

z/OS  
Resource Measurement Facility

The Latest and Greatest:  
z13 GA2 and z/OS V2R2



# 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 more complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml):

\*BladeCenter®, CICS®, DataPower®, DB2®, e business (logo)®, ESCON, eServer, FICON®, IBM®, IBM (logo)®, IMS, MVS, OS/390®, POWER6®, POWER6+, POWER7®, Power Architecture®, PowerVM®, PureFlex, PureSystems, S/390®, ServerProven®, Sysplex Timer®, System p®, System p5, System x®, z Systems®, System z9®, System z10®, WebSphere®, X-Architecture®, z13™, z Systems™, z9®, z10, z/Architecture®, z/OS®, z/VM®, z/VSE®, zEnterprise®, zSeries®

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 Sync 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.

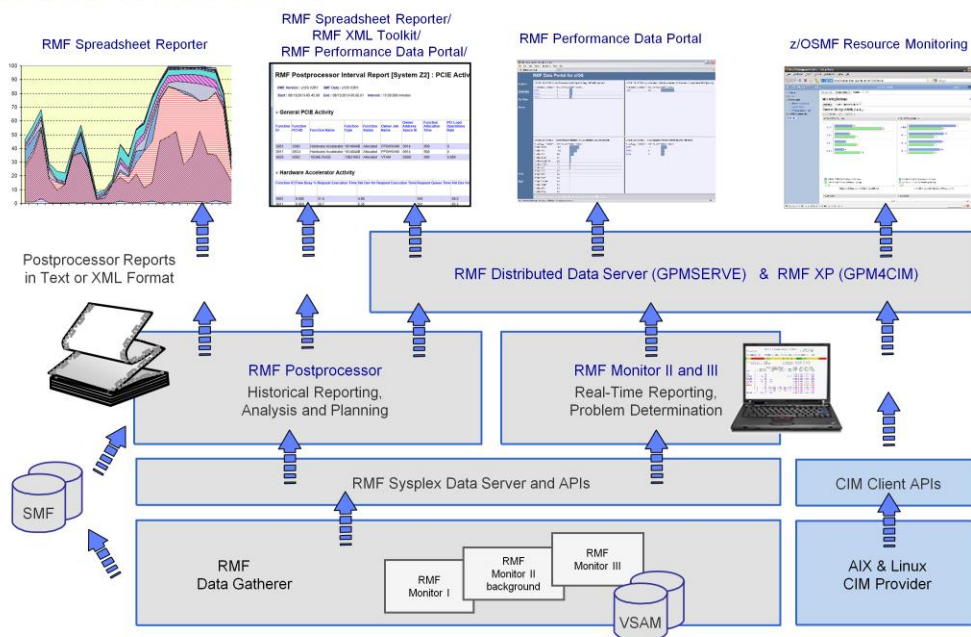
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available 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 Sync 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). Beginning 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.
- Since z/OS 1.12 there's another exploiter of the RMF DDS data: The z/OSMF Resource Monitoring plugin of the z/OS 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 at a Glance

- IBM z13 and z13s Support
  - RMF Statistics for Simultaneous Multithreading (SMT)
  - Extended ICSF Measurements for Crypto Express5S
  - Monitoring of new capping types



- SCM I/O Adapter Performance Reporting
  - Statistics for EADM Subchannel Activities
  - Performance and Thruput on SCM Card Level

- PCIE Activity Reporting

- z/OS V2.2 Monitor III PCIE Report
- Additional Measurements for PCIe attached RoCE and zEDC Devices on z13
- RMF support for SMC-D over ISM on z13 and z13s



- z/OS V2.2 Monitor III Job Resource Consumption Reporting
  - CPU, I/O and Storage Consumption Data at a Glance
  - Detailed Statistics for Job related GQSCAN Activities

- z/OS V2.2 zFS Reporting Enhancements

- Sysplex-wide Statistics for zFS Usage and Performance
- Improved zFS Data Gathering Performance
- Additional zFS Statistics for Shared File System environments



- DS8000 Support

- RMF Monitoring of SuperPAV capability

- Mobile Workload Reporting

- Statistics about total and mobile CPU consumption



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

With the availability of the IBM z13 servers, RMF provides first day support for a couple of notable hardware features.

- Comprehensive Statistics for Simultaneous Multithreading (SMT)
- Extended ICSF Measurements for Crypto Express5S
- Support for LPARs with up to 4TB Real Storage
- Monitoring of new capping types

Storage Class Memory – aka Flash Memory – is a new tier within the memory hierarchy of the zSeries family.

RMF provides detailed usage statistics for Storage Class Memory related operations by means of a new Monitor III SCM Activity report.

The z13 server family as well as z/OS 2.2 introduces various PCIe Activity reporting enhancements:

- With z/OS V2R2 RMF a complete new Monitor III PCIe Report allows online monitoring of PCIe related operations
- For z13, the existing Postprocessor PCIe Report has been extended with additional measurements
- With z13 GA2 and z13s a new SMC solution is available: SMC-Direct Memory Access (SMC-D) over Internal Shared Memory (ISM). RMF supports SMC-D with new PCIe statistics

The Monitor III Job Usage Report complements the Monitor III reporting suite with detailed statistics about address space resource consumption.

- The top resource consumers in terms of CPU, I/O and Storage can now be identified at a glance
- The report can serve as an excellent starting point for further drill-down and analysis
- Job related GQSCAN activities have been invisible in the past. With the new report detailed statistics with regard to GQSCAN usage can now be obtained

RMF z/OS 2.2 introduces new Monitor III Sysplex reports to monitor sysplex-wide z/OS Distributed File system (zFS) usage and performance.

SuperPAV is a new functionality of the IBM DS8000 server that allows to share Aliases among multiple control units.

RMF supports the monitoring of the DS8000 SuperPAV capability with new statistics in SMF records and report enhancements in the RMF Postprocessor I/O Queuing Activity report.

The new WLM support for mobile pricing allows to identify Mobile sourced transactions. RMF supports the WLM enhancements for mobile workloads with new statistics about total and mobile CPU consumption.

Last not Least the zEvent Mobile Application:

- This mobile App is currently under development and all details and capabilities – including the application name – are subject to change
- In a nutshell the app provides the following two main features to system administrators and performance analysts:
  - Receive push messages based on critical system events instantly
  - Access to the RMF Data Portal and z/OSMF Resource Monitoring anywhere and every time

## RMF z13 Support - Overview


- RMF is enhanced to monitor performance of logical zIIP cores and their threads configured in a Simultaneous Multithreading (SMT) environment on z13. **OA44101**
- RMF provides new measurements for PCIe-attached RoCE and zEDC devices on z13:
  - Supports new z13 10GbE shared RoCE (SR-IOV) express feature **OA44524**
- RMF support for Crypto Express5S (CEX5) card and new ICSF service measurements:
  - RSA Digital Signature Generate and Verify callable services **OA43493**
  - ECC Digital Signature Generate and Verify callable services
  - AES MAC Generate and Verify callable services
  - FPE Encipher, Decipher and Translate callable services
- RMF support for LPARs with up to 4 TB real storage. **OA44503**
- RMF support for z13 IBM Integrated Coupling Adapter (ICA SR) that provides PCIe based short-distance coupling links of type CS5 **OA44502**
- RMF support of SMC-D over ISM on z13 GA2 and z13s **OA49113**
- RMF support for new capping types **OA48688**



- With various new function APARs, RMF exploits the new functionality of the IBM z13:
  - OA44101: RMF support for the Simultaneous Multithreading (SMT) environment on z13 PTF available for z/OS 2.1
  - OA44524: RMF PCIE enhancements for RoCE and zEDC devices on z/OS 2.1.
  - OA43493: RMF support for the Crypto Express5S (CEX5) card and new ICSF service measurements.  
The support is available for z/OS 1.13 and z/OS 2.1.
  - OA44503: RMF support for z/OS 2.1 LPARs on z13 with up to 4TB real storage.
  - OA44502: RMF support for z13 IBM Integrated Coupling Adapter (ICA SR) that provides PCIe based short-distance coupling links of type CS5.  
PTFs available for z/OS 1.13 and z/OS 2.1.
  - OA49113: RMF support for SMC-Direct Memory Access (SMC-D) over Internal Shared Memory (ISM)  
PTF available for z/OS 2.2.
  - OA48688: RMF reporting enhancements for new capping types  
PTFs available for z/OS 2.1 and z/OS 2.2.
- RMF toleration support for IBM z13:
  - OA45890: z13 toleration for z/OS 1.10 and z/OS 1.11
  - OA45833: z13 toleration for z/OS 1.12 and z/OS 1.13

## z13 - Simultaneous Multithreading (SMT)

- “Simultaneous multithreading (SMT) permits multiple independent threads of execution to better utilize the resources provided by modern processor architectures.”\*
- With z13, SMT allows up to two instructions per core to run simultaneously to get better overall throughput
- SMT is designed to make better use of processors
- On z/OS, SMT is available for zIIP processing:
  - Two concurrent threads are available per core and can be turned on or off
  - Capacity (throughput) usually increases
  - Performance may in some cases be superior using single threading




**Two lanes process more traffic overall**

*Note: Speed limit signs for illustration only*

\*Wikipedia®

- Simultaneous multithreading (SMT) allows two active instruction streams (threads) per core, each dynamically sharing the core's execution resources. SMT will be available in IBM z13 for workloads running on the Integrated Facility for Linux (IFL) and the IBM z Integrated Information Processor (zIIP).
- SMT utilizes the core resources more efficiently: When a thread running on a core encounters a cache miss and can no longer make progress, the core switches to run a different thread that is ready to execute.
- Each thread runs slower than a non-SMT core, but the combined ‘threads’ throughput is higher. The overall throughput benefit depends on the workload.

## z13 –SMT Exploitation



- SMT Aware OS informs PR/SM that it intends to exploit SMT
  - PR/SM can dispatch any OS core to any physical core
  - OS controls the whole core – must follow rules
    - Maximize core throughput (Drive cores with high Thread Density [2] )
    - Maximize core availability (Meet workload goals using fewest cores )
- SMT is transparent to applications
- LOADxx and IEAOPTxx parmlib options to enable SMT on z/OS:
  - LOADxx: PROCVIEW CORE|CPU
  - IEAOPTxx: MT\_ZIIP\_MODE={1 | 2}

6

The Latest and Greatest

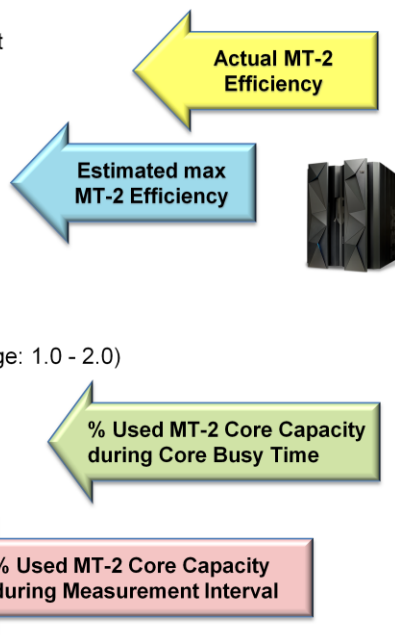
10/30/2016

© 2016 IBM Corporation

- The use of SMT mode can be enabled on an LPAR by LPAR basis via operating system (OS) parameters
- Once the OS switches to SMT mode, the only way back to single thread (ST) mode is via a disruptive action (re-activate the partition or re-IPL it).
- With the SMT enabled mode it is possible to dynamically switch between MT-1 (multi thread) and MT-2 mode for the processor types that support MT-2
- z/OS introduces new options for the LOADxx and IEAOPTxx parmlib members that are used to enable/disable SMT support and specify the MT mode of a processor class:
  - LOADxx parmlib option PROCVIEW CORE|CPU enables/disables SMT for the life of the IPL
    - PROCVIEW CORE on z13 enables SMT support
    - IPL required to switch between PROCVIEW CPU and CORE
  - New IEAOPTxx parameter to control the MT mode for zIIP processors
    - MT\_ZIIP\_MODE=1 specifies MT-1 mode for zIIPs (one active thread per online zIIP core)
    - MT\_ZIIP\_MODE=2 specifies MT-2 mode for zIIPs (two active threads per online zIIP core)
    - When PROCVIEW CPU is specified the processor class MT mode is always 1
    - SET OPT=xx operator command allows to switch dynamically between MT-1 and MT-2 mode
    - MT-2 mode requires HiperDispatch to be in effect
- z/OS SMT Terminology:
  - z/OS logical processor (CPU) → Thread
    - A thread implements (most of) the System z processor architecture
    - z/OS dispatches work units on threads
    - In MT mode two threads are mapped to a logical core
  - Processor core → Core
  - PR/SM dispatches logical core on a physical core
    - Thread density 1 (TD1) when only a single thread runs on a core
    - Thread density 2 (TD2) when both threads run on a core


## z13 - z/OS SMT Metrics

- **Capacity Factor (CF)**
  - How much work core actually completes for a given workload mix at current utilization - relative to single thread
  - MT-1 Capacity Factor is 1.0 (100%)
  - MT-2 Capacity Factor is workload dependent
- **Maximum Capacity Factor (mCF)**
  - How much work a core can complete for a given workload mix at most
- **Core Busy Time**
  - Time any thread on the core is executing instructions when core is dispatched to physical core
- **Average Thread Density**
  - Average number of executing threads during Core Busy Time (Range: 1.0 - 2.0)
- **Productivity**
  - Core Busy Time Utilization (percentage of used capacity) for a given workload mix
  - Productivity represents capacity in use (CF) relative to capacity total (mCF) during Core Busy Time.
- **Core Utilization**
  - Capacity in use relative to capacity total over some time interval
  - Calculated as Core Busy Time x Productivity



- z/OS SMT introduces several new metrics to describe how efficiently the core resources could be utilized and how efficiently they are actually utilized.






## z13 – RMF and SMT

- RMF enhanced with new metrics to monitor MT-2 efficiency and core utilization

- Re-interpret the meaning of existing RMF metrics:

- CPU metrics on core granularity (e.g. APPL%/EAPPL%)
- CPU metrics on thread granularity (e.g. MVS BUSY%)
- SMT updates in RMF Documentation



- Enhanced metrics descriptions
- General terminology:
  - „Processor“ → logical Core
  - „logical Processor“ → Thread

- MT-1 Equivalent Time

- z/OS CPU time consumed by work units (TCBs, SRBs) provided in terms of MT-1 equivalent time
- Time it would have taken to run same work in MT-1 mode
- Reflected in all RMF metrics reporting CPU consumption of workloads as CPU times or service units

- The RMF support for SMT provides new SMT related metrics to allow capacity planning and performance analysis in SMT environments.
- RMF supports SMT environments by extending the
  - Postprocessor CPU activity report
  - Monitor III CPC capacity report
  - Overview Conditions based on SMF 70.1
- RMF new function APAR OA44101 provides the SMT support for z/OS 2.1.
- The architecture introduced with SMT requires a reinterpretation of existing RMF metrics:
  - CPU metric data can now be on core or thread level granularity
  - z/OS charges CPU time consumed by work units (TCBs, SRBs) in terms of MT-1 equivalent time. The MT-1 equivalent time is the time it would have taken to run the same work in MT-1 mode. All RMF metrics reporting CPU consumption of workloads as CPU time or service units reflect MT-1 equivalent time.

# z13 – SMT: Postprocessor CPU Activity Report



- PP CPU activity report displayed in “old” format when SMT is inactive
- PP CPU activity report provides new metrics when SMT is active
  - MT Productivity and Utilization of each logical core
  - Multi-Threading Analysis section displays MT Mode, MT Capacity Factors and average Thread Density
- One data line in PP CPU activity report represents one thread (CPU)
  - CPU NUM designates the logical core
- Some metrics like TIME % ONLINE and LPAR BUSY provided at core granularity only

CPU ACTIVITY											
z/OS V2R1		SYSTEM ID CB8B				DATE 02/02/2015		INTERVAL 15.00.004			
		RPT VERSION V2R1 RMF				TIME 11.00.00		CYCLE 1.000 SECONDS			
---CPU---		TIME %				MT %		LOG PROC		--I/O INTERRUPTS--	
NUM	TYPE	ONLINE	LPAR BUSY	MVS BUSY	PARKED	PROD	UTIL	SHARE %	RATE	% VIA	TPI
0	CP	100.00	68.07	67.94	0.00	100.00	68.07	100.0	HIGH	370.1	13.90
1	CP	100.00	46.78	46.78	0.00	100.00	46.78	52.9	MED	5.29	16.93
... TOTAL/AVERAGE			8.66	54.17		100.00	8.66	152.9		375.3	13.95
A	IIP	100.00	48.15	41.70	0.00	85.84	41.33	100.0	HIGH		
				35.66	0.00						
B	IIP	100.00	38.50	32.81	0.00	85.94	33.09	100.0	HIGH		
				26.47	0.00						
... TOTAL/AVERAGE			29.48	23.23		86.47	25.39	386.7			
MULTI-THREADING ANALYSIS											
CPU TYPE	MODE	MAX CF	CF	AVG TD							
CP	1	1.000	1.000	1.000							
IIP	2	1.485	1.279	1.576							

MT-2 core capacity used:  
MT Core Productivity  
x TIME % LPAR BUSY

Productivity of logical core while  
dispatched to physical core

- The CPU Activity section reports on logical core and logical processor activity. For each processor, the report provides a set of calculations that are provided at a particular granularity that depends on whether multithreading is disabled (LOADxx PROCVIEW CPU parameter is in effect) or enabled (LOADxx PROCVIEW CORE parameter is in effect).
- If multithreading is disabled for a processor type, all calculations are at logical processor granularity.
- If multithreading is enabled for a processor type, some calculations are provided at logical core granularity and some are provided at logical processor (thread) granularity. The CPU Activity section displays exactly one report line per thread showing all calculations at logical processor granularity. Those calculations that are provided at core granularity are only shown in the same report line that shows the core id in the CPU NUM field and which is representing the first thread of a core.
- The following calculations are on a per logical processor basis when multithreading is disabled and on a per logical core basis when multithreading is enabled
  - Percentage of the interval time the processor was online
  - LPAR view of the processor utilization (LPAR Busy time percentage)
  - Percentage of a physical processor the logical processor is entitled to use
  - Multithreading core productivity (only reported when multithreading is enabled)
  - Multithreading core utilization (only reported when multithreading is enabled)
- The following calculations are on a per logical processor basis regardless whether multithreading is enabled or disabled:
  - MVS view of the processor utilization (MVS Busy time percentage)
  - Percentage of the online time the processor was parked (in HiperDispatch mode only)
  - I/O interrupts rate (general purpose processors only)
  - Percentage of I/O interrupts handled by the I/O supervisor without re-enabling (general purpose processors only)

## z13 – SMT: Monitor III CPC Report



MT Mode and Productivity for zIIP processors

RMF V2R1 CPC Capacity Line 1 of 50  
 Samples: 60 System: CB88 Date: 02/02/15 Time: 11.00.00 Range: 60 Sec  
 Partition: CB88 2964 Model 731  
 CPC Capacity: 3935 weight % of Max: 50.1 4h Avg: 138 Group: N/A  
 Image Capacity: 1777 WLM Capping %: 0.0 4h Max: 177 Limit: N/A  
 MT Mode IIP: 2 Prod % IIP: 80.9

Partition	--- MSU --- Def Act	Cap Def	Proc Num	Logical Effect	Util % Total	- Physical LPAR	Util % - Effect	Total
*CP			390			0.8	43.7	44.5
CB8B	0 192	NO	10.0	15.0	15.1	0.0	4.8	4.9
CB8D	0 134	NO	15.0	7.0	7.0	0.0	3.4	3.4
CB8E	0 330	NO	14.0	18.4	18.6	0.1	8.3	8.4
CB88	0 182	NO	14.0	10.2	10.3	0.0	4.6	4.6
C05	0 140	NO	14.0	7.9	7.9	0.0	3.5	3.6
C06	0 150	NO	14.0	8.4	8.4	0.0	3.8	3.8
LP1	0 507	NO	4.0	100	100	0.0	12.9	12.9

**SMT mode enabled:**  
 Processor data at logical core granularity

**SMT mode disabled:**  
 Processor data at logical processor granularity

- RMF Monitor III CPC report displays performance data for all partitions belonging to the CPC
- If multithreading is enabled the processor data is reported at logical core granularity, otherwise processor data is reported at logical processor granularity
- The report header is enhanced with the information about MT Mode and Productivity for the zIIP processors.
- Additional SMT metrics are available as hidden report header fields:
  - Multi-Threading Maximum Capacity Factor for IIP
  - Multi-Threading Capacity Factor for IIP
  - Average Thread Density for IIP
- These hidden report header fields can be displayed, if the CPC report is invoked in the RMF Data Portal for z/OS web browser frontend.

