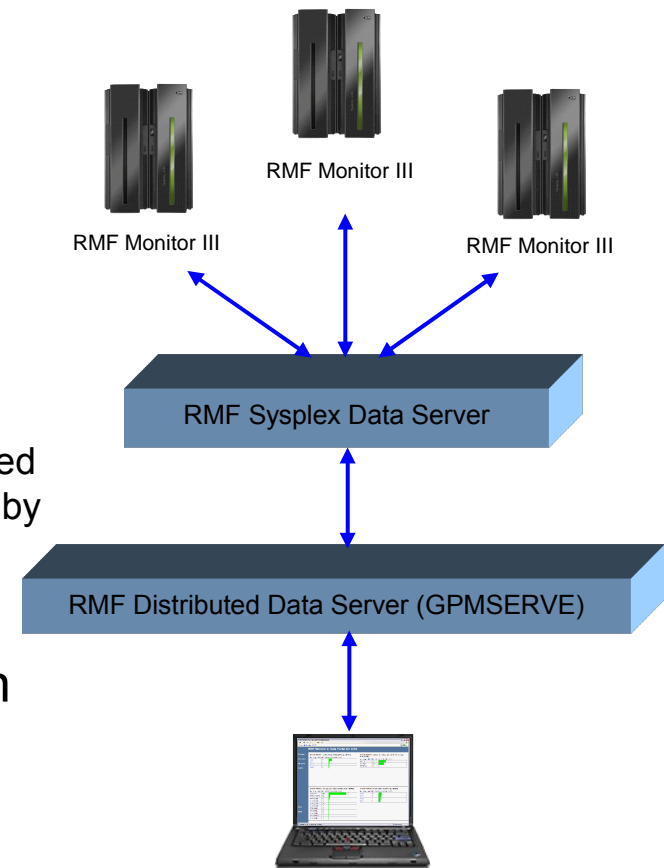
Available metrics for: ,SYSDPLEX,SYSPLEX

Distributed Data Server Report Selection




SHARE
Technology • Connections • Results

- The DDS retrieves Monitor III report data from all RMF instances in the Sysplex
- With z/OS V1.11 RMF it is possible to exclude specific Monitor III reports from being retrieved by the DDS
 - Users can exclude reports which are not required by any monitoring clients in their environment
 - In case of abnormal termination of reporter modules, operation of the DDS is affected. With this function, those reporters can be excluded temporarily. This allows continued operation of the DDS until a fix for the problem is provided by IBM
- A new DDS option is introduced for PARMLIB member GPMSRVxx. This option defines a list of Monitor III reports to be excluded from DDS processing



Distributed Data Server Report Selection




SHARE
Technology • Connections • Results

- New option EXCLUDE_REPORTS in DDS parmlib member
- The option is not mandatory
- If not specified, all Monitor III reports supported by the DDS are available
- If specified, the option contains a comma-separated list of report names which should be excluded from DDS processing
- Specify multiple options if the list of reports does not fit on one line
- In this example, the Cache Detail, Cache Summary, zFS Activity, and zFS Summary reports are excluded

```
/******  
/*  
/* NAME:          GPMSRV00  
/*  
/* DESCRIPTION:  PARMLIB MEMBER FOR THE RMF DISTRIBUTED  
/*              DATA SERVER HOST ADDRESS SPACE (GPMSERVE)  
/*  
/* COPYRIGHT:    LICENSED MATERIALS - PROPERTY OF IBM  
/*              "RESTRICTED MATERIALS OF IBM"  
/*              5694-A01  
/*              (C) COPYRIGHT IBM CORP. 1998, 2009  
/*  
/*              STATUS=HRM7760  
/******  
  
...  
  
/******  
/*  
/* List of reports to be excluded  
/*  
/******  
EXCLUDE_REPORTS (CACHDET,CACHSUM)  
EXCLUDE_REPORTS (ZFSACT,ZFSSUM)
```

Distributed Data Server Report Selection



SHARE
Technology • Connections • Results


- When a client requests a report from the DDS that has been deactivated by EXCLUDE_REPORTS the following message is generated:
GPM0737I The requested report <reportname> is deactivated
- When a client requests a metric from the DDS that is based on a deactivated report the following message is generated:
GPM0738I The requested metric is based on the deactivated report <reportname>

Never deactivate the report types needed by other components !!

