

z/OS
Resource Measurement Facility

The Latest and Greatest:
z14 and z/OS V2R3




RMF Enhancements at a Glance

- IBM z14 Support
 - Monitoring of Virtual Flash Memory
 - RMF support for for Crypto Express6S
 - IBM 10GbE RoCE Express2 support



- DS8880 Support
 - RMF Monitoring of SuperPAV capability
 - RMF Enhancements for zHyperWrite
 - Monitor Performance of PCIe Synchronous I/O links

- RMF Enhancements for Coupling Facility
 - Monitor Throughput and Delays of Asynchronous CF duplexing for lock structures
 - Predict overhead of CF encryption and report encryption state of CF structures



- RMF Enhancements for Spark
 - Report on new WLM functionality for Spark

- RMF WLM Reporting Enhancements
 - RMF support for Shorter WLM Response Time Goals



- Large Page Enhancements
 - Statistics for 2 GB Large Pages

- In accordance with the availability of new z/OS releases and new hardware functionality, the capabilities of RMF are enhanced consecutively
- With availability of IBM z14, RMF provides first day support for a set of new functions:
 - Monitoring of Virtual Flash Memory (replacement of zFlash Express)
 - RMF support for CEX6S crypto adapter
 - IBM 10GbE RoCE Express2 support
 - Monitor performance of PCIe Synchronous I/O links provided by the zHyperLink Express I/O feature
- SuperPAV is a new functionality of the IBM DS8880 series that allows to share Aliases among multiple control units. RMF supports the monitoring of the DS8880 SuperPAV capability with new statistics in SMF records and report enhancements in the RMF Postprocessor I/O Queuing Activity report.
- The zHyperWrite technology enables DB2 to perform parallel log writes to PPRC primary and secondary volumes. This improves both DB2 log throughput and transactional response time. RMF is enhanced to monitor I/O to PPRC secondary devices.
- Asynchronous CF duplexing for lock structures is a solution for continuous availability. RMF provides new statistics that helps to monitor throughput and delays of asynchronous CF duplexing for lock structures.
- z/OS V2.3 provides support for end-to-end encryption for both CF data in flight and data at rest in CF structures. RMF is enhanced to report the encryption state per CF structure and collects new CF structure statistics that can be used to predict the overhead of CF encryption and decryption.
- z/OS Platform for Apache Spark is the IBM solution to enable Apache Spark natively on z/OS. RMF reports on new WLM functionality that allows a better management of Spark on z/OS.
- With z/OS V2R3, WLM enhances the definition of response time goals for service classes. The current lower bound of 15 milliseconds for a response time goal is replaced by one millisecond. RMF is enhanced to report on shorter response times in various RMF reports.
- RMF provides new support for monitoring and reporting of 2 GB memory frames.

RMF z14 Support - Overview

- RMF support for Virtual Flash Memory (VFM) OA50761
- RMF support for 10GbE RoCE Express2 card OA50762
- RMF support for Crypto Express6S (CEX6) card OA50693
- RMF support for zHyperLink OA50755
 - Monitor Performance of PCIE Synchronous I/O links
- RMF z14 Toleration support OA51913



- With various new function APARs, RMF exploits the new functionality of the IBM z14:
 - OA50761: RMF support for Virtual Flash Memory (VFM) on z14
 - The APAR provides the support for Virtual Flash Memory on z/OS 2.2
 - For z/OS 2.1 and z/OS 2.2 the APAR provides the support of SMT for SAPs
 - OA50762: RMF z/OS 2.2 support for 10 GbE RoCE Express2 card.
 - OA50693: RMF support for the Crypto Express6S (CEX6) card.
 - PTFs available for z/OS 2.1 and z/OS 2.2.
 - OA50755: RMF support for zHyperLink
 - PTFs available for z/OS 2.1 and z/OS 2.2.
- RMF toleration support for IBM z14:
 - OA51913: z14 toleration for z/OS 1.13, z/OS 2.1 and z/OS 2.2

z14 - Virtual Flash Memory

- "storage class memory" provided by zFlash Express adapters is replaced with main memory, called Virtual Flash Memory (VFM)
- Each VFM Feature takes 1.5 TB (1536 GB) of memory. Up to four VFM features can be ordered.
- Much simpler management of VFM resource (HMC task)
 - Allocation of VFM storage moves to LPAR activation since LPAR hypervisor "owns" management of partition memory.
 - Customer specifies initial and maximum amount of VFM
 - VFM allocations and definitions for all partitions can be viewed through "Storage Information" panel
- No hardware repair and verify (no cables, no adapters)
- Better performance since no "I/O" to attached adapter takes place.