# z13 – SMT: Postprocessor Workload Activity Report

WORKLOAD ACTIVITY									
z/OS V2R1	SYSPLX	UTCPLXCB	DATE 02/02/2015	INTERVAL 15.00.004	MODE = GOAL				
	RPT VERSION V2R1	RMF	TIME 11.00.00						
REPORT BY: POLICY=BASEPOL									
-TRANSACTIONS-	TRANS-TIME	HHH.MM.SS.TTT	--DASD	I/O--	---SERVICE---	SERVICE TIME	---APPL %---		
AVG 790.12	ACTUAL	27.787	SSCHRT	3975	IOC 4233K	CPU 4659.039	CP 542.89		
MPL 790.12	EXECUTION	15.761	RESP	2.8	CPU 306468K	SRB 1513.624	AAPCP 0.00		
ENDED 9173	QUEUED	1	CONN	1.4	MSO 0	RCT 0.622	IIPCP 2.00		
END/S 10.19	R/S AFFIN	0	DISC	1.2	SRB 97415K	IIT 21.116	AAP N/A		
#SWAPS 6087	INELIGIBLE	0	Q+PEND	0.1	TOT 408116K	HST 0.000	IIP 97.84		
EXCTD 15860	CONVERSION	0	IOSQ	0.0	/SEC 453461	AAP N/A			
AVG ENC 4.00	STD DEV								
REM ENC 0.00									
MS ENC 0.00									
SMT active: MT-1 equivalent service units and service times					ABSRPTN 574				
					TRX SERV 574				

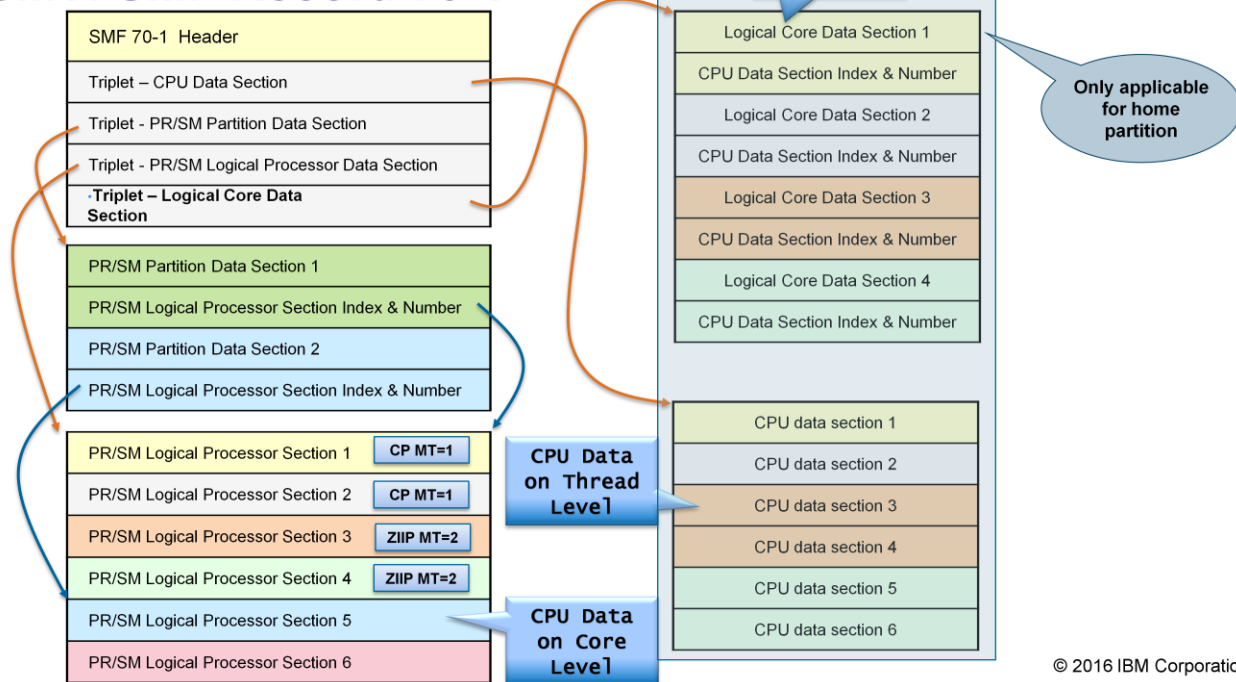
- Pre SMT/ SMT mode inactive:
  - Service time : Logical processor CPU time
  - APPL% : Percentage of logical processor capacity
- SMT mode active:
  - Service time: MT-1 equivalent CPU time
  - APPL% : Percentage of maximum core capacity used as:

$$APPL\% = \frac{MT-1 \text{ equivalent CPU time}}{\text{Interval length} \times mCF} \times 100$$

For MT Mode = 1 ⇔ mCF = 1

- The RMF Postprocessor Workload Activity report (WLMGL) reports the CPU time used by a workload in units of service times and service units. The APPL% metric shows the percentage of logical processor capacity used by the workload.
- With active SMT mode, the service time charged to the workload is based on the MT-1 equivalent CPU time (the CPU time that would have been used in MT-1 mode) . Service units are calculated from MT-1 equivalent CPU time, too. The APPL% now represents the percentage of maximum core capacity used by the workload.
- MT-2 APPL% numbers can continue to be used to understand relative core utilization in a given interval, or at times of comparable Maximum Capacity Factors. However, the Maximum Capacity Factor (mCF) needs to be considered when comparing APPL% across different workloads or times with different mCF values.
- If multithreading is disabled for a processor type, an mCF of 1 is used for the APPL% calculation so that the calculation is the same as before introduction of SMT.

## z13 – SMT: SMF Record 70-1



12

© 2016 IBM Corporation

- The RMF SMT support enhances the SMF 70-1 record with SMT related fields.

- **CPU Control Data Section:**

- Processor Class Maximum Capacity Factor Metrics:
  - SMF70MCF (CP), SMF70MCF5 (zIIP)
- Processor Class Capacity Factor Metrics:
  - SMF70CF (CP), SMF70CFS (zIIP)
- Processor Class Average Thread Density Metric:
  - SMF70ATD (CP), SMF70ATDS (zIIP)

- **PR/SM Partition Data Section:**

- Maximum Thread Id and MT enabled (SMF70MTID)

- **PR/SM Logical Processor Data Section:**

- MT inactive → section is on a per logical processor/CPU basis
- MT active → section is on a per logical core basis
- MT Idle Time (SMF70MTIT)

- With SMT active, there is no longer a 1:1 mapping between PR/SM Logical Processor data sections and CPU data sections. The PR/SM Logical Processor data sections now represent CPU data on logical Core level, the CPU data sections represent CPU data on logical thread level. To identify the CPU data sections belonging to a logical Core, a new Logical Core data section is introduced.

## z13 – SMT: SMF Record 70-1...

SMF record type 70 subtype 1 – Logical Core data section

Offsets	Name	Len	Format	Description
0 0	SMF70_CORE_ID	2	Binary	Core identification.
2 2	SMF70_CORE_FLG	1	Binary	Logical Core Information Bit Meaning when Set 0 Core LPAR Busy time is valid. 1-7 Reserved.
3 3		1		Reserved.
4 4	SMF70_CPU_SKIP	2	Binary	The CPU data sections for this core are grouped together in the record. To get to the first CPU data section associated with this logical core, skip over the number of CPU data sections specified by this field, starting at the first CPU data section in the record.
6 6	SMF70_CPU_NUM	2	Binary	Number of CPU data sections for this core. This value represents the number of threads that are active on this core.
8 8	SMF70_PROD	4	Binary	Multithreading core productivity numerator. Divide this value by 1024 to get the multithreading core productivity. A zero value is reported if the core was not configured ONLINE for the complete interval. If SMF70_CPU_NUM is greater than 1, the core productivity represents the percentage of how much work the core resources accomplished while dispatched to physical hardware over the maximum amount of work the core resources could have accomplished while dispatched to physical hardware.
12 C	SMF70_LPAR_BUSY	4	Binary	Multithreading core LPAR Busy Time in milliseconds. This field is valid if bit 0 of SMF70_CORE_FLG is set.

13

© 2016 IBM Corporation

- A new **Logical Core data section** is added to the SMF record 70 subtype 1 when the SMT is active (LOADxx PROCVIEW CORE parameter is in effect).
- The SMF 70 Subtype 1 Individual Header Extension is extended by a new triplet that describes the new Logical Core data sections.
- **Logical Core Data Section Fields:**
  - SMF70\_CPU\_SKIP and SMF70\_CPU\_NUM can be used to identify the CPU data sections with the thread data of a logical core:
    - Navigate to first CPU on core via a number of CPU data sections to skip (SMF70\_CPU\_SKIP)
    - Number of CPU data sections for this core (SMF70\_CPU\_NUM)
  - Core Productivity Metric (SMF70\_PROD)
  - Core LPAR Busy Metric (SMF70\_LPAR\_BUSY)

## z13 – SMT: SMF 70-1 based Overview Conditions

- New qualifier *coreid* is added to support overview reports for core metrics
  - *coreid* is a processor identifier (one or two hexadecimal digits) that either identifies a logical core (when LOADxx PROCVIEW CORE is in effect) or a logical processor (when LOADxx PROCVIEW CPU is in effect).
  - If the qualifier is omitted, the values represent the average of all logical processors/cores

Condition	Name	Qualifier	Source	Algorithm
Percent Multi-Threading core productivity for zIIPs	IIPPROD	coreid	SMF70_PROD	PROD / 1024
Percent Multi-Threading core utilization for zIIPs	IIPUTIL	coreid	SMF70_PROD SMF70PDT SMF70ONT SMF70_LPAR_BUSY	PROD / 1024 multiplied by value of Overview Condition IIPBSY

- For SMF 70-1 based overview reporting there are new overview conditions that can be used to display the MT Core Productivity and MT Core Utilization for a reporting interval.
- The new overview qualifier *coreid* allows overview reporting at core granularity.

## z13 – SMT: SMF 70-1 based Overview Conditions ...

- Existing qualifier *cpuid* is changed to support overview reports for logical processor and thread metrics
- *cpuid* is a processor identifier which must be in the format *procid[.threadid]*
  - *procid* is a processor identifier (one or two hexadecimal digits) that either identifies a logical core (when LOADxx PROCVIEW CORE is in effect) or a logical processor (when LOADxx PROCVIEW CPU is in effect)
  - *threadid* is an optional thread identifier (0 or 1) that identifies a thread that is executing on the logical core designated by *cpuid*. It is ignored when LOADxx PROCVIEW CPU is in effect
  - If LOADxx PROCVIEW CORE is in effect and *threadid* is omitted, the values represent the average of all threads executing on the logical core
- Examples with PROCVIEW CORE active
  - 0A specifies logical core id 0A
  - 3F.0 specifies thread id 0 of logical core id 3F
  - A.1 specifies thread id 1 of logical core id 0A
- Examples with PROCVIEW CPU active
  - 0A specifies logical processor id 0A
  - 3F.0 specifies logical processor id 3F

- To support overview reporting for logical processor and thread metrics the *cpuid* qualifier can now be specified in the format *cpuid[.threadid]*:
  - *cpuid* is a processor identifier (one or two hexadecimal digits) that either identifies a logical core (when LOADxx PROCVIEW CORE is in effect) or a logical processor (when LOADxx PROCVIEW CPU is in effect).
  - *threadid* is an optional thread identifier (0 or 1) that identifies a thread that is executing on the logical core designated by *cpuid*. It is ignored when LOADxx PROCVIEW CPU is in effect .
  - If LOADxx PROCVIEW CORE is in effect and *threadid* is omitted, the values represent the average of all threads executing on the logical core.



# z13 – SMT: RMF Distributed Data Server Metrics


**Sysplex**

- MVS Image
- I/O Subsystem
- Processor
- Storage
- Enqueue
- Operator
- Subsystems
- CPC
- LPAR**
- Coupling Facility
- CF Structure

Resource	New metric
<b>Sysplex</b>	<ul style="list-style-type: none"> <li>▶ MT mode for IIP by partition</li> <li>▶ MT capacity factor for IIP by partition</li> <li>▶ MT maximum capacity factor for IIP by partition</li> <li>▶ % MT IIP core productivity by partition</li> <li>▶ average thread density for IIP by partition</li> <li>▶ MT mode for CP by partition</li> <li>▶ MT capacity factor for CP by partition</li> <li>▶ MT maximum capacity factor for CP by partition</li> <li>▶ % MT CP core productivity by partition</li> <li>▶ average thread density for CP by partition</li> </ul>
<b>LPAR</b>	<ul style="list-style-type: none"> <li>▶ MT Mode IIP</li> <li>▶ MT capacity factor for IIP</li> <li>▶ MT maximum capacity factor for IIP</li> <li>▶ % MT IIP core productivity</li> <li>▶ average thread density for IIP</li> <li>▶ MT Mode CP</li> <li>▶ MT capacity factor for CP</li> <li>▶ MT maximum capacity factor for CP</li> <li>▶ % MT CP core productivity</li> <li>▶ average thread density for CP</li> </ul>

- The RMF Distributed Data Server (DDS) provides a new set of MT metrics for each processor class:
  - MT Mode
  - MT Core Productivity
  - MT Maximum Capacity Factor
  - MT Capacity Factor
  - Average Thread Density
- These metrics are available as single valued metrics for the LPAR resource and as list valued metrics for the SYSPLEX resource.

## z13 - Support of Crypto Express5S Card



```

CRYPTO HARDWARE ACTIVITY
z/OS V2R2          SYSTEM ID TRX2          DATE 02/01/2014          INTERVAL 14.59.998
RPT VERSION V2R2  RMF          TIME 11.00.00          CY
----- CRYPTOGRAPHIC CCA COPROCESSOR -----
----- TOTAL ----- KEY-GEN -----
TYPE  ID  RATE  EXEC TIME  UTIL%  RATE
CEX5C 0   2.15  227.8    48.9    2.15
----- CRYPTOGRAPHIC PKCS11 COPROCESSOR -----
----- TOTAL ----- OPERATIONS DETAILS -----
TYPE  ID  RATE  EXEC TIME  UTIL%  FUNCTION  RATE  EXEC TIME  UTIL%
CEX5P 1  446.5  0.243    8.3    ASYM FAST  274.3  0.175    2.4
      ASYM GEN   0.00  0.000    0.0
      ASYM SLOW  120.3  0.405    5.3
      SYMM COMPLETE 0.00  0.000    0.0
      SYMM PARTIAL 51.89  0.398    0.6
----- CRYPTOGRAPHIC ACCELERATOR -----
----- TOTAL ----- ME-FORMAT RSA OPERATIONS -- CRT-FORMAT RSA OPERATIONS --
TYPE  ID  RATE  EXEC TIME  UTIL%  KEY  RATE  EXEC TIME  UTIL%  RATE  EXEC TIME  UTIL%
CEX5A 2 1335.5  0.151    30.3  1024 678.2  0.225    14.2  544.4  0.145    5.8
      2048 0.00  0.000    0.0  22.6  0.465    4.8
      4096 0.00  0.000    0.0  90.3  0.378    5.5
----- ICSF SERVICES -----
----- ENCRYPTION ----- DECRYPTION ----- HASH ----- PIN -----
SDES  TDES  AES  SDES  TDES  AES  SHA-1  SHA-256  SHA-512  TRANSLATE  VERIFY
RATE  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
SIZE  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
----- MAC ----- AES MAC ----- RSA DSIG ----- ECC DSIG ----- F-PRESERVING ENCRYPTION -----
GENERATE  VERIFY  GENERATE  VERIFY  GENERATE  VERIFY  GENERATE  VERIFY  ENCIPHER  DECIPHER  TRANSLATE
RATE  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
SIZE  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00

```

Crypto Express5 Card reported with same set of measurements as Crypto Express4

Additional ICSF service statistics

- RMF new function APAR OA43493 introduces the RMF support of the Crypto Express5S feature on the z13
- In detail, RMF collects and reports performance measurements for operations executed on Crypto Express5S CCA coprocessors (CEX5C), PKCS11 coprocessors (CEX5P) and accelerators (CEX5A).
- The crypto measurements are stored to SMF 70 subtype 2 data sections:
  - CEX5C measurements are stored in the Cryptographic Coprocessor data section,
  - CEX5P measurements are stored in the Cryptographic PKCS11 Coprocessor data section and
  - CEX5A measurements are stored in the Cryptographic Accelerator data section.
- The Postprocessor Crypto Activity report provides the crypto measurements from the SMF 70 subtype 2 data sections in the corresponding report sections.
- The ICSF SERVICES report section displays request rates for the new ICSF activities:
  - RSA Digital Signature Generate callable services
  - RSA Digital Signature Verify callable services
  - ECC Digital Signature Generate callable services
  - ECC Digital Signature Verify callable services
  - AES MAC Generate callable services
  - AES MAC Verify callable services
  - FPE Encipher callable services
  - FPE Decipher callable services
  - FPE Translate callable services
- With new function APAR OA43493 the RMF support is available for z/OS V1.13 and z/OSV2.1.

# z13/z13s - RMF Enhancements for new Capping types

## New Capping Techniques:

- LPAR Absolute Group Capping
  - Limits CPU consumption of LPAR groups to a specified physical hardware group capping limit
  - Extends LPAR Absolute Capping type to group of LPARs
  - LPAR Group limit defined on HMC
  - Allows physical partitioning of CPC
- Absolute MSU Capping
  - Similar to WLM defined capacity or group capacity but LPAR will be always capped
  - Independent of 4 hour rolling average consumption.
  - General purpose processor
  - Controlled by new IEAOPTxx option `AbsMSUcapping = YES|NO`

## Comparison of capping types

Type of capping	Scope	Specification unit	Proc types	Suitable to isolate LPARs or LPAR groups against others	Control point
Initial (hard) capping	LPAR	LPAR share of CPC capacity	Any	+	SE/HMC
LPAR Absolute capping	LPAR	Fractional #processors		+	
LPAR Absolute Group Capping	Group of LPARs	Fractional #processors		+	
Defined capacity (DC, soft capping)	LPAR	MSU (4HRA)	CP	-	SE/HMC IEAOPT
LPAR group capacity (GC, soft capping)	Group of LPARs	MSU (4HRA)		-	
Absolute MSU Capping	LPAR or Group	MSU		+(CP only)	
Resource group capping	Groups of service classes in Sysplex or per LPAR	Unweighted CPU SU/sec, fraction of LPAR share, or fractional #CPs	CP*	N/A	WLM Service Definition
Logical configuration	LPAR	Integer #processors	Any	+ but coarse grain	HMC + OS

PR/SM controlled
  WLM controlled, PR/SM enforced
  WLM controlled

With availability of IBM z13 GA2 and z13s two new capping types are introduced:

- Physical Hardware Group Capping (also called LPAR Absolute Group Capping)
- Absolute MSU Capping

### LPAR Absolute Group Capping

- LPAR Absolute Group Capping allows to specify a limit for the physical usage of processors within a group of logical partitions. This concept is similar to the Absolute LPAR Capping function that was introduced with zEC12.
- The LPAR Absolute Group Capping limit is specified on HMC.
- With LPAR Absolute Group Capping, software charges can be based on the enforced capacity limits for a fixed group of images
- Feature of PR/SM as of z13 GA2, and z13s

### Absolute MSU Capping

- Absolute MSU capping can be activated on a per system level by specifying `AbsMSUcapping=YES` in the IEAOPTxx member. It influences how WLM enforces the LPAR defined capacity limit or group capacity limit specified at the Support Element (SE) or HMC.
- With absolute MSU capping WLM applies *always* a cap to the partition to limit its consumption to the effective limit, independently of the *four-hour rolling average* consumption. Therefore, absolute MSU capping is an effective means to permanently limit the consumption of an LPAR to a specific MSU figure, including times when the *four-hour rolling average* does not exceed the defined limit.
- When absolute MSU capping is used with an LPAR capacity group, the limit on behalf of the group entitlement will always be enforced, regardless of the *four-hour rolling group average* consumption.
- z13 GA2 is not a prerequisite for using absolute MSU capping.

# z13/z13s- Capping: Postprocessor CPU Activity Report

- New Hardware Group Report section in RMF PP CPU Activity report
- Displays LPAR Absolute Group Capping settings as defined on HMC



H A R D W A R E   G R O U P   R E P O R T

z/OS V2R2      SYSTEM ID R74      DATE 11/19/2015      INTERVAL 04.59.999  
 RPT VERSION V2R2 RMF      TIME 06.21.00      CYCLE 1.000 SECONDS

HW GROUP NAME	PARTITION	SYSTEM	CP	IIP	ICF	IFL
BOEB1	R74	R74	1.50	2.00	0.00	0.00
BOEB2	R75	R75	1.00	2.85	0.00	0.00

HW group name as defined on HMC

HW group members

Capacity limit that can be used by the LPAR group. Limit specified as number of processors.

also available in XML version of CPU report

- With APAR OA48688, RMF provides the support for the gathering and reporting of LPAR Absolute Group Capping limits on the IBM z13 GA2 and z13s server. It also adds support for the new Absolute MSU Capping provided by new WLM OPT parameter ABSMSUCAPPING.
- PTFs are available for z/OS V2.1 and z/OS V2.2
- The RMF support for LPAR Absolute Group Capping introduces a new Hardware Group Report section in the RMF Postprocessor CPU Activity report. This report section displays hardware group capping settings of the hardware groups and their partitions.
- Fields in the Hardware Group Report section:

HW GROUP NAME	Name of the hardware group.
PARTITION	Name of the logical partition.
SYSTEM	Name of the z/OS system.
HW GROUP LIMIT	Absolute limit on partition usage of all CPs / zIIPs / ICFs / IFLs which are members of the same hardware group, in terms of numbers of CPUs. If the hardware group name or the limit changed during the reporting interval, an '*' is appended.

## z13/z13s- Capping: Postprocessor CPU Activity Report

Image capacity considers an LPAR Absolute Group limit

At a glance: Capping options that are active for the partition

INITIAL CAP NO  
LPAR HW CAP YES  
HW GROUP CAP YES  
ABS MSU CAP YES

Sum of current weightings of shared processors

Combination of HW capping options:  
• Initial capping  
• LPAR Absolute capping  
• LPAR Absolute Group capping

```

PARTITION DATA REPORT
z/OS V2R2          SYSTEM ID R75          DATE 11/05/2015      INTERVAL 00.26.421
                   RPT VERSION V2R2 RMF    TIME 07.40.33       CYCLE 1.000 SECONDS

MVS PARTITION NAME R75          PHYS PROC NUM 141     GROUP NAME BOEB2
IMAGE CAPACITY     7777        CP            117     LIMIT      0*
NUMBER OF CONFIGURED PARTITIONS 51          ICF           8      AVAILABLE  N/A
WAIT COMPLETION    NO              IIP          16
DISPATCH INTERVAL DYNAMIC

----- PARTITION DATA -----
NAME  S  WGT  DEF  ACT  -DEF-  WLM%  NUM  TYPE  EFFECTIVE  TOTAL  EFFECTIVE  TOTAL  LPAR  MGMT  EFFECTIVE  TOTAL
-----MSU----- --CAPPING-- -- LOGICAL PARTITION PROCESSOR DATA -- -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --
PROCESSOR  ---DISPATCH TIME DATA--- LOGICAL PROCESSORS  --- PHYSICAL PROCESSORS ---
R75  A  10   0    4  N Y Y  0.0  80.0  CP  00.00.00.788  00.00.01.083  0.04  0.05  0.01  0.03  0.04
S5A  A  525  0   143 N N N  0.0  108.0 CP  00.00.37.413  00.00.38.872  1.31  1.36  0.05  1.21  1.26
S5D  A  500  0    62 N N N* 0.0  64.0  CP  00.00.15.851  00.00.16.857  0.94  1.00  0.03  0.51  0.55
S5H  A  475  0   531 N*N N  0.0  80.0  CP  00.02.23.383  00.02.24.436  6.78  6.83  0.03  4.64  4.67
R74  A   73  0    15 Y N Y  0.0   15  CP  00.00.03.951  00.00.03.996  1.00  1.01  0.00  0.13  0.13
S5B  A  500  0   168 N N N  0.0  32.0  CP  00.00.44.737  00.00.45.778  5.29  5.41  0.03  1.45  1.48
*PHYSICAL*
TOTAL 2083  00.04.06.124  00.04.16.025  0.31  7.96  8.28

```

- The Partition Data Section in the Postprocessor CPU Activity report displays whether an Initial Capping option, an absolute LPAR capping limit or an absolute hardware group capping limit was active during a reporting interval.
- In the report header you will also see, if absolute MSU capping is active (OPT parameter ABSMSUCAPPING=YES) on the reported system.
- Even when an LPAR is not limited by its weight, its defined capacity or a group capacity limit, it can be limited by this hardware group capping value. The CPU capacity available to an MVS image (Image capacity) is the minimum of the following:
  - The capacity based on the partition's logical CP configuration (both online and offline)
  - The defined capacity limit of the partition, if available (image softcap)
  - The capacity limit of the related WLM capacity group, if the partition belongs to a capacity group
  - The absolute physical hardware capping limit
  - The capacity based on the hardware group capping limit
- New and changed fields in the Partition Data Section of Postprocessor CPU Activity report

## Header fields:

INITIAL CAP	This field indicates whether the operator has set 'Initial Capping ON' in the logical partition controls of the Hardware Management Console (HMC) for the partition.
LPAR HW CAP	This field indicates whether an absolute physical hardware capping limit has been defined in the logical partition controls of the HMC for any processor type of the partition.
HW GROUP CAP	This field indicates whether an absolute hardware group capping limit has been defined in the logical partition group controls of the HMC for any processor type of the partition.
ABS MSU CAP	This field indicates whether the ABSMSUCAPPING parameter has been set in the active IEAOPTxx parmlib member for the partition.