Caution:

- ⇒ z/OS Capacity Provisioning requires RMF DDS and CIM server to obtain performance data from the CPC, SYSINFO and SYSSUM report.
- ⇒ OMEGAMON XE on z/OS can be optionally configured to require RMF DDS to obtain performance data from the following reports: CFACT, CFOVER, CFSYS, XCFOVW, XCFGROUP, XCFPATH, XCFSYS, LOCKSP, LOCKSU

Work Unit Queue Distribution




SHARE
Technology • Connections • Results



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

Work Unit Queue Distribution Example



SHARE
Technology • Connections • Results



0000





0000




0000



0000




Work Unit Queue Distribution: Spreadsheet Reporter




SHARE
Technology • Connections • Results


CPU Contention for System CB8D Date: 07/09/2010



CPU Work Unit Contention for CB8D Date: 07/09/2010

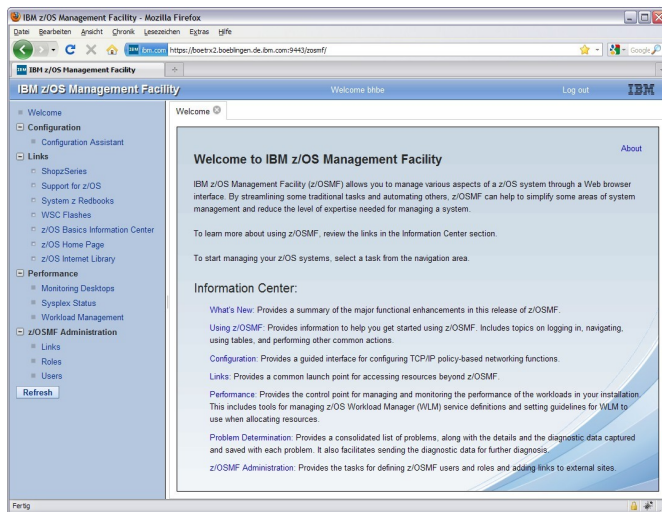


IBM z/OSMF Management Facility




SHARE
Technology • Connections • Results

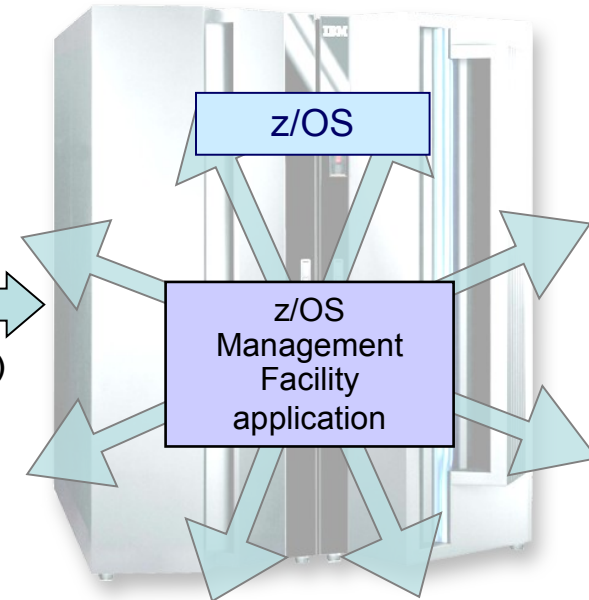
z/OS application, browser access



Browser




HTTP(s)



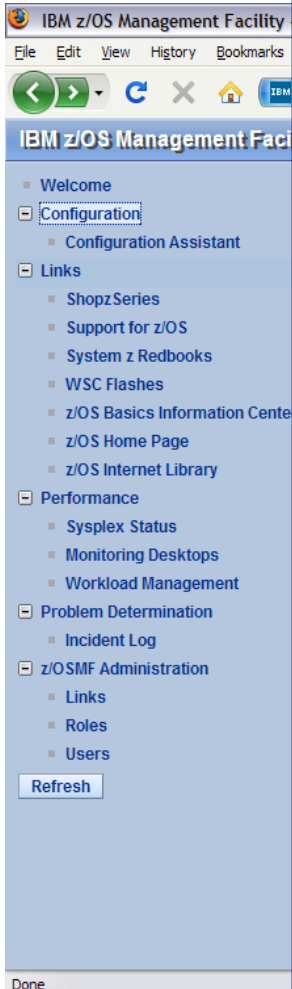
- *z/OS Management Facility is a Web 2.0 application on z/OS*
 - Manages z/OS from z/OS
 - Browser systems communicate with z/OSMF via secure connection, anywhere, anytime

IBM z/OSMF Management Facility ...