## Logical Partition Processor Data fields:

CAPPING DEF	The hardware capping option of the partition: This is a string denoting whether hardware capping mechanisms have been applied in the logical partition controls of the HMC for the partition. The values in the first, second and third position of the string are either Y (Yes) or N (NO) and have the following meaning: <ol style="list-style-type: none"> <li>The value 'Y' identifies that 'Initial Capping ON' has been set.</li> <li>The value 'Y' identifies that an absolute physical hardware capping limit (maximal number of CPUs) has been defined.</li> <li>The value 'Y' identifies that an absolute hardware group capping limit (maximal number of CPUs) has been defined.</li> </ol> The field is only useful to logical partitions with shared processors.
-------------	--

## z13/z13s- Capping: Monitor II OPT report

RMF - OPT Settings Line 1 of 37

CPU= 3/ 1 UIC= 65K PR= 0 System= SYSE Total

OPT: P0 Time: 02/05/16 12:30:01

Parameter	Default	Value	Unit	Description
ABNORMALTERM	Yes	Yes	Y/N	Abnormal terminations in routing
ABSMSUCAPPING	No	Yes	Y/N	Absolute, permanent MSU capping
BLWLINTHD	20	20	sec	Time blocked work waits for help
BLWLTRPCT	5	5	0/00	CPU cap. to promote blocked work
CCCAWMT	3200	3200	usec	Alternate wait management time
CCCSIGUR	45	24	msec	Min. mean-time-to-wait threshold
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
FULLPRESYSTEM	No	No	Y/N	System AS can preempt other work
HIPERDISPATCH	Yes	Yes	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

Display current setting of IEAOPTxx parmlib parameter ABSMSUCAPPING.

- The RMF Monitor II OPT Settings report is enhanced to display the information about the new IEAOPTxx parmlib parameter ABSMSUCAPPING.

## z13/z13s- Capping: Monitor III CPC report

RMF V2R2 CPC Capacity Line 1 of 72

Samples: 120 System: SYSF Date: 11/17/15 Time: 17.02.00 Range: 120 Sec

Partition: SYSF 2827 Model 732  
 CPC Capacity: 3665 weight % of Max: \*\*\*\* 4h Avg: 3 Group: RMFGRP  
 Image Capacity: 458 WLM Capping %: 0.0 4h Max: 7 Limit: 4000  
 MT Mode IIP: N/A Prod % IIP: N/A AbsMSUCap: Y

Partition	--- MSU ---	Cap Def	Proc Num	Logical Effect	Util % Total	- Physical Util % -	LPAR Effect	Total
*CP			419			0.9	26.7	27.6
COH2	0 6	N N N	4.0	1.3	1.4	0.0	0.2	0.2
COH3	0 8	N Y Y	4.0	1.6	1.6	0.0	0.2	0.2
IRD7CFE	0 0	Y Y N	1.0	0.1	0.1	0.0	0.0	0.0
R35LP43	0 1	N N N	2.0	0.1	0.2	0.0	0.0	0.0
R35LP44	0 270	N N N	2.0	93.1	93.3	0.0	16.9	17.0
R35LP45	0 1	N N N	3.0	0.2	0.3	0.0	0.1	0.1
R35LP56	0 0	N N N	2.0	0.0	0.0	0.0	0.0	0.0
TRX1	0 4	N N N*	3.0	0.8	0.9	0.0	0.2	0.2

Combination of HW capping options:

- Initial capping
- LPAR Absolute capping
- LPAR Absolute Group capping

Current setting of ABSMSUCAPPING parameter

- The Monitor III CPC Capacity report is modified to provide complete information about the hardware capping mechanisms defined for a partition. The CPC report header displays the setting of the ABSMSUCAPPING OPT parameter.
- The 'Image Capacity' header field takes the hardware group capping limits into account when reporting about the processor capacity available to the z/OS image (measured in MSUs per hour).
- The meaning of field 'Cap Def' has changed:

## Cap Def

The hardware capping option of the partition: This is a string denoting whether hardware capping mechanisms have been applied in the logical partition controls of the Hardware Management Console (HMC) for the partition.

The values in the first, second and third position of the string are either Y (Yes) or N (NO) and have the following meaning:

1. The value 'Y' identifies that 'Initial Capping ON' has been set.
2. The value 'Y' identifies that an absolute physical hardware capping limit (maximal number of CPUs) has been defined.
3. The value 'Y' identifies that an absolute hardware group capping limit (maximal number of CPUs) has been defined.

An asterisk (\*) to the right of a value indicates that the capping status is currently changing

# z13/z13s- Capping: RMF Data Portal CPC Report

Current setting of ABSMSUCAPPING parameter

**RMF Report [R75,MVS\_IMAGE] : CPC (Central Processor Complex)**

Time Range: 11/19/2015 08:25:00 - 11/19/2015 08:26:00

Partition Name: R75	CPU Type: 2964	CPU Model: 7B7	CPC Capacity (MSU/h): 11374
Weight % of Max: 20.0	4h MSU Average: 4	Capacity Group Name: BOEB2	Image Capacity: 5555
WLM Capping %: 0.0	4h MSU Maximum: 29	Capacity Group Limit: 5555	Less than 4h in Capacity Group: Y
MT Mode IIP: 1	MT IIP Core Productivity: 100	Absolute MSU Capping: N	Proj Time until Capping: 14400
Proj Time until Group Capping: 14400	4h Unused Group Capacity Average: 154	CPC sequence number: 0000000000819E7	CPC name: S89
# CP Processors: 117	# ICF+IFL+AAP Processors: 0	# AAP Processors: 0	# ICF Processors: 8
# IFL Processors: 0	# IIP Processors/Cores: 16	Configured Partitions: 51	Wait Completion: NO
			# Dedicated IIPs: 0
			Vary CPU management available: NO
			MT CP Core Productivity: 100
			MT Capacity Factor CP: 1.00
			Physical Total % of shared CPs: 1.4
			Physical Total % of shared IFLs: 0.0

Group Capacity Name	Group Capacity Limit	Group Capacity Min Entitlement	Group Capacity Max Entitlement	Central Storage Online (MB)	User Partition ID	Initial Capping Option	Absolute Capping Limit (CPUs)	Hardware Group Name	Hardware Group Capping Limit (CPUs)
				8192	45	NO			
				393216	53	NO			
				65536	58	NO			
BOEB1	0	N/A	N/A	20480	68	YES		BOEB1	33.33
BOEB1	0	0	0	10240	85	NO		BOEB1	33.33
				307200	10	NO			
				524288	75	NO			
				3145728	29	NO			

HW group name and capping limit for CPs, zIIPs, ICFs and IFLs, if set on HMC



# RMF Enhancements at a Glance

- IBM z13 and z13s Support
  - RMF Statistics for Simultaneous Multithreading (SMT)
  - Extended ICSF Measurements for Crypto Express5S
  - Monitoring of new capping types



- SCM I/O Adapter Performance Reporting
  - Statistics for EADM Subchannel Activities
  - Performance and Thruput on SCM Card Level

- PCIe Activity Reporting
  - z/OS V2.2 Monitor III PCIe Report
  - Additional Measurements for PCIe attached RoCE and zEDC Devices on z13
  - RMF support for SMC-D over ISM on z13 and z13s



- z/OS V2.2 Monitor III Job Resource Consumption Reporting
  - CPU, I/O and Storage Consumption Data at a Glance
  - Detailed Statistics for Job related GQSCAN Activities

- z/OS V2.2 zFS Reporting Enhancements
  - Sysplex-wide Statistics for zFS Usage and Performance
  - Improved zFS Data Gathering Performance
  - Additional zFS Statistics for Shared File System environments



- DS8000 Support
  - RMF Monitoring of SuperPAV capability

- Mobile Workload Reporting
  - Statistics about total and mobile CPU consumption



# z/OS 2.2 - SCM I/O Adapter Performance Reporting


## What is Flash Express?

- ▶ New memory hierarchie of the zSeries family
- ▶ Delivers tier within the fast Solid State Drive (SSD) technology
- ▶ Also denoted as Storage Class Memory (SCM)
- ▶ Integrated on PCI Express attached RAID 10 Cards
  - ⇒ Packaged as two card pair
  - ⇒ Each card holds 1.4 TB of memory per mirrored card pair
  - ⇒ Maximum value of four card pairs delivers up to 5.6 TB of memory



## How is Flash Express exploited by z/OS ?

- ▶ Designed for improved paging performance
- ▶ Compelling addition to traditional auxilliary storage
- ▶ Supports Pageable Large Pages, e.g. with DB2 and Java workloads
- ▶ Eliminates delays from SVC or standalone dump processing
- ▶ Provides overflow capability for MQSeries list structures




25

- The Flash Express feature, introduced with the IBM zEnterprise EC12 (zEC12) server, is a new memory hierarchie of the zSeries family.
- It consists of non-volatile storage using solid state devices on a PCIe card form factor.
- Flash Express implements a new tier of memory, called Storage Class Memory (SCM).
- Flash Express Cards are installed in pairs, which provides mirrored data to ensure a high level of availability and redundancy.
- In each Flash Express card, the data is stored in four solid-state disks in a RAID configuration. If a solid-state disk fails, the data are reconstructed dynamically. The cards in a pair mirror each other over a pair of cables, in a RAID 10 configuration. If either card fails, the data is available on the other card.
- Each Flash Express card has a capacity of 1.4 TB of usable storage.
- A maximum of four pairs of cards can be installed on a zEC12, for a maximum capacity of 5.6 TB of storage.
- Flash Memory is assigned to partitions like Main Memory from the allocation panel on the zEC12 Service Element (SE)
- The Flash Express technology is way faster then SSD technology. It provides access times within the microsecond range.
- z/OS can use Flash Express storage as Storage Class Memory (SCM) for paging.
- Flash Express helps to improve paging performance since page access time from Flash Express is faster than from DASD devices.
- The z/OS paging subsystem can work with a mix of Flash Express storage and External Disk.
- z/OS detects whether Flash Express storage is assigned to the LPAR and will try to page to Flash Express before using paging datasets on DASD.
- In combination with the new pageable 1MB pages, Flash Express helps to improve the performance of DB2 and Java workloads.
- Latency delays in SVC or standalone dump processing caused by page-ins from DASD can be significantly reduced by Flash Express.
- The Coupling Facility can exploit SCM as overflow capacity for list structure data. This functionality can be used by MQSeries to avoid structure-full conditions.

## EADM Subchannels

- I/O is managed by Extended Asynchronous Data Mover subchannels (EADM)
- Similar to standard I/O subchannels, but no channel path or device number assigned
- Created automatically at IPL time
- Not tied to a particular device, no association between subchannel and SCM card pair
- Any EADM subchannel can be used to access storage on any SCM card assigned to the LPAR




- When z/OS needs to read or write data from Storage Class Memory, it creates a new type of channel program and issues a SSCH to a special subchannel called an Extended Asynchronous Data Mover (EADM) subchannel. EADM subchannels are similar to I/O subchannels where you can issue I/O instructions such as SSCH to run channel programs and receive I/O interrupts. However, they do not have channel paths or device numbers assigned, and they are not defined in the I/O configuration. They are created automatically at IPL time.
- Unlike I/O subchannels, which are tied to a particular device, there is no association between an EADM subchannel and a storage increment or SCM card pair. Any EADM subchannel can be used to access storage on any SCM card assigned to the LPAR. The EADM subchannel is simply a certain kind of vehicle for accessing Storage Class Memory

## z/OS 2.2 - SCM Data Gathering

- RMF Monitor III Data Gatherer collects SCM I/O performance statistics frequently

- New Monitor III SCM Activity report

- Data gathering for Monitor III report controlled by Monitor III gatherer option SCM | NOSCM



```

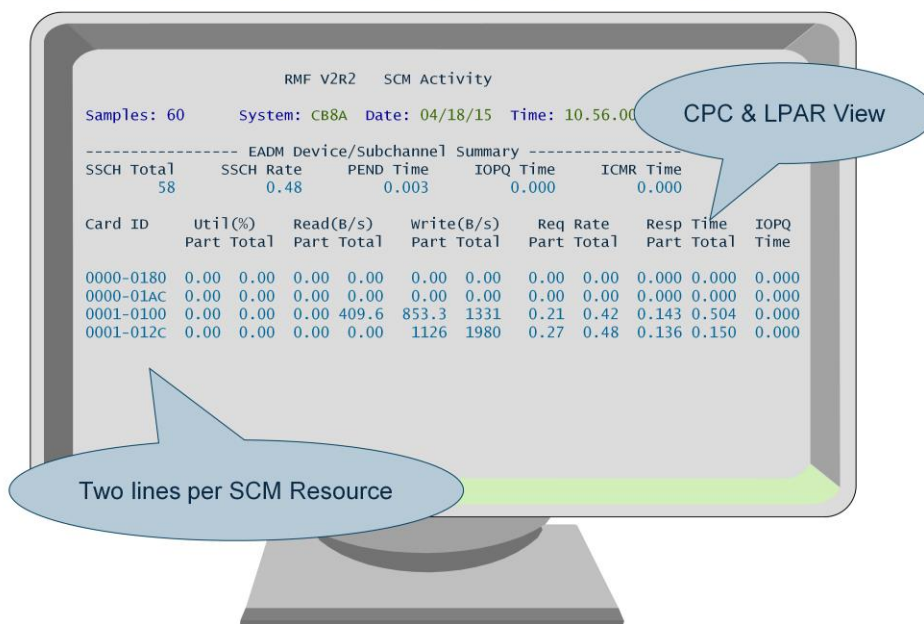
/*****
/* NAME:          ERBRMF04
/* DESCRIPTION:  PARMLIB MEMBER WITH RMF MONITOR III GATHERER OPTIONS
/*              (ALL OPTIONS ARE SET TO DEFAULTS)
/*****
CYCLE(1000)          /* SAMPLE EVERY SECOND (1000 MSEC)
DATASET(STOP)       /* NO DATASET SUPPORT
DATASET(NOSWITCH)   /* APPEND TO LAST NON-FULL DATASET
DATASET(WHOLD(7))   /* CONTROLS BUFFER PAGES IN STORAGE
MINTIME(100)        /* LENGTH OF MINTIME
NOOPTIONS           /* DO NOT DISPLAY OPTIONS
RESOURCE(*JES2,JES2) /* SPECIFIES JES STARTED TASK NAME
NOSTOP              /* RUN UNTIL OPERATOR ISSUES STOP
SYNC(00)           /* MINTIME SYNCHRONIZATION
. . .
PCIE                /* ACTIVATE PCIE DATA GATHERING
SCM                 /* ACTIVATE SCM DATA GATHERING
ZFS                 /* ZFS DATA GATHERING

```

- SCM performance data can be written to the new SMF 74 subtype 10 record
- SMF 74-10 data collection controlled by active SMFPRMxx parmlib settings
- New RMF Postprocessor SCM Activity report created by REPORTS(SCM) control statement
- Only available in XML format

- With RMF for z/OS 2.2 RMF Monitor III collects new SCM I/O performance statistics
- The data can be reported in the new RMF Monitor III SCM Activity report
- Data gathering option SCM / NOSCM was added to RMF Monitor III and controls data collection for the Monitor III SCM Activity report
- Default value SCM is set in shipped Monitor III PARMLIB member ERBRMF04
- If the currently active SMFPRMxx parameter settings indicate that SMF record type 74 subtype 10 is to be collected, the SCM performance data collected by RMF Monitor III is written as new SMF 74 subtype 10 record
- A new Postprocessor (PP) PCIE report can be created by use of REPORTS(PCIE) in the RMF PP JOB control statements. The report is only available in XML format

## z/OS 2.2 - Monitor III SCM Activity Report



The Monitor III SCM Activity report consists of two sections:

- Header section with global EADM subchannel statistics
- Tabular section with performance and throughput statistics on SCM Card level

The EADM subchannel activities are balanced by the system and no tuning is applicable for this kind of resource. Hence, the report provides just the accumulated statistics for all subchannels, no statistics on subchannel level.

One SCM Card which is plugged into the PCIe I/O drawer is also denoted as SCM resource.

Since one SCM resource can hold two internal cards, the tabular part of the SCM Activity report shows two lines per SCM resource.

In other words, the term *Card ID* consists of a prefix which identifies the SCM resource and a suffix which identifies the internal card.

The card mirrors are not visible in any way on the SCM Activity report.

The SCM card statistics are available for the local partitions as well as for the entire CPC.

# z/OS 2.2 - RMF Data Portal SCM Activity Report

[http://hostname:8803/gpm/rmf3.xml?report=SCM&resource=\\_sysid.MVS\\_IMAGE](http://hostname:8803/gpm/rmf3.xml?report=SCM&resource=_sysid.MVS_IMAGE)

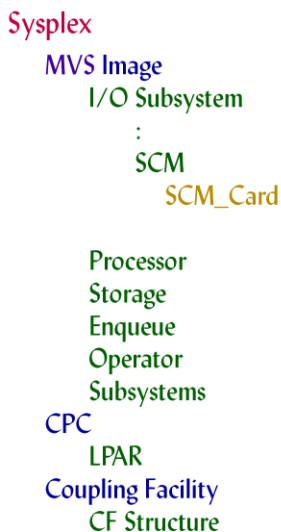
Card ID	Util% (LPAR)	Util% (Total)	Read B/Sec (LPAR)	Read B/Sec (Total)	Write B/Sec (LPAR)	Write B/Sec (Total)	Request Rate (LPAR)	Request Rate (Total)	Avg Response Time (LPAR)	Avg Response Time (Total)	Avg IOP Queue Time (Total)	Requests (LPAR)	Requests (Total)
0001-0100	0.00	0.00	136.5	136.5	68.27	273.1	0.05	0.10	0.256	0.213	0.000	3	6
0000-0180	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.000	0.000	0.000	0	0
0000-014C	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.000	0.000	0.000	0	0
0001-012C	0.00	0.00	68.27	136.5	0.00	0.00	0.02	0.03	0.128	0.256	0.000	1	2

Sort Rows

Fields not available in ISPF Report

- The statistics of the Monitor III SCM ISPF Report are available also by means of the Monitor III Data Portal without limitations.
- The SCM Report can be selected from the report list and basically all report columns can be displayed in the browser window.

# z/OS 2.2 – RMF Distributed Data Server SCM Metrics



Resource	New metric
SCM	<ul style="list-style-type: none"> <li>▶ # of SSCH</li> <li>▶ function pending time</li> <li>▶ initial command response time</li> <li>▶ IOP queue time</li> <li>▶ SSCH rate</li> </ul>
SCM (by card) SCM_CARD	<ul style="list-style-type: none"> <li>▶ % partition utilization</li> <li>▶ % total utilization</li> <li>▶ average partition response time</li> <li>▶ average total response time</li> <li>▶ average total IOP queue time</li> <li>▶ partition byte read rate</li> <li>▶ partition byte write rate</li> <li>▶ partition request rate</li> <li>▶ total byte read rate</li> <li>▶ total byte write rate</li> <li>▶ total request rate</li> </ul>

- The RMF DDS Resource Model represents a composition of resources that can exist in a Parallel Sysplex environment
  - MVS\_IMAGE is a child resource of resource SYSPLEX
  - I/O\_SUBSYSTEM is a child resource of resource MVS\_IMAGE
  - Child resource SCM is added to resource I/O\_SUBSYSTEM
  - Child resource SCM\_CARD is added to resource SCM
  
- A variety of metric values that are related to resource types SCM and SCM\_CARD can be requested from the RMF Distributed Data Server (DDS)
  
- Alternatively, the browser based version of the report can be requested from the RMF Distributed Data Server (DDS) by using the following URL: [http://hostname:8803/gpm/rmf3.xml?report=SCM&resource=,sysname,MVS\\_IMAGE](http://hostname:8803/gpm/rmf3.xml?report=SCM&resource=,sysname,MVS_IMAGE)

## RMF Enhancements at a Glance

- IBM z13 and z13s Support
  - RMF Statistics for Simultaneous Multithreading (SMT)
  - Extended ICSF Measurements for Crypto Express5S
  - Monitoring of new capping types
- SCM I/O Adapter Performance Reporting
  - Statistics for EADM Subchannel Activities
  - Performance and Thruput on SCM Card Level
- PCIE Activity Reporting
  - z/OS V2.2 Monitor III PCIE Report
  - Additional Measurements for PCIe attached RoCE and zEDC Devices on z13
  - RMF support for SMC-D over ISM on z13 and z13s
- z/OS V2.2 Monitor III Job Resource Consumption Reporting
  - CPU, I/O and Storage Consumption Data at a Glance
  - Detailed Statistics for Job related GQSCAN Activities
- z/OS V2.2 zFS Reporting Enhancements
  - Sysplex-wide Statistics for zFS Usage and Performance
  - Improved zFS Data Gathering Performance
  - Additional zFS Statistics for Shared File System environments
- DS8000 Support
  - RMF Monitoring of SuperPAV capability
- Mobile Workload Reporting
  - Statistics about total and mobile CPU consumption





## z/OS V2.2 – RMF Monitor III PCIE Activity Report

- RMF Monitor III Data Gatherer collects PCIe performance statistics frequently
- z/OS V2.1 introduced RMF Postprocessor PCIe Activity Report for after-the-facts analysis with SMF 74.9 data

Default value PCIE set in shipped PARMLIB member ERBRMF04

### z/OS V2.2

- The new RMF Monitor III PCIE Activity Report provides detailed short-term information about PCIe Express based functions. Currently supported functions are:
  - z Enterprise Data Compression (zEDC)
  - Shared Memory Communication via RDMA (SMC-R)
- New Monitor III data gatherer option  
PCIE | NOPCIE

```

/*****/
/* NAME:          ERBRMF04 */
/* DESCRIPTION:  PARMLIB MEMBER WITH RMF MONITOR III GATHERER OPTIONS */
/*              (ALL OPTIONS ARE SET TO DEFAULTS) */
/*****/
CYCLE(1000)      /* SAMPLE EVERY SECOND (1000 MSEC) */
DATASET(STOP)   /* NO DATASET SUPPORT */
DATASET(NOSWICH) /* APPEND TO LAST NON-FULL DATASET */
DATASET(WHOLD(7)) /* CONTROLS BUFFER PAGES IN STORAGE */
MINTIME(100)    /* LENGTH OF MINTIME */
NOOPTIONS       /* DO NOT DISPLAY OPTIONS */
RESOURCE(*JES2,JES2) /* SPECIFIES JES STARTED TASK NAME */
NOSTOP         /* RUN UNTIL OPERATOR ISSUES STOP */
SYNC(00)       /* MINTIME SYNCHRONIZATION */
.....
VSAMRLS        /* ACTIVATE VSAM RLS GATHERING */
OPD           /* ACTIVATE OMVS PROCESS DATA GATHERING */
PCIE          /* ACTIVATE PCIE DATA GATHERING */
ZFS           /* ACTIVATE ZFS DATA GATHERING */

```

- New Monitor III PCIE Activity report allows short-term performance analysis of RoCE devices and zEDC hardware accelerators. Before z/OS 2.2, performance problems on PCIe devices and hardware accelerators could only be analyzed after the facts using SMF 74.9 / RMF Postprocessor, now RMF online monitoring can be used to identify performance problems on short notice when they appear.
- Users can control whether or not they want Monitor III to collect PCIe activity data by specifying data gathering option: PCIE | NOPCIE
- Default value PCIE is set in shipped PARMLIB member ERBRMF04

## z/OS V2.2 – RMF Monitor III PCIE Activity Report

- Request the Monitor III PCIE Activity Report by selection **3** from the Primary Menu & selection **14** from the Resource Report Selection Menu or enter command: **PCIE** or **PCI**
- Metrics on the report main panel are independent of the type of the exploited hardware feature and reflect the activity of the z/OS system on which RMF data collection took place.


 ISPFR Report

```

RMF V2R2  PCIE Activity
Command ==>
Line 1 of 3
Scroll ==> CSR
Samples: 300  System: TA2  Date: 02/18/15  Time: 22.20.00  Range: 300  Sec
----- Function ----- Alloc - PCI Operations Rate - -Xfer Rate -
ID  PCID Type Jobname ASID Status Time% Load Store Block Refr Read Write
006C 0204 zEDC FPGHWAM 0013 Alloc 100 0 102 0 14.0
007C 025C zEDC FPGHWAM 0013 Alloc 100 0 102 0 14.0
00A2 013C RoCE VTAM390 002C Alloc 100 0.113 5999 0 1.69 0.267 288
  
```

- The Monitor III PCIE Activity Report can be used to investigate performance problems that are related to PCI Express based functions.
- On the main panel, metrics are displayed that are independent of the type of the exploited hardware feature and reflect the activity of the z/OS system on which RMF data collection took place.

# z/OS V2.2 – RMF Monitor III PCIE Activity Report

- Additional metrics are displayed for **zEDC Accelerators** on pop-up panel HW Accelerator And Compression Activity

Command ==> RMF V2R2 PCIE Activity Line 1 of 3  
Scroll ==> CSR

Samples: 300 System: TA2 Date:

FUNCTION					
ID	PCID	Type	Jobname	ASID	Status
006C	0204	ZEDC	FPGHWAM	0013	Alloc
007C	025C	ZEDC	FPGHWAM	0013	Alloc
00A2	013C	RoCE	VTAM390	002C	Alloc

RMF Hardware Accelerator And Compression Activity

Press Enter to return to the Report panel.

Function ID : 006C Alloc Time % : 99.7 More: +  
Allocated : 22.03.22 on 02/18/15

**Hardware Accelerator**  
Time Busy % : 0.286 Transfer Rate : 4.87

**Request**  
Execution Time : 28.0 Std. Deviation: 8.07  
Queue Time : 65.7 Std. Deviation: 140  
Size : 47.6

**Buffer Pool**  
Memory Size : 16 utilization : 0

	Compression	Decompression
Request Rate	: 102	0.437
Throughput	: 2.91	0.009
Ratio	: 2.79	0.652

**DELAYJ**

- No additional metrics, no pop-up panel for **RoCE devices**

- If cursor is placed on one of the cursor sensitive fields Function ID, Function PID, or Function Type, additional metrics are displayed for zEDC Accelerators on pop-up panel HW Accelerator And Compression Activity
- If cursor is placed on one of the cursor sensitive fields Function ID, Function PID, or Function Type, message “No additional information available” is displayed when selected PCIE function is a RoCE device

# z/OS V2.2 – RMF Data Portal PCIE Activity Report

RMF Report - One Row [.,SYSF,MVS\_IMAGE] : PCIE (PCIE Activity)

Time Range: 05/04/2015 10:02:00 - 05/04/2015 10:03:00

Function Id	0022
Function PID	037C
Function Type	zEDC
Jobname	FPGHWAM
ASID	0012
Function Status	Alloc
% Alloc Time	100
Allocation Date	04/30/15
Allocation Time	16.02.43
Load Operations Rate	0
Store Operations Rate	0
Store Block Operations Rate	0
Refresh PCI Translations Operations Rate	0
Xfer Read Rate	0
Xfer Write Rate	0
HWA Time % Busy	0
HWA Transfer Rate	
Request Execution Time	
Request Queue Time	
Request Size	
Request Execution Time StdDev	
Request Queue Time StdDev	
Compression Request Rate	0
Decompression Request Rate	0
Compression Throughput	0
Decompression Throughput	0
Compression Ratio	
Decompression Ratio	
Buffer Pool Memory Size	16
Buffer Pool % Utilization	
# DMA AS	1
HWA Type	zCompression
Received Packets Rate	
Transmitted Packets Rate	
Work Unit Rate	
Adapter Utilization	

[http://hostname:8803/gpm/rmf3.xml?report=PCIE&resource=,sysid,MVS\\_IMAGE](http://hostname:8803/gpm/rmf3.xml?report=PCIE&resource=,sysid,MVS_IMAGE)

RMF Data Portal for z/OS Home Explore Overview My View ?

RMF Report [.,SYSF,MVS\_IMAGE] : PCIE (PCIE Activity)

Time Range: 05/04/2015 10:02:00 - 05/04/2015 10:03:00

Function Id	Function PID	Function Type	Jobname	ASID	Function Status	% Alloc Time	Allocation Date	Allocation Time	Load Operations Rate
0022	037C	zEDC	FPGHWAM	0012	Alloc	100	04/30/15	16.02.43	0
0026	037C	zEDC	FPGHWAM	0012	Alloc	100	04/30/15	16.02.43	0
0029	03BC	zEDC	FPGHWAM	0012	Alloc	100	04/30/15	16.02.43	0
002C	03BC	zEDC	FPGHWAM	0012	Alloc	100	04/30/15	16.02.43	0

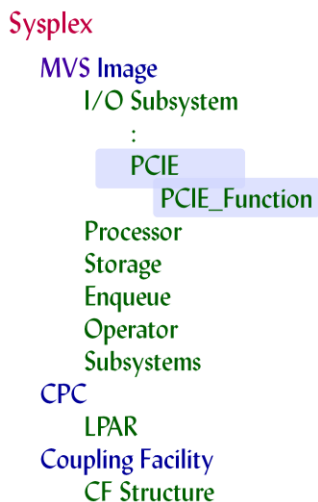
ISPF Fields

Sort Rows

Hyperlink to vertical View

- The statistics of the Monitor III PCIE ISPF Report are available also by means of the Monitor III Data Portal without limitations.
- The PCIE Report can be selected from the report list and basically all report columns can be displayed in the browser window.
- But due to the high number of report columns of you need to use the slider in order to see the rightmost columns as well.
- As an alternative, you can also switch to the vertical report view to see all measurements for a selected device at a glance.


## z/OS V2.2 – Distributed Data Server PCIE Metrics



Resource	New metric
PCIE	<ul style="list-style-type: none"> <li>▶ ... by PCIE function</li> </ul>
PCIE_Function	<ul style="list-style-type: none"> <li>▶ % allocation time</li> <li>▶ % buffer pool utilization</li> <li>▶ % time busy</li> <li>▶ buffer pool memory size</li> <li>▶ compression ratio</li> <li>▶ compression request rate</li> <li>▶ compression throughput</li> <li>▶ decompression ratio</li> <li>▶ decompression request rate</li> <li>▶ decompression throughput</li> <li>▶ received packets rate</li> <li>▶ request execution time</li> <li>▶ request execution time standard deviation</li> <li>▶ request queue time</li> <li>▶ request queue time standard deviation</li> <li>▶ request size</li> <li>▶ transfer rate</li> <li>▶ transfer read rate</li> <li>▶ transfer write rate</li> <li>▶ transmitted packets rate</li> <li>▶ work unit rate</li> <li>▶ PCI adapter utilization</li> <li>▶ PCI load operations rate</li> <li>▶ PCI refresh operations rate</li> <li>▶ PCI store block operations rate</li> <li>▶ PCI store operations rate</li> </ul>

- RMF Distributed Data Server (DDS) supports new resource types PCIE and PCIE Function to allow performance analysis of RoCE devices and zEDC hardware accelerators by DDS API exploiters
- RMF DDS Resource Model represents a composition of resources that can exist in a Parallel Sysplex environment
  - MVS\_IMAGE is a child resource of resource SYSPLEX
  - I/O\_SUBSYSTEM is a child resource of resource MVS\_IMAGE
  - Child resource PCIE is added to resource I/O\_SUBSYSTEM
  - Child resource PCIE\_FUNCTION is added to resource PCIE
- A variety of metric values that are related to resource types PCIE and PCIE\_FUNCTION can be requested from the RMF Distributed Data Server (DDS)
- The metrics are available as single valued metrics for the PCIE\_FUNCTION resource and as list valued metrics for the PCIE resource.

## z13 - RMF PCIE Enhancements



- z13 introduces new PCIE performance measurements for RDMA-over-converged-ethernet (RoCE Express) and zEnterprise data compression (zEDC) devices
  - Existing DMA read/write measurements are replaced by new PCIE function type specific measurements
  - For RoCE Express devices, there are four new measurements
    - **Received-Bytes** No. of bytes received on the external ethernet interface
    - **Transmitted-Bytes** No. of bytes transmitted on the external ethernet interface
    - **Received-Packets** No. of packets received on the external ethernet interface
    - **Transmitted-Packets** No. of packets transmitted on the external ethernet interface
  - For zEDC devices, there are two new measurements
    - **Consumed-Work-Units** No. of work units processed by the PCI function
    - **Maximum Work Units** Maximum no. of work units that the PCI function is capable of processing per second
- With zEC12 / zBC12, the existing DMA Read/Write metrics are still maintained
  - **DMA Reads** No. bytes transferred from DMA address spaces to PCIE function
  - **DMA Writes** No. bytes transferred from PCIE function to DMA address spaces

- RMF PCIE monitoring and reporting functionality is enhanced to support new measurements for PCIe-attached RoCE and zEDC devices configured on z13.

# z13 - RMF Postprocessor PCIE Activity Report

**RMF Postprocessor Interval Report [System Z2] : PCIE Activity Report**

RMF Version : z/OS V2R1 SMF Data : z/OS V2R1  
 Start : 08/13/2013-05:45:00 End : 08/13/2013-06:00:01 Interval : 15:00:00 minutes

**General PCIe Activity**

Function ID	Function PCHID	Function Name	Function Type	Function Status	Owner Job Name	Owner Address Space ID	Function Allocation Time	PCI Lead Operations Rate	PCI Store Operations Rate	PCI Store Block Operations Rate	Read Transfer Rate	Write Transfer Rate	Packets Received Rate	Packets Transmitted Rate	Work Units Processed Rate	Adapter Utilization
0001	0380	Hardware Accelerator	1014044B	Allocated	FPQHWAM	0014	900	0	0.091	0	2.91	1	62.0	0		
0011	05C4	Hardware Accelerator	1014044B	Allocated	FPQHWAM	0014	900	0	0.091	0	2.92	1	14.3	19.1		
0020	038C	10GbE RoCE	15B31003	Allocated	VTAM	00DE	900	0.889	0	0	0	1	42.9	0		

**Hardware Accelerator Activity**

Function ID	Time Busy %	Request Execution Time	Std Dev for Request Execution Time	Request Queue Time	Std Dev for Request Queue Time	Request Size	Transfer Rate	Total
0001	0.005	31.4	4.88			545	68.0	75.2
0011	0.004	30.7	5.35			541	93.3	74.4

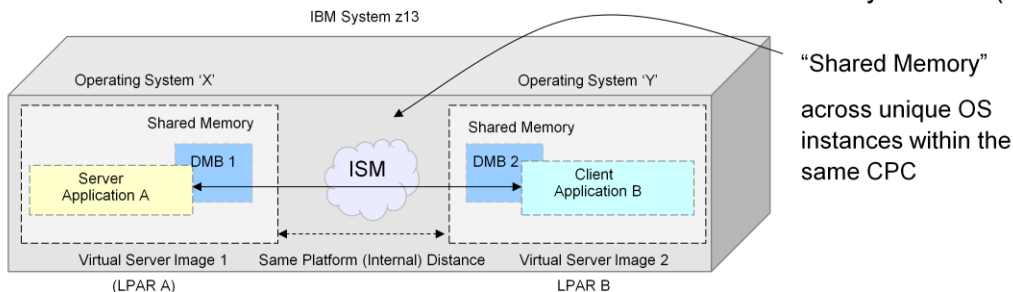
**Hardware Accelerator Compression Activity**

Function ID	Compression Request Rate	Compression Throughput	Compression Ratio	Decompression Request Rate	Decompression Throughput	Decompression Ratio	Buffer Pool Size	Buffer Pool Utilization
0001	1.46	0.088	4.11	0	0		64	0
0011	1.46	0.087	3.94	0	0		64	0

- Due to this architectural change with z13, PCIe Read and Write statistics can now be obtained independent of the DMA address space.
- The PCIe related thruput is now reported in terms of Byte Transfer Rates and additionally as Packet Rates.
- Since on z13 the actual number of processed work units as well as the theoretical maximum number of work units can be retrieved, the PCIe adapter utilization can be calculated and reported as well.

## z13/z13s - RMF PCIE Enhancements

- Shared Memory Communication via Remote Direct Memory Access (SMC-RDMA or SMC-R) is a zEC12 and z13 feature that provides high performance CPC to CPC communication
- SMC-R uses RDMA enabled RoCE PCIE Functions and serves as accelerator for OSA traffic
- z13 GA2 and z13s introduces a new variation of the z PCIe architecture called virtual PCI express (vPCIe) for Internal Shared Memory (ISM).
- Based on vPCIe for ISM a new SMC solution is available: SMC-Direct Memory Access (SMC-D)



SMC-D (over ISM) is very similar to SMC-R (over RoCE). SMC-D extends the benefits of SMC-R to same CPC operating system instances without requiring physical resources (RoCE adapters, PCI bandwidth, ports, I/O slots, network resources, 10GbE switches etc.).

- Shared Memory Communications over Remote Direct Memory Access (SMC-RDMA or SMC-R) is a protocol that allows TCP sockets applications to transparently exploit RDMA (RoCE).
- SMC-R actually offers the benefits of HiperSockets across processor boundaries. It takes advantage of high speed protocols and direct memory placement of data.
- With z13 GA2 and z13s a new SMC solution is available: SMC-Direct Memory Access (SMC-D) over Internal Shared Memory (ISM)
- SMC-D (over ISM) is very similar to SMC-R (over RoCE). SMC-D extends the benefits of SMC-R to system instances running on the same CPC without requiring physical resources.
- ISM is defined as a PCIe device. Function ID(s) / Virtual Function ID(s) must be defined in HCD (or IOCDs)
- An ISM PCIe function must be associated with a channel, either:
  - IQD (a single IQD / HiperSocket) channel or...
  - OSD channels
- The association of ISM Function ID(s) to the channel(s) is created by defining (HCD) matching Physical Network IDs (PNet IDs)
- PNet IDs are dynamically discovered by the Operating System
- The channel devices (OSD or IQD) provide IP connectivity



# z13/z13s - RMF Postprocessor PCIE Activity Report



## RMF Postprocessor Interval Report [System S09] : PCIE Activity Report

RMF Version : z/OS V2R2 SMF Data : z/OS V2R2  
 Start : 01/26/2016-16.30.00 End : 01/26/2016-17.00.00 Interval : 30.00.000 minutes

### General PCIE Activity

Function ID	Function CHID	Function Name	Function Status	Owner Job Name	Owner Address Space	Function Allocation Time	PCI Load Operations Rate	PCI Store Operations Rate	PCI Store Block Operations Rate	Refresh PCI Translations Operations Rate	DMA Address Space Count	Read Transfer Rate	Write Transfer Rate	Packets Received Rate	Packets Transmitted Rate
0249	0204	10GbE RoCE Express	Allocated	VTAM	002B	1800	0.107	0	0	0.871	1	0	0	0	0
0342	027C	10GbE RoCE Express	Allocated	VTAM	002B	1800	0.114	16123	0	0.871	1	200	212	295486	294741
1501	07D0	ISM	Allocated	VTAM	002B	1800	0	0	13907	0	1		43.9		

Physical or virtual channel identifier for the PCIE function.

SMC-D

Function ID	Function CHID	Packets Transmitted Rate	Work Units Processed Rate	Adapter Utilization	Physical Network ID Port 1	Physical Network ID Port 2
0249	0204				NETWORK1	NETWORK1
0342	027C	294741			NETWORK7D0	NETWORK7D1
1501	07D0				NETWORK7D0	

Physical-network identifiers (PNET IDs) of an Ethernet network that is accessible from the ports of the channel path

RoCE: Up to 2 PNET IDs  
 ISM: 1 PNET ID

- With new function APAR OA49113 RMF supports SMC-D on z13 GA2 and z13s. The support is available for z/OS 2.2 and enhances the following RMF reports:
  - Postprocessor PCIE Activity report
  - Postprocessor Channel Path Activity report
  - Monitor III PCIE Activity (PCIE) report
  - Monitor III Channel Path Activity (CHANNEL) report
- The following SMF records are extended:
  - SMF 73 Channel Path Activity
  - SMF 74 subtype 9: PCIE Function Data section and PCIE Function Type Data section.
- Existing SMF 74-9 based overview condition PCIBYTT is changed to support SMC-D.
- RMF Postprocessor PCIE Activity report:
  - PCIE activity data for SMC-D over Internal Shared Memory (ISM) virtual PCIe functions are now reported as new PCIE function name 'ISM'
  - Report field 'Function PCHID' is renamed to 'Function CHID' since it can now report a physical channel identifier or in case of SMC-D a virtual channel identifier.
  - Report field 'Function Type' is removed
  - Two new report fields 'PHYSICAL NETWORK ID PORT 1' and 'PHYSICAL NETWORK ID PORT 2' are added to report the Physical Network IDs (PNet IDs) :
    - For SMC-R (over RoCE) PCIE functions there can be up to two PNet IDs.
    - For SMC-D (over ISM) PCIE functions there can be only one PNet ID

# z13/z13s - RMF Postprocessor Channel Activity Report

CHANNEL PATH ACTIVITY															
z/OS V2R2				SYSTEM ID S09				DATE 01/26/2016				INTERVAL 30.00.375			
				RPT VERSION V2R2 RMF				TIME 16.30.00				CYCLE 1.000 SECONDS			
IODF = 98		CR-DATE: 01/25/2016		CR-TIME: 14.33.03		ACT: ACTIVATE		MODE: LPAR		CPMF: EXTENDED MODE					
CHANNEL PATH		UTILIZATION(%)				READ(MB/SEC)		WRITE(MB/SEC)		PHYSICAL NETWORK IDS					
ID	TYPE	G	SPEED	SHR	PART	TOTAL	PART	TOTAL	PART	TOTAL	PORT 1	PORT 2			
21	OSD		10G	Y	0.00	0.00	0.00	7.32	14.63	7.21	14.40	NETWORK7D0			
23	OSD		10G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	NETWORK1			
27	OSD		1G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	NETWORK1	NETWORK1		
28	OSD		1G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	NETWORK1	NETWORK1		
29	OSD		1G	Y	0.00	25.58	1.11	0.00	67.43	0.00	66.55	NETWORK7CD	NETWORK7CD		
CHANNEL PATH		WRITE(B/SEC)				MESSAGE RATE		MESSAGE SIZE		SEND FAIL		RECEIVE FAIL		PHYSICAL NETWORK ID	
ID	TYPE	G	SHR	PART	TOTAL	PART	TOTAL	PART	TOTAL	PART	TOTAL	PART	TOTAL	NETWORK ID	
F4	IQD			0	0	0	0	-----	-----	0	0	0	0	NETWORK7DD	
F6	IQD			0	0	0	0	-----	-----	0	0	0	0	NETWORK7D1	

Physical network identifier (PNET ID) of the Ethernet network(s) that are accessible from the port(s) of the channel path. Up to two PNET IDs for OSD channels

Physical network identifier (PNET ID) of an Ethernet network that is accessible from the channel path. Only one PNET ID for IQD channel

**RMF Postprocessor Channel Activity report:**  
 The RMF Postprocessor Channel Activity report shows the new field 'PHYSICAL NETWORK ID(S)'. Here you see the PNET ID of the Ethernet network that is accessible from the port(s) of an OSD or IQD channel. An OSD channel can have up to two PNET IDs, reported in columns 'PORT 1' and 'PORT 2'. There can be only one PNET ID for an IQD channel.

- RMF Monitor III enhancements:**
- The Monitor III PCIE Activity report shows measurement data for SMC-Direct over Internal Shared Memory (ISM) virtual PCIe functions. The PNET IDs are added as hidden fields of the PCIE report table. You can display this information if you use the full Monitor III PCIE report in the RMF Data Portal for z/OS or use the RMF Utility to customize the Monitor III PCIE ISPF report.
  - In the RMF Monitor III Channel Activity report, the new PNET ID information for OSD and IQD channels is now available as hidden fields of the Channel report table.