SHARE
Technology • Connections • Results


Applications / R12 plugins



- **Configuration category**
 - **Configuration Assistant for z/OS Communication Server** – Simplified configuration and setup of TCP/IP policy-based networking functions (with z/OS V1.11)
- **Links** Links to resources - provides common launch point for non-z/OSMF resources
- **Performance category (with z/OS V1.12)**
 - **Sysplex status** – single view of sysplex performance
 - **Monitor desktops** – dynamic real time metrics for system performance, customizable for correlations
 - **Workload Manager Policy Editor** -- Simplified creation, editing, installation, and activation of WLM service definitions and policies
- **Problem Determination category**
 - **Incident Log** -- Simplified capture, packaging, and sending of SVC dump diagnostic data. (Available with z/OS V1.10)
- **z/OSMF Administration** Authorization services, add users, define roles, add links.

IBM z/OSMF Management Facility ...

Resource Monitoring – Sysplex Status Task



SHARE
Technology • Connections • Results

The screenshot shows the IBM z/OSMF Management Facility interface in a Mozilla Firefox browser. The page title is "Sysplex Status" and it includes a "Refresh" button in the left sidebar. A table lists resources with their connectivity and performance index status. A "Refresh" button is visible at the bottom of the table area, and a checkbox for "Automatic refresh" is checked.

Automatic refresh


Enterprise-wide health check

Resource	Connectivity	Performance Index Status	Related Service Definition	Active WLM Policy
LOCALPLEX	Connected	PI <= 1 for all periods	RTDST3	RTDST
SYSDPLEX	Connected	PI > 1 for important periods	system	POLICY01
SCLMPLEX	Connected	PI <= 1 for all periods	Default	STANDARD
TRXPLEX	Connected	PI > 1 for important periods	Ralfpol	BASEPOL
SHARPLEX	Connected	PI <= 1 for all periods	SHARPLEX	SHAREPOL

Total: 5
Refresh Last refresh: 22.02.2010 17:57:24 local time (22.02.2010 16:57:24 GMT)
 Automatic refresh

IBM z/OSMF Management Facility ...

Resource Monitoring – Monitoring Desktops Task



SHARE
Technology • Connections • Results

The screenshot displays the IBM z/OSMF Management Facility interface in a Mozilla Firefox browser window. The main content area is titled "Monitoring Desktops" and shows "CPU & Workload Activity (Running)". It is divided into two panels: "Execution Velocity" and "CPU Consumption".

Execution Velocity Panel:

Job Name	Blue Bar Value	Green Bar Value
CICSWEB.1	70	93
PAYROLL.1	50	32
BTCHCR1.1	50	0
STCHIGH.1	40	0
USSDEF.1	40	0
STCDEF.1	30	0.3
BTCHHIGH.1	20	0
DB2DEF.1	20	0
WEBDEF.1	20	0
TSODEF.3	20	0
BILLING.1	15	9.8


CPU Consumption Panel:

Job Name	Blue Bar Value	Green Bar Value
WLM2.QSERCICS [0066]	109.3	0.1
WLM2.QSERPAY [0068]	40	0
WLM2.QSERCICS [0067]	14.8	0
WLM2.QSERPAY [0065]	9.7	0
WLM2.BILL01 [0036]	7.4	7.4
WLM2.BILL00 [0035]	7.1	7.1
WLM2.DB2Q01 [0038]	0.8	0.8
WLM2.DEF01 [003A]	0.7	0.7
WLM2.RMFGAT [0030]	0.6	0.6
WLM2.WLM [000A]	0.2	0.2
WLM2.CONSOLE [0009]	0.1	0.1

Annotations on the screenshot:

- Metric Group:** A blue callout bubble pointing to the "CPU & Workload Activity (Running)" title.
- Metric:** A blue callout bubble pointing to a specific data point in the CPU Consumption panel.
- Interval Control:** A blue callout bubble pointing to the time range selector at the bottom of the panels.

Information and Tools



SHARE
Technology • Connections • Results

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/r11pdf/#rmf






SHARE
Technology • Connections • Results

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

Redbooks

ibm.com/redbooks



Appendix: Function Reference



SHARE
Technology • Connections • Results

Function	Availability
Enhanced Group Capacity Reporting	z/OSV1.11 RMF
Monitor II OPT Report	z/OSV1.11 RMF
Postprocessor XML Formatted Reports	z/OSV1.11 RMF
HTTP API to access historical data	z/OSV1.12 RMF
Distributed Data Server Report Selection	z/OSV1.11 RMF
In-Ready work Unit Queue Distribution	z/OSV1.12 RMF
IBM z/OSMF Management Facility: Resource Monitoring Plugin	z/OSV1.12 zOSMF