# RMF Enhancements at a Glance



- IBM z13 and z13s Support
  - RMF Statistics for Simultaneous Multithreading (SMT)
  - Extended ICSF Measurements for Crypto Express5S
  - Monitoring of new capping types



- SCM I/O Adapter Performance Reporting
  - Statistics for EADM Subchannel Activities
  - Performance and Thruput on SCM Card Level

- PCIE Activity Reporting

- z/OS V2.2 Monitor III PCIE Report
- Additional Measurements for PCIe attached RoCE and zEDC Devices on z13
- RMF support for SMC-D over ISM on z13 and z13s



- z/OS V2.2 Monitor III Job Resource Consumption Reporting
  - CPU, I/O and Storage Consumption Data at a Glance
  - Detailed Statistics for Job related GQSCAN Activities

- z/OS V2.2 zFS Reporting Enhancements

- Sysplex-wide Statistics for zFS Usage and Performance
- Improved zFS Data Gathering Performance
- Additional zFS Statistics for Shared File System environments



- DS8000 Support

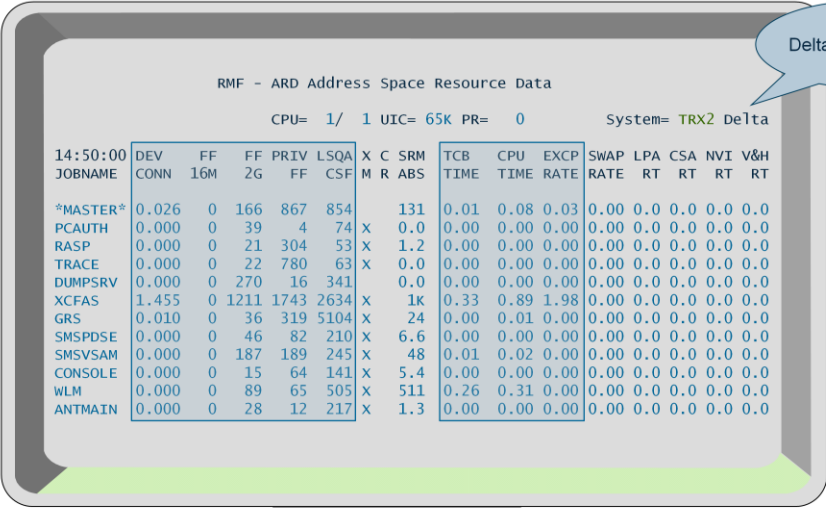
- RMF Monitoring of SuperPAV capability

- Mobile Workload Reporting

- Statistics about total and mobile CPU consumption



## Job Resource Consumption Reporting: Monitor II ARD Report



RMF - ARD Address Space Resource Data

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

14:50:00 JOBNAME	DEV CONN	FF 16M	FF 2G	PRIV FF	LSQA CSF	X M	C R	SRM ABS	TCB TIME	CPU TIME	EXCP RATE	SWAP RATE	LPA RT	CSA RT	NVI RT	V&H RT
*MASTER*	0.026	0	166	867	854			131	0.01	0.08	0.03	0.00	0.0	0.0	0.0	0.0
PCAUTH	0.000	0	39	4	74	x		0.0	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0
RASP	0.000	0	21	304	53	x		1.2	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0
TRACE	0.000	0	22	780	63	x		0.0	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0
DUMPSRV	0.000	0	270	16	341			0.0	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0
XCFAS	1.455	0	1211	1743	2634	x		1K	0.33	0.89	1.98	0.00	0.0	0.0	0.0	0.0
GRS	0.010	0	36	319	5104	x		24	0.00	0.01	0.00	0.00	0.0	0.0	0.0	0.0
SMSPDSE	0.000	0	46	82	210	x		6.6	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0
SMSVSAM	0.000	0	187	189	245	x		48	0.01	0.02	0.00	0.00	0.0	0.0	0.0	0.0
CONSOLE	0.000	0	15	64	141	x		5.4	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0
WLM	0.000	0	89	65	505	x		511	0.26	0.31	0.00	0.00	0.0	0.0	0.0	0.0
ANTMAIN	0.000	0	28	12	217	x		1.3	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0

Delta or Total Mode

RMF Monitor II provides job related resource consumption by means of the following reports

- Address Space Resource Data
- Address Space State Data
- Address Space SRM Data

The data is reported on snapshot base for a certain interval (Delta mode) or it is accumulated for the entire life cycle of the address space (Total mode)

However, there is no comprehensive report for job related resource consumption, the data is spread on three different reports.

For this reason, with z/OS V2R2 RMF the new Monitor III Job Usage Report is introduced.

The ARD report gives information on the system resources that are used by an address space. The information provided in this report includes, for example, information on processor time, paging, and central storage.

The ARD report enables you to determine which jobs are creating performance problems.

From the Monitor II ARD Report, the following statistics are now also available in Monitor III

- Device Connect Time
- Fixed Frame Counts on Virtual Storage location level
- TCB Times and Total CPU Times
- EXCP Rates

## Job Resource Consumption Reporting: Monitor II ASD Report

RMF - ASD Address Space State Data

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

14:50:00

JOBNAME	SRVCLASS	P	L	S	C	R	DP	CS	CS	TAR	X	PIN	TX	SWAP	WSM
							PR	F	TAR	WSS	M	RT	SC	RV	RV
*MASTER*	SYSTEM	1	NS	FF	6353					0		0	0	0	
PCAUTH	SYSTEM	1	NS	FF	129					0	X	0	0	0	
RASP	SYSTEM	1	NS	FF	363					0	X	0	0	0	
TRACE	SYSTEM	1	NS	FF	872					0	X	0	0	0	
DUMPSRV	SYSTEM	1	NS	FF	583					0		0	0	0	
XCFAS	SYSTEM	1	NS	FF	7814					0	X	0	0	0	
GRS	SYSTEM	1	NS	FF	14.1K					0	X	0	0	0	
SMSPDSE	SYSTEM	1	NS	FF	4320					0	X	0	0	0	
SMSVSAM	SYSTEM	1	NS	FF	3310					0	X	0	0	0	
CONSOLE	SYSTEM	1	NS	FF	4477					0	X	0	0	0	
WLM	SYSTEM	1	NS	FF	29.0K					0	X	0	0	0	
ANTMAIN	SYSTEM	1	NS	FF	1511					0	X	0	1	998	

The ASD report gives an overview of the current state of an address space. Basically, the report tells you where each address space is and what it is doing.

You can use the ASD report, for example, to determine which jobs are using large amounts of central storage or which jobs are being swapped excessively and why the swapping is occurring.

If you have a workload delaying your application, you can check the workloads dispatching priority (**DP PR**) on the ASD report, and change it if necessary.

From the Monitor II ASD Report, the following statistics are now also available in Monitor III

- Service Class and Service Class Period
- Dispatching Priority
- Central Storage Frame Counts

# Job Resource Consumption Reporting: Monitor II ASRM Report

RMF - ASRM Address Space SRM Data

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

14:33:58

JOBNAME	SRVCLASS	S P	TRANS ACTIVE	TRANS CUR RES	TX CT	TX SC	TX CPU	TX MSO	TX IOC	TX SRB	SESS TOTAL
*MASTER*	SYSTEM	1	31:10	31:10:51	1	0	5183K	0	51347	17549K	22784K
PCAUTH	SYSTEM	1	31:10	31:10:54	1	0	42	0	5	158	205
RASP	SYSTEM	1	31:10	31:10:54	1	0	20855	0	0	172162	193017
TRACE	SYSTEM	1	31:10	31:10:54	1	0	374	0	80	357	811
DUMPSRV	SYSTEM	1	31:10	31:10:36	1	0	3296	0	2910	1711	7917
XCFAS	SYSTEM	1	31:10	31:10:41	1	0	87313K	0	1181K	142.2M	230.7M
GRS	SYSTEM	1	31:10	31:10:54	1	0	17289K	0	180	3351K	20640K
SMSPDSE	SYSTEM	1	31:10	31:10:54	1	0	1111K	0	0	172760	1283K
SMSVSAM	SYSTEM	1	31:10	31:10:41	1	0	3157K	0	2945	4832K	7992K
CONSOLE	SYSTEM	1	31:10	31:10:54	1	0	3670K	0	2815	897165	4570K
WLM	SYSTEM	1	31:10	31:10:41	1	0	70296K	0	910	13016K	83312K
ANTMAIN	SYSTEM	1	31:10	31:10:41	1	1	217423	0	3360	49190	269973
ANTAS000	SYSSTC	1	31:10	31:10:36	1	1	21204	0	2945	2657	26806

The ASRM report gives an overview of the system resources that are used by an address space.

The report gives, for example, information on processor service, storage service, and I/O service.

The report enables you to determine which jobs are using which services and whether certain jobs are creating performance problems by making excessive use of system services.

From the Monitor II ASRM Report, the following statistics are now also available in Monitor III

- Service Class and Service Class Period
- Transaction Active Times
- Transaction Current Resident Times
- Transaction Counts

## z/OS V2.2 – Monitor III USAGE Report

- Request the Monitor III USAGE Report by selection **1** from the Primary Menu & selection **4A** from the Overview Report Selection Menu or enter command:  
*USAGE [job\_class, service\_class] or USG*

Identify Top Resource Consumers at a Glance

RMF V2R2 Job Oriented Usage

Samples: 60 System: TRX1 Date: 04/18/15 Time: 10.56.00 Range: 60 Sec

Jobname	Service CX Class	--- I/O ---		--- CPU ---		- Storage -		----- QScan -----		
		Conn	EXCP	Total	TCB	Total	Fixed	Total	Resct	Time
XCFAS	S SYSTEM	0.446	1.97	0.25	0.11	7754	2384	0	0.0	0
BHBE	T TSOCLASS	0.103	3.90	0.07	0.07	21	623	1	0.0	2581
*MASTER*	S SYSTEM	0.042	0.00	0.02	0.00	6323	1107	0	0.0	0
SMF	S SYSTEM	0.028	0.00	0.00	0.00	900	210	0	0.0	0
CATALOG	S SYSTEM	0.027	0.17	0.03	0.03	1824	228	0	0.0	0
GRS	S SYSTEM	0.020	0.00	0.01	0.01	14136	451	0	0.0	0
JES2	S SYSSTC	0.010	0.38	0.03	0.02	9277	1041	0	0.0	0
NET	S SYSSTC	0.010	0.00	0.01	0.00	3050	138	0	0.0	0
DFSZFS	S SYSSTC	0.008	0.60	0.00	0.00	30660	499	0	0.0	0
OMVS	S SYSTEM	0.006	0.17	0.00	0.00	16098	356	0	0.0	0
SMS	S SYSSTC	0.004	0.93	0.00	0.00	548	89	0	0.0	0
PAGENT	SO SYSSTC	0.003	9.45	0.01	0.01	2978	18072	0	0.0	0
HZSPROC	SO SYSSTC	0.000	0.00	0.00	0.00	5125	183	0	0.0	0

46

© 2016 IBM Corporation

The new Monitor III Job Usage Report provides a comprehensive overview about job related resource consumption. The report is arranged accordingly to the primary resource categories:

- I/O
- CPU
- Storage

By default the report is sorted by descending I/O Connect Time

In the ISPF Version of the report the counts in the visible columns are related to the current Monitor III Range.

Furthermore the total accumulated counts (as applicable by means of the Monitor II Total Mode) are also stored as hidden fields in the Monitor III ISPF table.

In addition to the general resource consumption statistics, the new report displays statistics about job related QSCAN activities:

From the API perspective, QSCAN requests can be initiated by GQSCAN macro or ISGQUERY REQINFO=QSCAN. (ISGQUERY REQINFO=QSCAN is the IBM recommended replacement of GQSCAN)

The following counts are reported:

- QScan Total: Total requests is the accumulated number of QScan requests for the address space
- QScan Resct: Resct is the average number of resources returned by QScan requests for the address space
- Qscan Time: Time is the average QScan request time in microseconds for the address space

# z/OS V2.2 – Monitor III USAGE Report: Cursor Sensitivity Targets

RMF V2R2 Job Oriented Usage

Samples: 60    System: TRX1    Date: 04/18/15    Time: 10.56.00    Range: 60 Sec

Jobname	Service CX Class	--- I/O ---		--- CPU ---		- Storage -		----- QScan -----		
		Conn	EXCP	Total	TCB	Total	Fixed	Total	Resct	Time
XCFAS	S SYSTEM	0.446	1.97	0.25	0.11	7754	2384	0	0.0	0
*MASTER*	S SYSTEM	0.042	0.00	0.02	0.00	6323	1107	0	0.0	0
SMF	S SYSTEM	0.028	0.00	0.00	0.00	900	210	0	0.0	0
CATALOG	S SYSTEM	0.027	0.17	0.03	0.03	1824	228	0	0.0	0
GRS	S SYSTEM	0.020	0.00	0.01	0.01	14136	451	0	0.0	0
JES2	S SYSSTC	0.010	0.38	0.03	0.02	9277	1041	0	0.0	0
NET	S SYSSTC	0.010	0.00	0.01	0.00	3050	138	0	0.0	0
DFSZFS	S SYSSTC	0.008	0.60	0.00	0.00	30660	499	0	0.0	0
OMVS	S SYSTEM	0.006	0.17	0.00	0.00	16098	356	0	0.0	0
SMS	S SYSSTC	0.004	0.93	0.00	0.00	548	89	0	0.0	0
PAGENT	SO SYSSTC	0.003	9.45	0.01	0.01	2978	18072	0	0.0	0
HZSPROC	SO SYSSTC	0.000	0.00	0.00	0.00	5125	183	0	0.0	0

- In case of general resource contention, the Job Usage report can serve as an excellent starting point for further analysis and problem determination.
- In this context the Monitor III Cursor Sensitivity feature can be used to drill-down problems efficiently.



## z/OS V2.2 – RMF Monitor III USAGE Report

- The set of report options common to Monitor III address space related reports is applicable for the USAGE report
- Use RO command to filter by address space type or WLM service class

```

Command ==>          RMF Delay Report Options: USAGE          Line 1 of 2
                                                                Scroll ==> CSR
Change or verify parameters. To exit press END.
Changes will apply to DELAY, DEV, ENQ, HSM, JES, PROC, PROCU, STOR, STORC,
STORF, STORM, USAGE and XCF.

Class      ==> ALL      Classes: ALL TSO BATCH STC ASCH OMVS
Service class ==> *ALL  *ALL or one of available service classes below

Jobs       ==> NO      View job selection/exclusion panel next (YES NO)

Available Service classes
BTCHDEF   GPMSERVE   OE           OMVSKERN   STCDEF     TSODEF     SYSTEM
SYSS1C

Command ==>          RMF Delay Report Options: USAGE          Line 1 of 18
                                                                Scroll ==> CSR
Select (S), exclude (X), or fill-in jobs for report. Press END.

Sel  Jobname  Sel  Jobname  Sel  Jobname  Sel  Jobname  Sel  Jobname
-    -        -    -        -    -        -    -        -    -
S    *ALL      *MASTER*  ALLOCAS  ANTAS000  ANTMAIN
     APPC      ASCH      ASCHINT  AUXMON    AXR
     AXR04     BJAG      BKGE     BKGEDDSF  BLUCIUS
     BMAI      BPMU      BPXOINIT CATALOG   CAZO
     CBDQDISP CEA       CFZCIM   CONSOLE   CSF
  
```

The set of common report options for all Monitor III address space related reports will be applicable for the new USAGE report as well.

The RO command displays the Report Options panel which allows you to filter the report by address space type or WLM service class.

Once you specify **YES** in the *Jobs* input field, the Job Selection/Exclusion panel is displayed in order to filter the report by specific job names.

In addition, you can invoke the report directly by specifying an address space type or service class name as parameter (e.g. *USAGE S* or *USAGE SYSS1C*).

# z/OS V2.2 – RMF Data Portal Job Usage Report

## RMF Report - One Row [TRX2,MVS\_IMAGE] : USAGE

Time Range: 05/04/2015 19:15:00 - 05/04/2015 19:16:00

Jobname	DFSZFS
ASID (dec)	0036
Job Class	S
Job Class Ext	S
Service Class	SYSSTC
Period	1
Dispatching Priority	FE
Transaction Active Time	99:05:20
Transaction Resident Time	99:05:20
Transaction Count	2
Total Frames	31083
Fixed Frames	501
Fixed Frames High	470
Fixed Frames Above	31
Fixed Frames Below	0
Total Device Connect Time	27.20
Device Connect Time	0.006
EXCP Rate	0.60
Total CPU Time	19.44
CPU Time	0.00
Total TCB Time	16.37
TCB Time	0.00
QSCAN Requests	0
Specific QSCAN Requests	0
QSCAN Resource Count	0.0
QSCAN Resource Count Std.Dev	0.0
QSCAN Request Time	
QSCAN Request Time Std.Dev	

[http://hostname:8803/gpm/rmf3.xml?report=USAGE&resource=,sysid,MVS\\_IMAGE](http://hostname:8803/gpm/rmf3.xml?report=USAGE&resource=,sysid,MVS_IMAGE)

## RMF Data Portal for z/OS

Sort Rows

RMF Report [TRX2,MVS\_IMAGE] : USAGE (Job Oriented Usage)

Time Range: 05/04/2015 19:15:00 - 05/04/2015 19:16:00

Jobname	ASID (dec)	Job Class	Job Class Ext	Service Class	Period	Dispatching Priority	Transaction Active Time	Transaction Resident Time
*MASTER*	0001	S	S	SYSTEM	1	FF	99:24:20	99:24:20
XCFAS	0006	S	S	SYSTEM	1	FF	99:05:20	99:05:20
JES2	0052	S	S	SYSSTC	1	FE	99:05:20	99:05:20
CATALOG	0047	S	S	SYSTEM	1	FF	99:05:20	99:05:20
RMFGAT	0065	S	SO	SYSSTC	1	FE	99:05:20	99:05:20
GRS	0007	S	S	SYSTEM	1	FF	99:05:20	99:05:20
SMF	0030	S	S	SYSTEM	1	FF	99:05:20	99:05:20
SMS	0024	S	S	SYSSTC	1	FE	99:05:20	99:05:20
DFSZFS	0036	S	S	SYSSTC	1	FE	99:05:20	99:05:20
NET	0063	S	S	SYSSTC	1	FE	99:05:20	99:05:20
OMVS	0016	S	S	SYSTEM	1	FF	99:05:20	99:05:20
CONSOLE		S	S	SYSTEM	1	FF	99:05:20	99:05:20

Hyperlink to vertical View

ISPF Fields

The new Monitor III USAGE report will be added to the standard reports which are available in XML format by means of the RMF Distributed Data Server.

Therefore all statistics of the Monitor III Job Usage ISPF Report are available also by means of the Monitor III Data Portal without limitations.

The USAGE Report can be selected from the report list and basically all report columns can be displayed in the browser window, also by using the following URL:

[http://hostname:8803/gpm/rmf3.xml?report=USAGE&resource=,sysname,MVS\\_IMAGE](http://hostname:8803/gpm/rmf3.xml?report=USAGE&resource=,sysname,MVS_IMAGE)

Sorting is possible in the RMF Data Portal: Just one click on the column header brings the job with the highest consumption count to the top.

Due to the high number of report columns you need to use the slider in order to see the rightmost columns as well. So, as an alternative, you can also switch to the vertical report view to see all measurements for a selected device at a glance.

Clicking on a jobname will display all the values for this job in one single view.

# z/OS V2.2 – RMF USAGE Report: RMF Distributed Data Server Metrics

Sysplex

- MVS Image
- I/O Subsystem
- Processor
- Storage
- Enqueue
- Operator
- Subsystems
- CPC
- LPAR
- Coupling Facility
- CF Structure

Resource	New metric
MVS_IMAGE	<ul style="list-style-type: none"> <li>▶ transaction active time by job</li> <li>▶ transaction resident time by job</li> <li>▶ # transactions by job</li> <li>▶ # QSCAN generic requests by job</li> <li>▶ # QSCAN specific requests by job</li> <li>▶ # QSCAN resources by job</li> <li>▶ # QSCAN resources standard deviation by job</li> <li>▶ QSCAN request time by job</li> <li>▶ QSCAN request time standard deviation by job</li> </ul>
STORAGE	<ul style="list-style-type: none"> <li>▶ # frames fixed below 16 MB by job</li> <li>▶ # frames fixed above 16 MB by job</li> <li>▶ # frames fixed above 2 GB by job</li> </ul>
I/O_SUBSYSTEM	<ul style="list-style-type: none"> <li>▶ total device connect time by job</li> <li>▶ connect time by job</li> <li>▶ EXCP rate by job</li> </ul>
PROCESSOR	<ul style="list-style-type: none"> <li>▶ total CPU time by job</li> <li>▶ CPU time by job</li> <li>▶ total TCB time by job</li> <li>▶ TCB time by job</li> </ul>

DDS metrics extracted from the Job USAGE Report

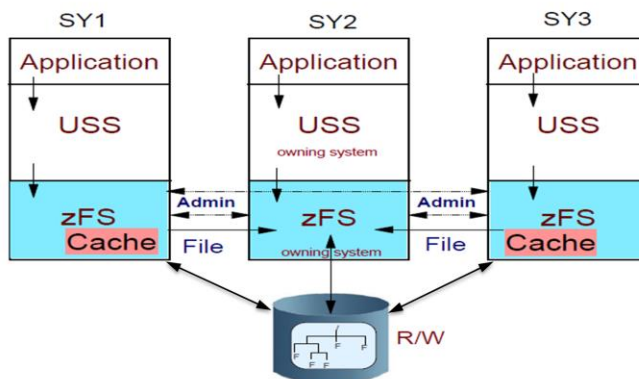
- In addition a subset of the metrics contained in the USAGE report were added to the job related DDS metrics. Hence, these metrics are supported by the RMF Performance Data Portal and the z/OSMF Resource Monitoring GUI.
- Since the Job Usage Report contains comprehensive address space related statistics, the new DDS metrics which are extracted from this report are spread across multiple resources.

# RMF Enhancements at a Glance

- **IBM z13 and z13s Support**
  - RMF Statistics for Simultaneous Multithreading (SMT)
  - Extended ICSF Measurements for Crypto Express5S
  - Monitoring of new capping types
- **SCM I/O Adapter Performance Reporting**
  - Statistics for EADM Subchannel Activities
  - Performance and Thruput on SCM Card Level
- **PCIE Activity Reporting**
  - z/OS V2.2 Monitor III PCIE Report
  - Additional Measurements for PCIe attached RoCE and zEDC Devices on z13
  - RMF support for SMC-D over ISM on z13 and z13s
- **z/OS V2.2 Monitor III Job Resource Consumption Reporting**
  - CPU, I/O and Storage Consumption Data at a Glance
  - Detailed Statistics for Job related GQSCAN Activities
- **z/OS V2.2 zFS Reporting Enhancements**
  - Sysplex-wide Statistics for zFS Usage and Performance
  - Improved zFS Data Gathering Performance
  - Additional zFS Statistics for Shared File System environments
- **DS8000 Support**
  - RMF Monitoring of SuperPAV capability
- **Mobile Workload Reporting**
  - Statistics about total and mobile CPU consumption



## z/OS 2.2 - zFS Reporting Enhancements



### Rationale:

- Existing RMF Monitor III single system reports ZFSSUM and ZFSACT offered no possibility to monitor details of zFS related to sysplex awareness of zFS file systems.
- Some customers observed performance problems when gathering zFS performance data.

The existing RMF Monitor III zFS reports (introduced with z/OS 1.7) has some disadvantages:

- Since the reports are single system reports, it is not possible to monitor zFS activity in shared file system environments:
  - There is no sysplex wide view of zFS activity
  - There are no zFS statistics related to sysplex awareness of zFS file systems
- The current Monitor III data gatherer uses multiple zFS interface calls to collect the zFS data. In case of system environments with a larger number of zFS aggregates and file systems, performance problems were observed that caused delays in overall RMF Monitor III data gathering

RMF for z/OS 2.2 introduces new Monitor III zFS sysplex reports that allow to monitor shared file system environments effectively. The Monitor III zFS data gatherer uses a new zFS interface that reduces the number of zFS interface calls so that RMF zFS data gathering performance is improved.

## z/OS 2.2 - zFS Reporting Enhancements

- z/OS V2.2 introduces three new Monitor III zFS reports:
  - zFS Overview report
  - zFS File System report
  - zFS Kernel report
- Sysplex-wide data on:
  - zFS response time / wait times
  - zFS XCF activity
  - zFS Cache activity
  - zFS activity / capacity by File System
- Data helps to control the zFS environment according to
  - Monitoring of zFS activity in Shared File System environments
  - Tuning of Cache sizes
  - Monitoring of File System Performance:
    - Discover Bottlenecks and
    - Balance File System I/O
  - Capacity Control for File Systems

```

RMF Sysplex Report Selection Menu
Selection =====>
Enter selection number or command for desired report.

Sysplex Reports
1 SYSSUM  Sysplex performance summary          (SUM)
2 SYSRTD  Response time distribution          (RTD)
3 SYSWKM  Work Manager delays                (WKM)
4 SYSEQ   Sysplex-wide Enqueue delays       (ES)
5 CFOVER  Coupling Facility overview        (CO)
6 CFSYS   Coupling Facility systems         (CS)
7 CFACT   Coupling Facility activity        (CA)
8 CACHSUM Cache summary                     (CAS)
9 CACHDET Cache detail                      (CAD)
10 RLSSC  VSAM RLS activity by storage class (RLS)
11 RLSDS  VSAM RLS activity by data set     (RLD)
12 RLRLRU VSAM LRU overview                 (RL)
13 ZFSOVW zFS Overview                      (ZFO)
14 ZFSFSS zFS File System                   (ZFF)
15 ZFSKN  zFS Kernel                        (ZFK)
  
```

RMF for z/OS V2.2 introduces three new RMF Monitor III zFS reports:

- zFS Overview report
- zFS File System report
- zFS Kernel report

The new reports are sysplex reports.

They can be invoked from the RMF Sysplex Report Selection Menu or by following commands:

- **ZFSOVW** or **ZFO**: zFS Overview report
- **ZFSFSS** or **ZFF**: zFS File System report
- **ZFSKN** or **ZFK**: zFS Kernel report

Gathering of zFS activity data is controlled by existing Monitor III gatherer option:

**NOZFS** | **ZFS**

With z/OS V2.2 the zFS data gathering default is changed back from NOZFS (z/OS 2.1) to ZFS.

## z/OS 2.2 – zFS Overview report

- Monitor III zFS Overview report (ZFSOVW) provides a summary of zFS activity, wait percentages and cache statistics on the current sysplex.
- Helps to discover bottlenecks and tune cache behaviour.

```

RMF V2R2   zFS Overview   - UTCPLXCB   Line 1 of 3

Samples: 60   Systems: 3   Date: 04/21/15   Time: 05.02.00   Range: 60   Sec

System      -----Wait%-----      -----Cache Activity-----
            I/O  Lock Sleep      ---User---      --Vnode---      -Metadata-
            Rate Hit%      Rate Hit%      Rate Hit%

CB8C        21.3  0.0  36.2      40.50  100      153.2  99.2      41.22  99.1
CB8D         0.0  0.0   9.2       3.500  100      16.05  100       19.80  99.8
CB86         6.8  0.0   0.1       5775  84.2     17238  87.7     59821  93.5
  
```

### Fields in Monitor III zFS Overview report (ZFSOVW) :

System	Name of the system running zFS.
Wait%	The following Wait percentages are reported: <b>I/O</b> Percentage of time that zFS requests had to wait for I/O completion. <b>Lock</b> Percentage of time that zFS requests had to wait for locks. <b>Sleep</b> Percentage of time that zFS requests had to wait for events.

Dashes (----) in these fields indicate that RMF is unable to calculate a reasonable value.

### Cache Activity section

User	The user file cache is for caching regular user files that are larger than 7K. The measured statistics have the following meanings: <b>Rate</b> Total number of read and write requests per second made to the user file cache. <b>Hit%</b> Percentage of read and write requests to the user file cache that completed without accessing the DASDs.
Vnode	The vnode cache is used to hold virtual inodes. An inode is a data structure related to a file in the file system, holding information about the file's user and group ownership, access mode and type. The measured statistics have the following meanings: <b>Rate</b> Number of read and write requests per second made to the vnode cache. <b>Hit%</b> Percentage of read and write requests to the vnode cache that completed without accessing the DASDs.
Metadata	The metadata cache is used for file system metadata and for files smaller than 7K. It resides in the primary z/FS address space. The measured statistics have the following meanings: <b>Rate</b> Number of read and write requests per second made to the metadata cache. <b>Hit%</b> Percentage of read and write requests to the metadata cache that completed without accessing the DASDs.

# z/OS 2.2 – zFS Overview report : Cursor Sensitivity Targets

RMF V2R2 zFS Overview - UTCPLXCB Line 1 of 3

Samples: 60 Systems: 3 Date: 04/21/15 Time: 05.02.00 Range: 60 Sec

System	----wait%----			-----User-----		--Vnode--		-Metadata-	
	I/O	Lock	Sleep	Rate	Hit%	Rate	Hit%	Rate	Hit%
CB8C	21.3	0.0	36.2	40.50	100	153.2	99.2	41.22	99.1
CB8D	0.0	0.0	9.2	3.500	100	16.05	100	19.80	99.8
CB8E	6.8	0.0	0.1	5775	84.2	17238	87.7	59821	93.5

User Cache Details

Vnode Cache Details

Metadata Cache Details

zFS Overview - I/O Details by Type

The following details are available for system CB8C  
Press Enter to return to the Report panel.

Count	waits	Cancl	Merge	Type
45	5	0	0	FILE SYSTEM METADATA
7	0	0	0	LOG
50	0	0	0	USER FILE DATA

Request breakdown into three major request types

From the **zFS Overview Report**, you can navigate to a variety of detail information using cursor-sensitive control.

The following details pop-up panels can be displayed:

- **I/O Details by Type** pop-up panel from **Wait% -I/O** field in the **Wait%** section.
- **User Cache Details** pop-up panel from any value in the **Cache Activity – User** section.
- **Vnode Cache Details** pop-up panel from any value in the **Cache Activity – Vnode** section
- **Metadata Cache Details** pop-up panel from any value in the **Cache Activity – Metadata** section

### Fields in Monitor III ZFSOVW report I/O Details pop-up panel:

System	Name of the system running zFS.
Count	Total number of I/O requests of the indicated type.
Waits	Number of zFS requests waiting for an I/O completion of the indicated I/O type.
Cancl	Number of cancelled zFS requests during an I/O request of the indicated type, for example, a user tried to delete a file during a pending I/O to this file's metadata.
Merge	Number of merges of two I/O requests into a single request because of better performance.
Type	Type of the I/O request (I/O for metadata, log data or user file data).



# z/OS 2.2 – zFS Overview Cache Details

**zFS Overview - User Cache Details**

The following details are available for system CB8C

Size : 1385M                      Storage fixed : NO  
 Total Pages : 3007K  
 Free Pages : 2997M  
 Segments : 0

----- Read -----				----- Write -----				Read%	Dly%
Rate	Hit%	Dly%	Async Rate	Rate	Hit%	Dly%	Sched Rate		
27.00	100	0.0	23.73	13.50	100	0.1	0.850	66.7	0.0

----- Misc -----  
 Page Reclaim Writes : 0  
 Fsyncs : 16

Cache statistics for 'regular' user files:

- Request rates, hit ratios, delays
- Storage statistics

Cache statistics for file system metadata (directory contents, file status information ...):

- Request rate, hit ratio
- Storage statistics

**zFS Overview - Vnode Cache Details**

The following details are available for system CB8C

Size : 32768

----- Vnodes -----					
Total	Size	Ext.#	Ext.Size	Open	Held
32768	224	32768	816	128	2119

----- Requests -----				
Total	Rate	Hit%	Alloc	Delete
9189	153.2	99.2	0	44

Cache statistics for file system related data structures (Vnodes):

- Request rate, hit ratio
- Vnode statistics

**zFS Overview - Metadata Cache Details**

The following details are available for system CB8C

Size : 1024M                      Storage fixed : YES  
 Buffers : 131K

----- Requests -----			----- Misc -----	
Total	Rate	Hit%	Updates	Partial Writes
2473	41.22	99.1	407	0

## Fields in Monitor III ZFSOVW report User Cache Details pop-up panel:

- System** Name of the system running zFS.
- Size** Total number of pages in the user file cache.
- Total Pages** Number of zFS requests waiting for an I/O completion of the indicated I/O type.
- Free Pages** Total number of free pages in the user file cache.
- Segments** Total number of allocated segments in the user file cache.
- Storage fixed** Shows whether the size of the user file cache storage is fixed. If the zFS parameter user\_cache\_size is set to 'fixed', then zFS reserves real storage for use by zFS only. The 'fixed' option helps to improve performance during data access and can be applied if you have enough real storage available.
- Read Rate** Number of read requests per second made to the user file cache.
- Read Hit%** Percentage of read requests to the user file cache that completed without accessing the DASD.
- Read Dly%** Percentage of delayed read requests to the user file cache. A read request is delayed if it must wait for pending I/O, for example, because the file is in a pending read state due to asynchronous read ahead from DASD to the user file cache.
- Async Read Rate** Number of readaheads per second.
- Write Rate** Number of write requests per second made to the user file cache.
- Write Hit%** Percentage of write requests to the user file cache that completed without accessing the DASD.
- Write Dly%** Percentage of delayed write requests to the user file cache. The following reasons are counted as write request delays:  
 Write wait: a write must wait for pending I/O.  
 Write faulted: a write to a file needs to perform a read from DASD. If a write-only updates a part of a file's page, and this page is not in the user file cache, then the page must be read from DASD before the new data is written to the cache.
- Scheduled Write Rate** Number of scheduled writes per second.
- Read%** Percentage of read requests, based on the sum of read and write requests.
- Dly%** Percentage of delayed requests.
- Page Reclaim Writes** Total number of page reclaim writes. A page reclaim write action writes one segment of a file from the user file cache to DASD. Page reclaim writes are performed to reclaim space in the user file cache. If page reclaim writes occur too often in relation to the write rate, then the user file cache may be too small.
- Fsyncs** Total number of requests for file synchronization (fsync) between user file cache and DASD.

### Fields in Monitor III ZFSOVW report Vnode Cache Details pop-up panel:

System Name of the system running zFS.

Size Total size of the vnode cache.

#### Vnodes

Total Number of currently allocated vnodes in the vnode cache. If more vnodes are requested than are currently available, then zFS dynamically allocates more vnodes.

Size Size of a vnode data structure in bytes.

Ext.# Number of extended vnodes.

Ext. Size Size of an extended vnode data structure in bytes.

Open Number of currently open vnodes.

Held Number of vnodes currently held in zFS by USS.

#### Requests

Total Number of requests to the vnode cache.

Rate Number of requests per second made to the vnode cache.

Hit% Percentage of requests to the vnode data that found the target vnode data structures in the vnode cache. High hit rates indicate a favorable zFS environment, because each miss involves initialization of vnode data structures in the vnode cache.

Alloc Number of requests to create new vnodes (for operations such as create or mkdir).

Delete Number of requests to delete vnodes (for operations such as remove or failed creates or mkdirs).

### Fields in Monitor III ZFSOVW report Metadata Cache Details pop-up panel:

System Name of the system running zFS.

Size Total size of the metadata cache.

Buffers Total number of buffers in the metadata cache. The buffer size is 8K.

Storage fixed Shows whether the size of the metadata cache storage is fixed. If the zFS parameter meta\_cache\_size is set to 'fixed', then zFS reserves real storage for use by zFS only.

The 'fixed option' helps to improve performance during data access and can be applied if you have enough real memory available

#### Requests

Total Number of requests made to the metadata cache.

Rate Number of requests per second made to the metadata cache.

Hit% Percentage of requests to the metadata cache completing without accessing the DASD.

#### Misc

Updates Number of updates made to buffers in the metadata cache.

Partial writes Number of times that only half of an 8K metadata block needed to be written.

# z/OS 2.2 – zFS File System report

- Monitor III zFS File System report (ZFSFS) provides detailed measurements of zFS file system activity for single file systems in compatibility mode aggregates.
- Helps to monitor performance and capacity limits of file systems.

```

RMF V2R2   zFS File System - UTCPLXCB           Line 1 of 568

Samples: 60   Systems: 3   Date: 04/21/15   Time: 05.02.00   Range: 60   Sec
----- File System Name -----
System      Owner      Mode      Size Usq%   I/O  Resp  Read  XCF
Rate        Time      %         Rate      Time  %     Rate
APIRWW.DB2  *ALL      CB86      RW S    1320M  1.0   0.000 0.000  0.0 0.000
NETVIEW.V6R1M0C.ZFS
*ALL      CB86      RW S     12M 49.8   0.000 0.000  0.0 0.000
OMVS.CB8C.JAVATEST.OUTPUT.ZFS
*ALL      CB8C      RW S    9849M  0.3   0.000 0.000  0.0 0.000
OMVS.CB8C.JAVATEST.ZFS
*ALL      CB8C      RW S    4932M  0.7   0.133 0.005  100 0.000
OMVS.CB8D.JAVATEST.OUTPUT.ZFS
*ALL      CB8D      RW S    9849M  3.8   0.000 0.000  0.0 0.000
OMVS.CB8D.JAVATEST.ZFS
*ALL      CB8D      RW S     13G 79.9   0.067 0.003  100 0.000
OMVS.CB86.JAVATEST.OUTPUT.ZFS
*ALL      CB86      RW S     12G  0.3   0.000 0.000  0.0 0.000
    
```

58

© 2016 IBM Corporation

## Fields in Monitor III zFS File System report (ZFSFS):

File System Name	File system name.
System	Name of the system connected to the file system. In the first data line for a file system, the name is '*ALL' to indicate that this line shows the SYSPLEX view of the data rather than a single system view.
Owner	Name of owning system.
Mode	Mount mode of the file system. Possible values are: <b>RW</b> mounted in read-write mode. <b>RO</b> mounted in read-only mode. <b>NM</b> not mounted. <b>QS</b> not available because the aggregate is quiesced. The mount mode is followed by an <b>S</b> if the file system is using zFS sysplex sharing (RWSHARE).
Size	Maximum logical size of the file system (in Bytes).
Usg%	Percentage of currently used space by the file system.
I/O Rate	The rate of read and write requests per second (directory and file) made by applications to this file system.
Resp Time	Average response time in milliseconds for read and write requests made by applications to this file system.
Read%	Percentage of read operations contained in 'I/O Rate'.
XCF Rate	The rate of read and write XCF calls per second to the server.

## z/OS 2.2 – zFS File System report: Report Options

- The contents of the Monitor III ZFSFS report can be tailored by report options (use RO command on ZFSFS panel).
- Example: Detailed statistics for one file systems

```

RMF ZFS File System Report Options          Line 1 of 284

Change or verify parameters. To exit press END.
Changes will apply to the ZFSFS report.

Name  ==> OMVS.CB8D.JAVATEST.ZFS
      ALL or one of the available zFS file systems below
Detail ==> YES    Show single system data (YES or NO) in ZFSFS report

      Available zFS File Systems
OMVS.CB8C.JAVATEST.OUTPUT.ZFS
OMVS.CB8C.JAVATEST.TESTSUITE.ZFS
OMVS.CB8C.JAVATEST.ZFS
OMVS.CB8D.JAVATEST.OUTPUT.ZFS
OMVS.CB8D.JAVATEST.ZFS
OMVS.CB8E.JAVATEST.OUTPUT.ZFS
OMVS.CB8E.JAVATEST.ZFS

RMF V2R2    zFS File System - UTCPLXCB          Line 1 of 5
Samples: 60   Systems: 3   Date: 04/21/15   Time: 05.02.00   Range: 60   Sec

----- File System Name -----
System      Owner      Mode      Size Usq%   I/O  Resp  Read  XCF
              Rate      Time %      Rate      %      Rate
OMVS.CB8D.JAVATEST.ZFS
*ALL        CB8D       RW S      13G 79.9   0.067 0.003 100 0.000
CB8C        CB8D       RW S              0.0 0.000 0.000 0.0 0.000
CB8D        CB8D       RW S      13G 79.9   0.067 0.003 100 0.000
CB8E        CB8D       RW S              0.0 0.000 0.000 0.0 0.000
  
```

The Report Options panel for the zFS File System report allows you to specify options for this report.

### Fields in Monitor III zFS File System report options panel:

- Name** Specify either ALL or the name of one of the zFS file systems available in the sysplex as shown in the field **Available File Systems**, which provides a list of all zFS file systems that are currently defined to the sysplex. You can use an asterisk (\*) as the last character of the file system name as a wild card. When a wild card is used, all file systems whose names start with the specified character sequence before the asterisk are reported on, no matter which characters follow.
- Detail** Specify the desired level of detail in the zFS File System report:
- NO** The report contains summary data for the sysplex only.
  - YES** The report contains data for the sysplex and all single systems.

If the list of file names is too long to fit on the first page, this report options panel can be scrolled up and down using function keys F7 and F8.

# z/OS 2.2 – zFS File System report: Report Options

- Example:

Using ALL for 'Name' and Details=YES, complete information for all file systems on all systems is provided.

RMF V2R2 zFS File System - UTCPLXCB Line 1 of 1420

Samples: 60 Systems: 3 Date: 04/21/15 Time: 05.02.00 Range: 60 Sec

File System Name	System	Owner	Mode	Size	Usg%	I/O Rate	Resp Time	Read %	XCF Rate
APIRWW.DB2									
	*ALL	CB86	RW S	1320M	1.0	0.000	0.000	0.0	0.000
	CB8C	CB86	RW S		0.0	0.000	0.000	0.0	0.000
	CB8D	CB86	RW S		0.0	0.000	0.000	0.0	0.000
	CB86	CB86	RW S	1320M	1.0	0.000	0.000	0.0	0.000
NETVIEW.V6R1M0C.ZFS									
	*ALL	CB86	RW S	12M	49.8	0.000	0.000	0.0	0.000
	CB8C	CB86	RW S		0.0	0.000	0.000	0.0	0.000
	CB8D	CB86	RW S		0.0	0.000	0.000	0.0	0.000
	CB86	CB86	RW S	12M	49.8	0.000	0.000	0.0	0.000
OMVS.CB8D.JAVATEST.OUTPUT.ZFS									
	*ALL	CB8D	RW S	9849M	3.8	0.000	0.000	0.0	0.000
	CB8C	CB8D	RW S		0.0	0.000	0.000	0.0	0.000
	CB8D	CB8D	RW S	9849M	3.8	0.000	0.000	0.0	0.000
	CB86	CB8D	RW S		0.0	0.000	0.000	0.0	0.000
OMVS.CB8D.JAVATEST.ZFS									
	*ALL	CB8D	RW S	13G	79.9	0.067	0.003	100	0.000
	CB8C	CB8D	RW S		0.0	0.000	0.000	0.0	0.000
	CB8D	CB8D	RW S	13G	79.9	0.067	0.003	100	0.000
	CB86	CB8D	RW S		0.0	0.000	0.000	0.0	0.000

## z/OS 2.2 – zFS File System report: Cursor Sensitivity Targets

```

RMF V2R2      zFS File System - UTCPLXCB      Line 1 of 5
Samples: 60   Systems: 3   Date: 04/21/15   Time: 05.02.00   Range: 60   Sec
----- File System Name -----
System      Owner      Mode      Size Usq%      I/O  Resp Read  XCF
OMVS.CB8D.JAVATEST.ZFS
*ALL        CB8D       RW S      13G 79.9      0.067 0.003 100 0.000
CB8C        CB8D       RW S           0.0 0.000 0.000 0.0 0.000
CB8D        CB8D       RW S      13G 79.9      0.067 0.003 100 0.000
CB8E
    
```



zFS File System Details

```

File System Name : OMVS.CB8D.JAVATEST.ZFS
Mount
Point : /CB8D/javatest

System : CB8D           Owner : CB8D           Mode : RW S

----- Read -----
--- Appl --- XCF --- Aggr
Rate  Resp  Rate  Resp  Rate
Time  Time
0.067 0.003 0.000 0.000 0.000

----- Write -----
--- Appl --- XCF --- Aggr
Rate  Resp  Rate  Resp  Rate
Time  Time
0.000 0.000 0.000 0.000 204.8

Vnodes           : 3      USS held vnodes   : 2
Open objects     : 0      Tokens            : 3
User cache 4k pages : 0      Metadata cache 8k pages : 10

ENOSPC errors    : 0      Disk I/O error    : 0
XCF comm. failures : 0      Cancelled operations : 0
    
```

- From the zFS File System Report, you can navigate to a variety of detail information using cursor-sensitive control. If you place the cursor on any of the lines with file system values, a pop-up window appears showing the details for this file system.

## z/OS 2.2 – zFS Kernel report

- Monitor III zFS Kernel report (ZFSKN) provides measurements counting the calls made to zFS from z/OS UNIX and the average response time of zFS requests.
- Helps to monitor zFS Kernel performance and determine appropriate tuning options

```

RMF V2R2      zFS Kernel      - UTCPXCB      Line 1 of 3

Samples: 60      Systems: 3      Date: 04/21/15      Time: 05.02.00      Range: 60      Sec

System Name      - Request Rate -      --- XCF Rate ---      - Response Time -
                  Local   Remote      Local   Remote      Local   Remote
CB8C              54.18   219.7      0.000   3.417      0.017   0.140
CB8D              9.150   49.52      0.000   0.150      0.149   2.996
CB86             35298   577.5      0.000   46.95      0.610   3.453

```

### Fields in Monitor III zFS Kernel report (ZFSKN):

System Name	Name of the system running zFS. In the context of requests against file systems, this is the name of the requesting system.
Request Rate	Rate of zFS requests during the report interval for file systems which are locally and remotely owned. A file system is locally owned if the requesting system is also the owner of the file system. It is remotely owned if the owner of the file system is not the requesting system.
XCF Rate	Rate of zFS requests during the report interval requiring data from another system via XCF, both for locally and remotely owned file systems.
Response Time	Average time in milliseconds required for the completion of the zFS requests during the report interval for locally and remotely owned file systems.

## z/OS V2.2 – RMF Data Portal ZFS Reports

- The browser based version of ZFSOVW, ZFSFS and ZFSKN can be requested from the RMF Distributed Data Server (DDS) by use of following URL (example is for ZFSKN):

<http://sysf.boebl.de:8803/gpm/rmfm3.xml?report=ZFSKN&resource=„,SYSPLEX“>

System Name	Request Rate Local	Request Rate Remote	XCF Rate Local	XCF Rate Remote	Response Time Local	Response Time Remote
CB8A	113.5	0.000	0.067	0.000	17269	0.000
CB8C	378.4	554.6	0.000	72.53	26.00	10473
CB86	26.40	1569	0.000	119.7	14.00	4302
CB89	21.80	524.0	0.000	72.42	11.00	10198

The new Monitor III ZFSOVW, ZFSFS and ZFSKN reports were added to the standard reports which are available in XML format by means of the RMF Distributed Data Server.

Therefore all statistics of the new Monitor III zFS ISPF reports are available also by means of the Monitor III Data Portal without limitations.

The zFS Reports can be selected from the list of full RMF reports that are available for the SYSPLEX resource and basically all report columns can be displayed in the browser window.

As alternative you can invoke a zFS report by use of following URLs :

ZFSOVW report: <http://hostname:8803/gpm/rmfm3.xml?report=ZFSOVW&resource=„,SYSPLEX“>

ZFSFS report: <http://hostname:8803/gpm/rmfm3.xml?report=ZFSFS&resource=„,SYSPLEX“>

ZFSKN report: <http://hostname:8803/gpm/rmfm3.xml?report=ZFSKN&resource=„,SYSPLEX“>



## z/OS V2.2 – zFS Reporting enhancements: RMF Distributed Data Server Metrics

- All performance metrics from the new zFS sysplex reports are added to the ZFS resource in the DDS and are promoted to the SYSPLEX resource as well.

**Resource Hierarchy:**

- Sysplex
- MVS Image
- I/O Subsystem
- zFS
- Aggregate

**RMF Data Portal for z/OS - Available metrics for: ,SYSDPLEX,SYSPLEX**

Metric description	Help	Id
by aggregate		
% read (in I/O rate) (sysplex) by aggregate	Explanation	8D4FF0
% read (in I/O rate) by aggregate	Explanation	8D5000
% space used (sysplex) by aggregate	Explanation	8D5020
% space used by aggregate	Explanation	8D5030
# cancelled operations (sysplex) by aggregate	Explanation	8D5050
# cancelled operations by aggregate	Explanation	8D5060
# disk I/O errors (sysplex) by aggregate	Explanation	8D5080
# disk I/O errors by aggregate	Explanation	8D5090

**RMF Data Portal for z/OS - Available metrics for: SYSF,BHBE.SYSF.OMVS,AGGREGATE**

Metric description	Help	Id
% read (in I/O rate)	Explanation	8D4FE0
% space used	Explanation	8D5010
% used space	Explanation	8D2AE0
# cancelled operations	Explanation	8D5040
# disk I/O errors	Explanation	8D5070
# open objects	Explanation	8D50A0
# tokens	Explanation	8D50D0
# vnodes	Explanation	8D5100
# ENOSPC errors	Explanation	8D5130
# USS held vnodes	Explanation	8D5160
# XCF communication failures	Explanation	8D5190
# 4K pages in user cache	Explanation	8D51C0
# 8K pages in metadata cache	Explanation	8D51FD
aggregate read rate	Explanation	8D5220
aggregate write rate	Explanation	8D5250
application read rate	Explanation	8D5280
application read response time	Explanation	8D52B0
application write rate	Explanation	8D52E0
application write response time	Explanation	8D5310

(Example is for AGGREGATE resource, a child of ZFS)

- The performance metrics for the new zFS sysplex reports were added to the zFS related DDS metrics. The metrics are associated to the Aggregate, ZFS and SYSPLEX resources of the DDS resource tree.
- DDS is the data source for z/OSMF Resource Monitoring, the state-of-the-art graphical workstation frontend for RMF performance data.
- All DDS metrics can be configured for continuous monitoring with z/OSMF in terms of metrics groups and dashboards.

# RMF Enhancements at a Glance



- IBM z13 and z13s Support
  - RMF Statistics for Simultaneous Multithreading (SMT)
  - Extended ICSF Measurements for Crypto Express5S
  - Monitoring of new capping types



- SCM I/O Adapter Performance Reporting
  - Statistics for EADM Subchannel Activities
  - Performance and Thruput on SCM Card Level

- PCIE Activity Reporting

- z/OS V2.2 Monitor III PCIE Report
- Additional Measurements for PCIe attached RoCE and zEDC Devices on z13
- RMF support for SMC-D over ISM on z13 and z13s




- z/OS V2.2 Monitor III Job Resource Consumption Reporting

- CPU, I/O and Storage Consumption Data at a Glance
- Detailed Statistics for Job related GQSCAN Activities

- z/OS V2.2 zFS Reporting Enhancements

- Sysplex-wide Statistics for zFS Usage and Performance
- Improved zFS Data Gathering Performance
- Additional zFS Statistics for Shared File System environments



- DS8000 Support

- RMF Monitoring of SuperPAV capability

- Mobile Workload Reporting

- Statistics about total and mobile CPU consumption



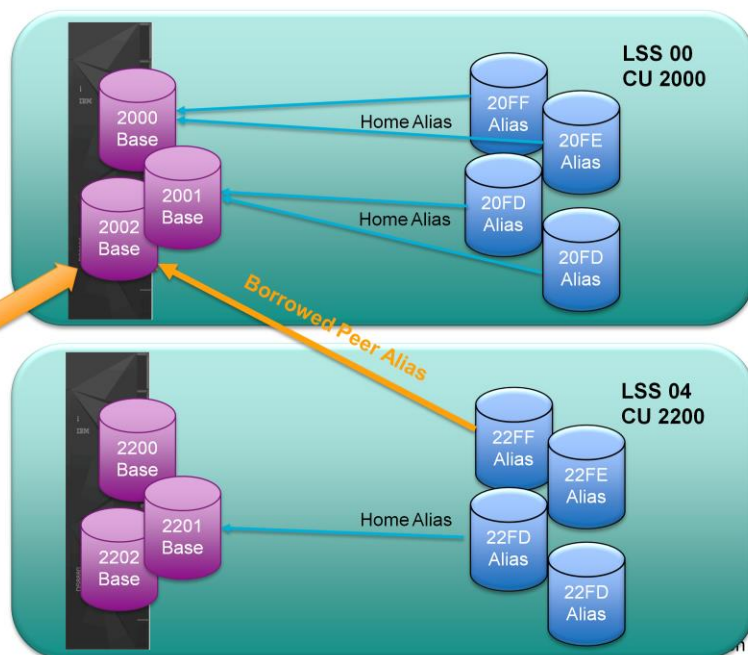


## DS8000 – SuperPAV

- Extension of HyperPAV
  - Aliases are used in an on demand fashion
- Allows aliases to be shared among “like” control units on a storage controller
  - Aliases are first selected from the “home” control unit
  - Aliases are borrowed from a “peer” when no “home” aliases are available



Device 2002 needs alias but none available in “home” alias pool within CU 2000. Device 2002 borrows an alias from peer CU 2200 and assigns alias 22FF to 2002.



Parallel Access Volumes (PAV) modes:

- **PAV base** mode is the mode when alias devices are assigned to one PAV base device. An I/O for a PAV base device is executed using aliases assigned to that PAV base device.
- **HyperPAV** mode is the mode when a pool of alias devices is assigned to one LCU. An I/O for a PAV base device can be executed using any alias device of that pool.
- **SuperPAV** mode is the mode when a pool of alias devices is assigned to one LCU and multiple LCUs are grouped into one Alias Management Group (AMG). An I/O for a PAV base device can be executed using any alias device of these multiple alias pools. The favored way is to use the alias device assigned to the same LCU (home LCU) that the PAV base device is assigned to.

## DS8000 – SuperPAV: What are “like” control units?

- Even and Odd numbered control units on a storage controller with like paths are grouped into Alias Management Groups (AMGs)
- When “home” aliases are exhausted, z/OS will look for free aliases in same AMG

### AMG

CUADDs 01, 03, 05

#### PATHS

- CHPID 10 SW 25 PORT 08
- CHPID 28 SW 33 PORT 17

### AMG

CUADDs 07, 0B, 21, 25

#### PATHS

- CHPID 10 SW 25 PORT 15
- CHPID 28 SW 33 PORT 0A

These odd CUs are in separate AMG groups because their destination ports are different

### AMG

CUADDs 00, 04, 10

#### PATHS

- CHPID 10 SW 25 PORT 08
- CHPID 28 SW 33 PORT 17

### AMG

CUADDs 02, 06, 08

#### PATHS

- CHPID 12 SW 25 PORT 08
- CHPID 2A SW 33 PORT 17

These even CUs are in separate AMG groups because their CHPIDs are different

# SuperPAV – RMF Postprocessor I/O Queuing Activity Report

- New Alias Management Groups section
- System-wide summary of LCU performance data on Alias Management Group (AMG) level

I/O QUEUING ACTIVITY														
z/OS V2R2			SYSTEM ID TAO				DATE 03/31/2016				INTERVAL 09.59.995			
			RPT VERSION V2R2 RMF				TIME 11.40.00				CYCLE 1.000 SECONDS			
ALIAS MANAGEMENT GROUPS														
AMG	DCM	GROUP	CHAN	CHPID	% DP	% CU	AVG	AVG	CONTENTION	DELAY	AVG	HPAV	AVG	DATA
	MIN	MAX	DEF	TAKEN	BUSY	BUSY	CUB	CMR	RATE	Q	CSS	WAIT	OPEN	XFER
			PATHS				DLY	DLY		LNTH	DLY		EXCH	CONC
00000001				90	741.32	0.00	0.00	0.0	0.2					
				91	741.62	0.00	0.00	0.0	0.2					
			*		1482.9	0.00	0.00	0.0	0.2	0.000	0.00	27.2	2.013	21
00000002				90	741.25	0.00	0.00	0.0	0.2					
				91	743.49	0.00	0.00	0.0	0.2					
			*		1484.7	0.00	0.00	0.0	0.2	0.000	0.00	27.1	1.973	24
00000003				5D	358.47	0.00	0.00	0.0	16.2					
				65	346.24	0.00	0.00	0.0	16.0					
				34	121.38	0.00	0.00	0.0	16.8					
				5E	332.06	0.00	0.00	0.0	16.2					
			*		1158.1	0.00	0.00	0.0	16.2	0.000	0.00	2.8	1.846	20
00000004				5D	343.01	0.00	0.00	0.0	9.5					
				65	397.37	0.00	0.00	0.0	9.3					



Maximum number of Aliases that are used concurrently by an LCU within the AMG

With RMF new function APAR OA49415, RMF supports the SuperPAV functionality of the DS8000. The support is available for z/OS 2.1 and z/OS 2.2.

New SuperPAV statistics are collected in SMF 74-1 (Device Activity) and SMF 78-3 (I/O Queuing Activity).

The RMF Postprocessor I/O queuing activity report is now grouped into 3 sections:

- Input/Output Processors
- Alias Management Groups
- Logical Control Units

With the new SuperPAV functionality, that allows the use of Aliases across control units, there's a need to look at data summed at the Alias Management Group (AMG) level. This is addressed by the new Alias Management Groups section.

This section is reported only, if the system is running in SuperPAV mode and there are AMGs defined. It shows accumulated values for CHPIDs and LCUs that are grouped into AMGs.

# SuperPAV – Postprocessor I/O Queuing Activity Report ...

LOGICAL CONTROL UNITS																
LCU/ AMG	CU	DCM	GROUP	CHAN	CHPID	% DP	% CU	AVG	AVG	CONTENTION	DELAY	AVG	HPAV	AVG	DATA	
		MIN	MAX	DEF	TAKEN	BUSY	BUSY	CUB	CMR	RATE	Q	CSS	WAIT	OPEN	XFER	
				PATHS				DLY	DLY		LNGLTH	DLY		EXCH	CONC	
0041	4002			90	247.26	0.00	0.00	0.0	0.2							
00000002				91	247.88	0.00	0.00	0.0	0.2							
				*	495.13	0.00	0.00	0.0	0.2	0.000	0.00	27.1	2.034	22	2.48	2.39
0043	4202			90	246.78	0.00	0.00	0.0	0.2							
00000002				91	247.90	0.00	0.00	0.0	0.2							
				*	494.68	0.00	0.00	0.0	0.2	0.000	0.00	27.1	1.985	24	2.53	2.45
0045	4402			90	247.21	0.00	0.00	0.0	0.2							
00000002				91	247.72	0.00	0.00	0.0	0.2							
				*	494.93	0.00	0.00	0.0	0.2	0.000	0.00	27.1	1.904	21	2.64	2.55
004A	4802			5D	82.401	0.00	0.00	0.0	16.3							
00000003				65	202.19	0.00	0.00	0.0	15.9							
				34	54.142	0.00	0.00	0.0	16.9							
				5E	58.144	0.00	0.00	0.0	16.3							
				*	396.87	0.00	0.00	0.0	16.1	0.000	0.00	1.8	1.867	18	15.4	8.95
004E	4A02			5D	175.19	0.00	0.00	0.0	16.2							
00000003				65	46.904	0.00	0.00	0.0	16.2							
				34	38.967	0.00	0.00	0.0	16.8							
				5E	117.55	0.00	0.00	0.0	16.2							
				*	378.62	0.00	0.00	0.0	16.3	0.000	0.00	3.6	1.901	19	14.6	8.49

The maximum number of concurrently used Aliases (including borrowed Aliases) 10/30/2016

- If the LCU is grouped to an AMG, the AMG number is reported under column LCU/AMG in the Logical Control Unit section.
- LCUs are now sorted by AMG and LCU number.
- LCUs not grouped to an AMG are displayed first.

## SuperPAV – New/changed Overview Conditions

- New qualifiers for CHPID (IOCHPID) and AMG (IOAMG) are introduced to support overview reports for AMG metrics

Condition	Condition Name	Qualifier	Source	Algorithm
Contention rate	IOCTR	lcluid <b>IOAMG(amg)</b>	R783QCT SMF78INT	QCT/INT
Average queue length of delayed I/O requests	IODLQ	lcluid <b>IOAMG(amg)</b>	R783QCT R783QSM	(QSM-QCT)/QCT
Channel path taken rate	IOART	lcluid <b>IOAMG(amg)</b> <b>IOCHPID(chpid)</b>	R783PT SMF78INT	(Pti)/INT
Percentage of requests caused by control unit busy	IOCUB	lcluid <b>IOAMG(amg)</b> <b>IOCHPID(chpid)</b>	R783DPB R783CUB R783PT	MAX(CUBi*100)/(PT+CUB+DPB)i MIN(CUBi*100)/(PT+CUB+DPB)i <b>When IOCHPID selected:</b> <b>CUBi*100/(PT+CUB+DPB)i</b>
Percentage of requests caused by director port busy	IODPB	lcluid <b>IOAMG(amg)</b> <b>IOCHPID(chpid)</b>	R783DPB R783CUB R783PT	MAX(DPBi*100)/(PT+CUB+DPB)i MIN(DPBi*100)/(PT+CUB+DPB)i <b>When IOCHPID selected:</b> <b>DPBi*100/(PT+CUB+DPB)i</b>

With the RMF support for SuperPAV, new overview conditions are added based on SMF record 78 subtype 3. These overview conditions can be used to report the new SuperPAV metric data in the SMF 78-3.

New qualifiers for CHPID (IOCHPID) and AMG (IOAMG) are introduced.

Here are some examples of using the SuperPAV (IOXxxxxx) conditions:

```
OVW(IOCUB(IOCUB(IOAMG(00000002),IOCHPID(30))))
```

```
OVW(IOCUBAI(IOCUB(IOAMG(00000002))))
```

In the **Algorithm** column:

**MAX** Applies to exception operator GE, and specifies the sum of each channel path taken, where i represents channel path 0 to channel path 7.

**MIN** Applies to exception operator LE, and specifies the sum of each channel path taken, where i represents channel path 0 to channel path 7.

## SuperPAV – New/changed Overview Conditions ...

Condition	Condition Name	Qualifier	Source	Algorithm
Average control unit busy delay time	IOCBT	Icuid <b>IOAMG(amg)</b> <b>IOCHPID(chpid)</b>	R783CBT R783PT	Sum(CBT)/Sum(PT)
Average initial command response time	IOCMR	Icuid <b>IOAMG(amg)</b> <b>IOCHPID(chpid)</b>	R783CMR R783PT	Sum(CMR)/Sum(PT)
Average channel subsystem delay time	IOCSS	Icuid <b>IOAMG(amg)</b>	R783CSST R783PT	CSST/Sum(PT)
HyperPAV wait ratio	IOHWAIT	Icuid <b>IOAMG(amg)</b>	R783HNAI R783HTIO	HNAI/HTIO <b>(Sum(HNAI))/(Sum(HTIO))</b>
Maximum number of in-use aliases	IOHMAX	Icuid <b>IOAMG(amg)</b>	R783HAIU R783XHBC	<b>HAIU + XHBC</b> <b>When IOAMG selected:</b> <b>Maximumf over all LCUs of that</b> <b>AMG</b> <b>MAX(HAIU+XHBC)</b>
Maximum number of in-use HyperPAV aliases for one device	IOHDMAX	Icuid <b>IOAMG(amg)</b>	R783HCAD	HCAD



## SuperPAV – New/changed Overview Conditions ...

Condition	Condition Name	Qualifier	Source	Algorithm
The high watermark of queued I/O requests	IOHIOQC	lcuid IOAMG(amg)	R783HIOQ	HIOQ
Ratio of successful alias requests	IOXSAREQ	IOAMG(amg)	R783XAUC R783XANC	Sum(XAUC)/Sum(XANC)
Ratio of unsuccessful alias requests in home LCU	IOXUAHRQ	IOAMG(amg)	R783XNHC R783XANC	Sum(XNHC)/Sum(XANC)
Rate of aliases borrowed from peer LCUs	IOXABC	IOAMG(amg)	R783XABC SMF78INT	Sum(XABC)/INT
High water mark of concurrently borrowed aliases	IOXHCBA	IOAMG(amg)	R783XHBC	XHBC
Rate of aliases loaned to a peer LCU	IOXALC	IOAMG(amg)	R783XALC SMF78INT	Sum(XALC)/INT
High water mark of concurrently loaned aliases to a peer LCU	IOXHCLA	IOAMG(amg)	R783XHLC	XHLC
Average queue length when an alias was needed	IOXCQD	IOAMG(amg)	R783XCQD R783XANC	Sum(XCQD)/Sum(XANC)
Average number of in use aliases when an alias was needed	IOXIUAC	IOAMG(amg)	R783XCIU R783XANC	Sum(XCIU)/Sum(XANC)

## RMF Enhancements at a Glance



- IBM z13 and z13s Support
  - RMF Statistics for Simultaneous Multithreading (SMT)
  - Extended ICSF Measurements for Crypto Express5S
  - Monitoring of new capping types



- SCM I/O Adapter Performance Reporting
  - Statistics for EADM Subchannel Activities
  - Performance and Thruput on SCM Card Level

- PCIE Activity Reporting

- z/OS V2.2 Monitor III PCIE Report
- Additional Measurements for PCIe attached RoCE and zEDC Devices on z13
- RMF support for SMC-D over ISM on z13 and z13s



- z/OS V2.2 Monitor III Job Resource Consumption Reporting

- CPU, I/O and Storage Consumption Data at a Glance
- Detailed Statistics for Job related GQSCAN Activities

- z/OS V2.2 zFS Reporting Enhancements

- Sysplex-wide Statistics for zFS Usage and Performance
- Improved zFS Data Gathering Performance
- Additional zFS Statistics for Shared File System environments



- DS8000 Support

- RMF Monitoring of SuperPAV capability

- Mobile Workload Reporting

- Statistics about total and mobile CPU consumption

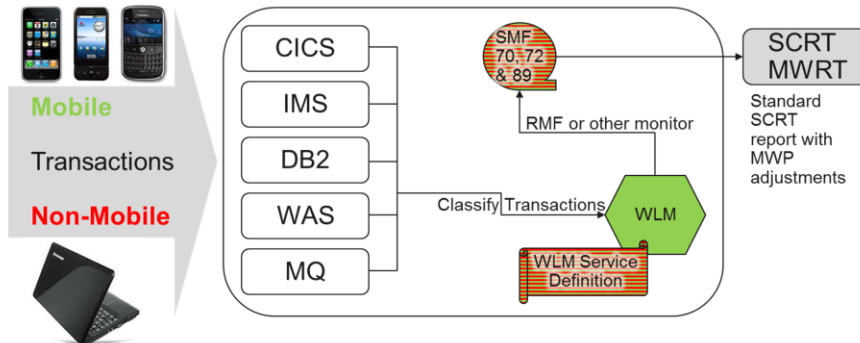


# RMF Mobile Workload Reporting Enhancements

## WLM Support for Mobile Workload Pricing(MWP):

Allow identification of mobile transactions in WLM service definition

- Via a **transaction level** attribute
- Processor consumption data aggregated by WLM
- Reporting integrated into standard performance monitors (RMF) and low volume SMF records
- Applicable to all workloads, including address space work, enclave work and CICS/IMS work




### WLM Policy

Classification rules with new reporting attributes:

- NONE for normal transactions (default)
- MOBILE for mobile transactions.
- CATEGORYA/ CATEGORYB for future use.

- Mobile Workload Pricing is an IBM Software Pricing Option, announced in May 2014
- It offers a discount on MSUs consumed by transactions that originated on a mobile device
- To take advantage of this discount, customers need a process, agreed upon by IBM, to identify (tag and track) their mobile-sourced transactions and their consumption
- Today, this type of mobile workload reporting requires product-specific tooling from high-volume transaction level accounting data, which induces significant overhead both for customers and product development
  - The new WLM support for mobile pricing allows to identify Mobile sourced transactions: Customers can use a new reporting attribute in the WLM classification rules of the WLM service definition to classify transactions as "mobile"
- WLM tracks and reports the total and the mobile CPU consumption for all service and report classes
  - Exploiters of the WLM Execution Delay Monitoring Services like CICS or IMS can provide CPU times for all their transactions
  - As soon as they do, total and mobile CPU consumption data is also available for CICS and IMS transaction service and report classes that previously did not report any CPU consumption data
- WLM also aggregates and reports the system-wide mobile consumption data
- Besides 'mobile' there are two more categories (CATEGORYA and CATEGORYB) that are currently unused but may be used in the future



## RMF MWP Support – Overview

- WLM tracks and reports the total and the mobile CPU consumption for all service and report classes. RMF retrieves and stores the data into new fields of SMF 72-3 Service/Report Class Period Data Section.
- WLM also aggregates and reports the system-wide mobile consumption data. RMF retrieves and stores the data into new fields of SMF 70-1 CPU Control Section
- RMF Postprocessor Workload Activity report displays new statistics about total and mobile CPU consumption.
- The Workload Activity Spreadsheet of the RMF Spreadsheet Reporter is enabled to visualize the new mobile workload measurements
- New Postprocessor overview conditions based on the new fields in SMF 70-1 and 72-3

- With RMF new function APAR OA48466, RMF supports the new WLM support for Mobile Pricing. PTFs are available for z/OS V2.1 and z/OS V2.2

# RMF MWP Support - Postprocessor Workload Activity report

Processor consumption data for transaction service / report classes !

WORKLOAD ACTIVITY									
z/OS V2R2		SYSPLEX UTCPLXCB		DATE 09/28/2015		INTERVAL 01.00.000		MODE = GOAL	
		RPT VERSION V2R2 RMF		TIME 14.15.00					
REPORT BY: POLICY= BASEPOL		WORKLOAD=CICSCPU		SERVICE CLASS=CICSLow		RESOURCE GROUP=*NONE		PERIOD=1	
				CRITICAL =NONE					
-TRANSACTIONS- TRANS-TIME HHH.MM.SS.TTT									
AVG	0.00	ACTUAL							14
MPL	0.00	EXECUTION							14
ENDED	68875	QUEUED							0
END/S	1150.39	R/S AFFIN							0
#SWAPS	0	INELIGIBLE							0
EXCTD	58263	CONVERSION							0
AVG ENC	0.00	STD DEV							28
REM ENC	0.00								
MS ENC	0.00								
TRANSACTION APPL% :									
TOTAL :		CP	58.18	AAP/IIP ON CP	0.00	AAP/IIP	0.00		
MOBILE :		CP	0.00	AAP/IIP ON CP	0.00	AAP/IIP	0.00		

New section: Total and Mobile CPU consumption for transactions

MOBILE is subset of TOTAL

- With the RMF support for Mobile Workload pricing the RMF Postprocessor Workload Activity Postprocessor (WLMGL) report is enhanced.
- A new TRANSACTION APPL% section with the Total and Mobile CPU consumption on the different processor types is available on the reporting category levels:
  - Service Class
  - Service Class Period
  - Report Class
  - Report Class Period
  - Workload
  - Policy
- New fields in the TRANSACTION APPL% section:

<b>TOTAL</b>	Total percentage of the processor time used by transactions running on the different processor types.
CP	Total percentage of general purpose processor time used by transactions.
AAP/IIP ON CP	Total percentage of general purpose processor time used by transactions eligible to run on specialty processors.
AAP/IIP	Total percentage of specialty processor time used by transactions.
<b>MOBILE</b>	Percentage of the processor time used by transactions classified with reporting attribute MOBILE running on the different processor types. MOBILE is a subset of TOTAL.
CP	Percentage of general purpose processor time used by transactions classified with reporting attribute MOBILE.
AAP/IIP ON CP	Percentage of general purpose processor time used by transactions classified with reporting attribute MOBILE, eligible to run on specialty processors.
AAP/IIP	Percentage of specialty processor time used by transactions classified with reporting attribute MOBILE.

**Note: Transaction Application Time %**

When transaction processor usage is reported to WLM through IWM4RPT or IWM4MNTF services, the consumed service units are accounted to the transaction service or report classes, and deducted from the region's service and report classes. If the number of transactions is very small and a single transaction reports high processor times, it can occur that processor times become negative. In such a case RMF, displays asterisk (\*).

## RMF MWP Support – SMF 70-1 based Overview Conditions

Condition	Name	Qualifier	Source	Algorithm
Long-term average of CPU service (millions of service units) consumed by transactions classified with reporting attribute MOBILE	LACSM	none	SMF70LACM	Value or comparison
Long-term average of CPU service (millions of service units) consumed by transactions classified with reporting attribute CATEGORYA	LACSA	none	SMF70LACA	Value or comparison
Long-term average of CPU service (millions of service units) consumed by transactions classified with reporting attribute CATEGORYB	LACSB	none	SMF70LACB	Value or comparison

- RMF retrieves and stores the system-wide mobile consumption data into new fields of SMF 70-1 CPU Control Section.  
A new set of overview conditions can be used to report the new SMF 70-1 data.

## RMF MWP Support – SMF 72-3 based Overview Conditions

Condition	Name	Qualifier	Source	Algorithm
Total service per second, consumed by transactions, executed on general purpose processors	TCPSRV	type	R723TSUCP Interval	$SUM(R723TSUCP) / Interval$
Total application execution time, consumed by transactions in seconds, executed on general purpose processors	TCPSEC	type	R723TSUCP R723MADJ R723MCPU	$SUM((R723TSUCP * R723MADJ) / (1600 * R723MCPU))$
Total percentage of general purpose processor time used by transactions	TAPPLCP	type	R723TSUCP R723MADJ R723MCPU R723MCF Interval	$SUM((R723TSUCP * R723MADJ) / (1600 * R723MCPU)) / (Interval * (R723MCF / 1024)) * 100$
Total service per second consumed by transactions, executed on specialty processors	TSPSRV	type	R723TSUSP Interval	$SUM(R723TSUSP) / Interval$
Total application execution time, consumed by transactions in seconds, executed on specialty processors	TSPSEC	type	R723TSUSP R723MADJ R723MCPU	$SUM((R723TSUSP * R723MADJ) / (1600 * R723MCPU))$
Total percentage of specialty processor time used by transactions	TAPPLSP	type	R723TSUSP R723MADJ R723MCPU R723MCF Interval	$SUM((R723TSUSP * R723MADJ) / (1600 * R723MCPU)) / (Interval * (R723MCF / 1024)) * 100$

78

© 2016 IBM Corporation

- RMF retrieves and stores the total and the mobile CPU consumption for service and report classes into new fields of SMF 72-3 Service/Report Class Period Data Section.
- A new set of overview conditions can be used to report the new SMF 72-3 data.
- The type qualifier in the overview condition can have one of the following values:
  - S.sname.period Service class period
  - S.sname Service class
  - R.rname.period Report class period
  - R.rname Report class
  - W.wname Workload
  - POLICY Policy

## RMF MWP Support – SMF 72-3 based Overview Conditions

Condition	Name	Qualifier	Source	Algorithm
Total service per second, consumed by transactions, eligible to run on specialty processors but executed on general purpose processors	TOCPSRV	type	R723TSUOCP Interval	$SUM(R723TSUOCP) / Interval$
Total application execution time, consumed by transactions in seconds, eligible to run on specialty processors but executed on general purpose processors	TOCPSEC	type	R723TSUOCP R723MADJ R723MCPU	$SUM((R723TSUOCP * R723MADJ) / (1600 * R723MCPU))$
Total percentage of general purpose processor time used by transactions eligible to run on specialty processors	TAPPLOCP	type	R723TSUOCP R723MADJ R723MCPU R723MCF Interval	$SUM((R723TSUOCP * R723MADJ) / (1600 * R723MCPU)) / (Interval * (R723MCF / 1024)) * 100$
Service per second, consumed by transactions classified with reporting attribute MOBILE, executed on general purpose processors	TMCPSRV	type	R723MSUCP Interval	$SUM(R723MSUCP) / Interval$
Application execution time, consumed by transactions classified with reporting attribute MOBILE in seconds, executed on general purpose processors	TMCPSEC	type	R723MSUCP R723MADJ R723MCPU	$SUM((R723MSUCP * R723MADJ) / (1600 * R723MCPU))$



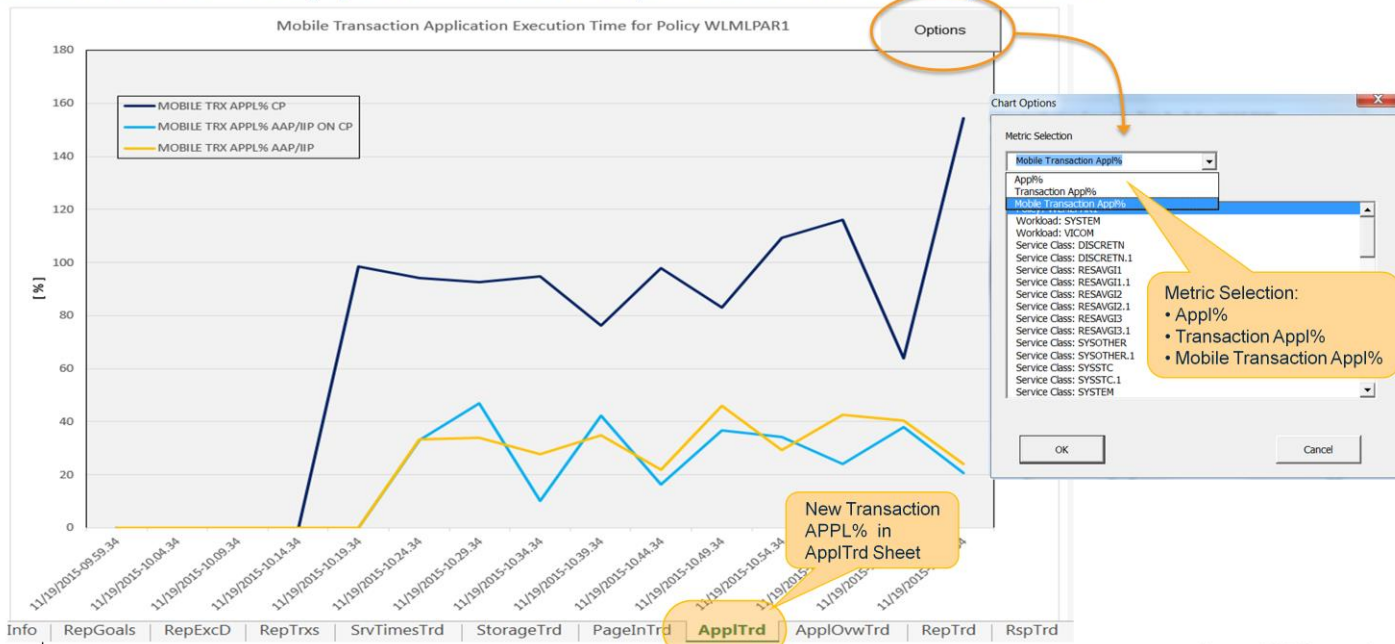
## RMF MWP Support – SMF 72-3 based Overview Conditions

Condition	Name	Qualifier	Source	Algorithm
Percentage of general purpose processors used by transactions classified with reporting attribute MOBILE	MAPPLCP	type	R723MSUCP R723MADJ R723MCPU R723MCF Interval	$SUM((R723MSUCP * R723MADJ) / (1600 * R723MCPU)) / (Interval * (R723MCF / 1024)) * 100$
Service per second, consumed by transactions classified with reporting attribute MOBILE, executed on specialty processors	TMSPSRV	type	R723MSUSP Interval	$SUM(R723MSUSP) / Interval$
Application execution time, consumed by transactions classified with reporting attribute MOBILE in seconds, executed on specialty processors	TMSPSEC	type	R723MSUSP R723MADJ R723MCPU	$SUM((R723MSUSP * R723MADJ) / (1600 * R723MCPU))$
Percentage of specialty processor time used by transactions classified with reporting attribute MOBILE	MAPPLSP	type	R723MSUSP R723MADJ R723MCPU R723MCF Interval	$SUM((R723MSUSP * R723MADJ) / (1600 * R723MCPU)) / (Interval * (R723MCF / 1024)) * 100$
Service per second, consumed by transactions classified with reporting attribute MOBILE, eligible to run on specialty processors but executed on general purpose processors	TMOCPSRV	type	R723MSUOCP Interval	$SUM(R723MSUOCP) / Interval$

## RMF MWP Support – SMF 72-3 based Overview Conditions

Condition	Name	Qualifier	Source	Algorithm
Application execution time, consumed by transactions classified with reporting attribute MOBILE in seconds, eligible to run on specialty processors but executed on general purpose processors	TMOCPSEC	type	R723MSUOCP R723MADJ R723MCPU	$SUM((R723MSUOCP * R723MADJ) / (1600 * R723MCPU))$
Percentage of general purpose processor time used by transactions classified with reporting attribute MOBILE eligible to run on specialty processors	MAPPLOCP	type	R723MSUOCP R723MADJ R723MCPU R723MCF Interval	$SUM((R723MSUOCP * R723MADJ) / (1600 * R723MCPU)) / (Interval * R723MCF / 1024) * 100$

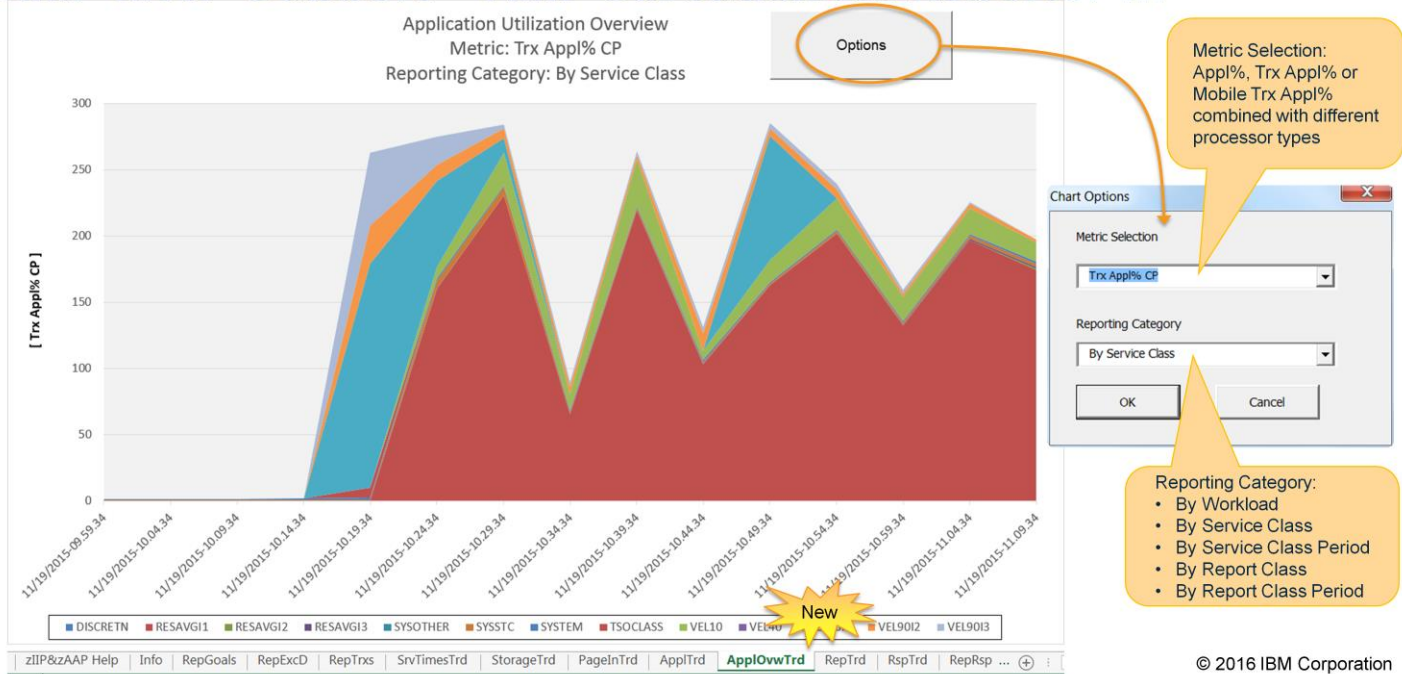
# RMF MWP Support – RMF Spreadsheet Reporter



© 2016 IBM Corporation

- In the RMF Spreadsheet Reporter, the Workload Activity Trend Report Spreadsheet provides additional functionality to visualize mobile workloads.
- The sheet AppITrd now offers enhanced chart options to select one of the following metric scopes:
  - App%
  - Transaction App% : Total percentage of the processor time used by transactions running on the different processor types
  - Mobile Transaction App% : Percentage of the processor time used by transactions classified with reporting attribute MOBILE running on the different processor types.

# RMF MWP Support – RMF Spreadsheet Reporter ...



- A new AppIOvwTrd sheet was added to the Workload Activity Trend Report Spreadsheet. The new sheet displays a selectable Application Utilization metric for a selectable workload group.
- Following metrics are supported:
  - Appl%: CP, AAP ON CP, IIP ON CP, AAP, IIP
  - Trx Appl%: CP, AAP/IIP ON CP, AAP/IIP
  - Mobile Trx Appl%: CP, AAP/IIP ON CP, AAP/IIP
- Following Workload Groups are supported
  - By Workload
  - By Service Class
  - By Service Class Period
  - By Report Class
  - By Report Class Period



**The zEvent Mobile Application:**  
Receive Push Messages based on critical  
System Events  
Access to RMF Data Portal and z/OSMF Resource  
Monitoring

# z/OSMF Resource Monitoring & Mobile Devices



Very convenient to know right now that the Sysplex of my boss has got enough CSA and SQA...

 **Cross platform performance monitoring anytime and everywhere!**



- z/OSMF supports by default the recent versions of Mozilla Firefox and Microsoft Internet Explorer
- However, z/OSMF can be used together with the most browsers available for mobile devices as well
- Cross platform performance monitoring anytime and everywhere: Try z/OSMF on your tablet PC or smartphone!

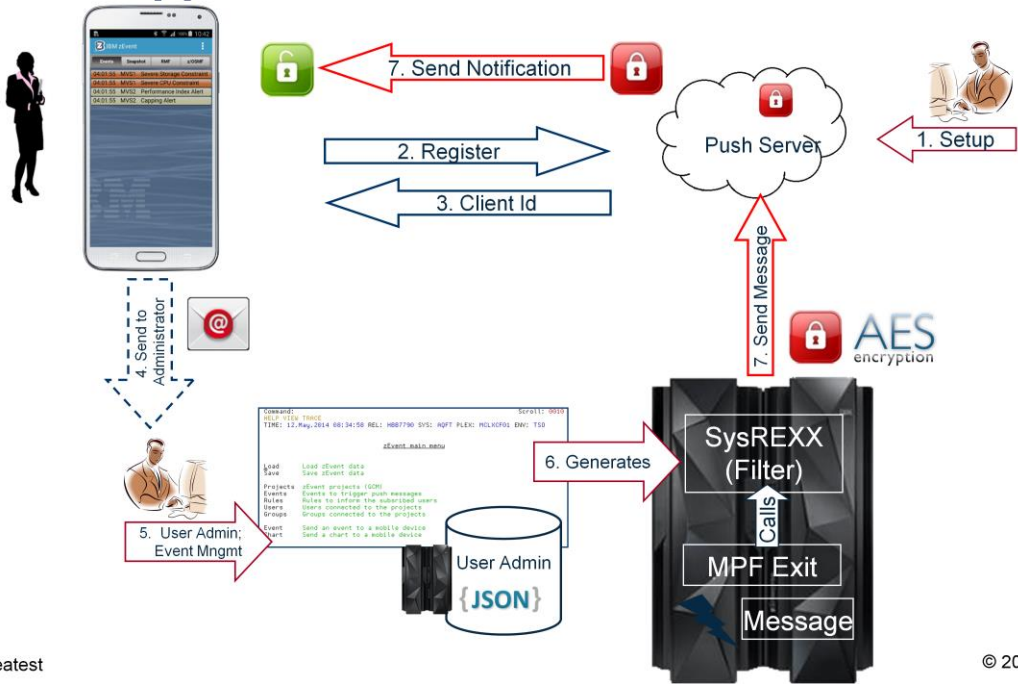
## zEvent Mobile Application

- Receive push messages based critical system events
- Access to z/OS performance data from mobile devices
- Single point of control for the enterprise
- Developed with IBM Mobile First Studio
- Prototype status (Android OS)
- Current availability target is 1Q 2016
- Supported Monitoring Facilities:
  - RMF Performance Data Portal
  - z/OSMF Resource Monitoring



- Last but not least, the zEvent Mobile Application:
- This mobile App is currently under development and all details and capabilities – including the application name – are subject to change
- In a nutshell, the app provides the following two main features to system administrators and performance analysts:
  - Receive push messages based on critical system events instantly
  - Access to the RMF Data Portal and z/OSMF Resource Monitoring anywhere and every time

# zEvent Mobile Application





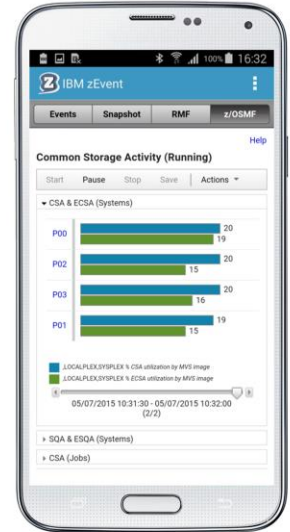
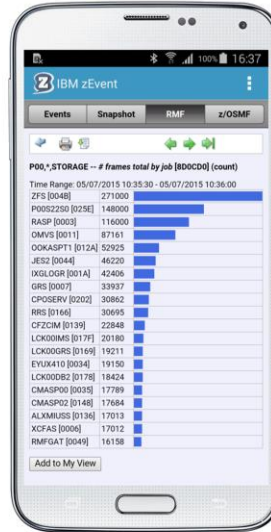
## zEvent Mobile Application

- Connect to all z/OS systems in the enterprise
- Direct connection from the App to the z/OS system or via Mobile First Server



# zEvent Mobile Application

- Monitor z/OS system performance with the RMF Data Portal
- Exploit the extended capabilities of z/OSMF Resource Monitoring



## Information and Tools

- website: [www.ibm.com/systems/z/os/zos/features/rmf/](http://www.ibm.com/systems/z/os/zos/features/rmf/)
  - Product information, newsletters, presentations, ...
  - Downloads
    - Spreadsheet Reporter
    - Postprocessor XML Toolkit
- RMF email address: [rmf@de.ibm.com](mailto:rmf@de.ibm.com)
- Documentation and news:
  - RMF Report Analysis, SC34-2665
  - RMF User's Guide, SC34-2664
  - RMF Programmer's Guide, SC34-2667
  - Latest version of PDF files can be downloaded from:  
[www.ibm.com/systems/z/os/zos/library/bkserv/v2r2pdf/#ERB](http://www.ibm.com/systems/z/os/zos/library/bkserv/v2r2pdf/#ERB)



# Function Reference

Function	Availability
Support for Simultaneous Multithreading (SMT) Extended Measurements for PCIe attached Devices Additional ICFS Statistics for Crypto Express5S Support for LPARs with up to 4 TB real storage Support for z13 IBM Integrated Coupling Adapter (ICA SR) RMF z13 Toleration for z/OS V1R10 and V1R11 RMF z13 Toleration for z/OS V1R12 and V1R13 RMF support for new capping types	APAR OA44101 APAR OA44524 APAR OA43493 APAR OA44503 APAR OA44502 APAR OA45890 APAR OA45833 APAR OA48688
SCM I/O Adapter Performance Reporting	z/OS 2.2 RMF
RMF Monitor III PCIe Online Reporting	z/OS 2.2 RMF
RMF support for SMC-D over ISM	APAR OA49113
RMF Monitor III Job Resource Consumption Reporting	z/OS 2.2 RMF
zFS Reporting Enhancements	z/OS 2.2 RMF
RMF support for SuperPAV	APAR OA49415
RMF Mobile Workload Reporting Enhancements	APAR OA48466
zEvent Mobile Application	APAR PI57136 <a href="http://www.ibm.com/systems/z/os/zos/features/zevent/zevent.html">http://www.ibm.com/systems/z/os/zos/features/zevent/zevent.html</a>

z13 and z13s Enhancements

z13 and z13s Enhancements