z14 – RMF VFM Support



OA50761

RMF Monitor III
Storage Class Memory
Activity report

```

RMF V2R2   SCM Activity

Samples: 60   System: TRX1   Date: 03/30/17   Time: 14.40.00   Range: 60

----- EADM Device/Subchannel Summary -----
SSCH Total    SSCH Rate    PEND Time    IOPQ Time    ICMR Time
 30541         509.02         0.003         0.001         0.000

Card ID      Util(%)      Read(B/s)    Write(B/s)    Req Rate    Resp Time    IOPQ
Part Total   Part Total   Part Total   Part Total   Part Total   Part Total   Time

VFM          0.00  0.00  14.4M  14.4M  16.0M  16.0M  3139  3139  0.003  0.003  0.000
  
```

Virtual Flash
Memory

Percentage SAP
processor time used
for SCM processing

- The RMF support for VFM will be available with z/OS V2.3 RMF. New function APAR OA50761 provides the support for z/OS 2.2.
- RMF reports VFM statistics in the existing RMF Postprocessor Storage Class Memory Activity and Monitor III SCM Activity reports.
- With VFM the meaning of following report fields will change:
 - The terms “Flash Adapter” and “Card ID” will also be used for z14 since VFM can be seen as simulated SCM adapter which is effectively all cache. When VFM is active, the Postprocessor and Monitor III SCM Activity reports display *VFM* in the Card ID report field.
 - *Util%* is the average utilization of Virtual Flash Memory, reported as the percentage of the time spent on System Assist Processors (SAP) for SCM processing compared to the total available SAP time in this reporting interval.

z14 – RMF VFM Support ...

RMF Postprocessor Interval Report [System TRX1] : Storage Class Memory Activity Report

RMF Version : z/OS V2R3 SMF Data : z/OS V2R2

Start : 03/30/2017-13.54.34 End : 03/30/2017-13.59.33 Interval : 05:00:000 minutes

▶ EADM Device/Subchannel Summary

Percentage SAP processor time used for SCM processing

▼ Flash Adapter Activity

| 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) |
|---------|--------------|---------------|-------------------|--------------------|--------------------|---------------------|---------------------|----------------------|--------------------------|---------------------------|----------------------------|-----------------|------------------|
| VFM | 0.00 | 10.49 | 0.00 | 3265M | 0.00 | 1376M | 0.00 | 139046.5 | 0.000 | 0.013 | 0.001 | 0 | 41713943 |

Virtual Flash Memory

z14 - 10GbE RoCE Express2




- z14 introduces next generation of RoCE technology : IBM 10GbE RoCE Express2
- 10GbE RoCE Express2 provides a technology refresh for RoCE on IBM Z.
- RoCE Express2 provides two physical 10GbE ports (no change)
- Key difference:
RoCE Express2 provides increased virtualization (sharing) capabilities allowing RoCE to be extended to more workloads:
 - RoCE Express2 supports 63 Virtual Functions (VFs) per physical port for a total of 126 VFs per PCHID.
 - RoCE Express supports 31 Virtual Functions (VFs) per PCHID.



- Shared Memory Communication via Remote Direct Memory Access (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
- With the IBM 10GbE RoCE Express2 of z14, IBM Z introduces the next generation of RoCE technology
- The 10GbE RoCE Express2 provides a technology refresh for RoCE on IBM Z.

z14 – RMF RoCE Express2 Enhancements



General PCIE Activity

| Function ID | Function CHID | Function Name | Function Status | Owner Job Name | Owner Address Space ID | Function Allocation Time |
|-------------|---------------|---------------|-----------------|----------------|------------------------|--------------------------|
| 2601 | 0204 | RoCE Express2 | Allocated | VTAM | 002B | 1800 |
| 2602 | 0204 | RoCE Express2 | Allocated | VTAM | 002B | 1800 |
| 2603 | 0204 | RoCE Express2 | Allocated | VTAM | 002B | 1800 |
| 2604 | 0204 | RoCE Express2 | Allocated | VTAM | 002B | 1800 |

RMF Postprocessor PCIE Activity report: General PCIE Activity section is enhanced to report RoCE Express-2 activity

RMF Monitor III PCIE Activity report: RoCE Express-2 activity displayed as function type 'RoCE2'

RMF V2R2 PCIE Activity

Samples: 60 System: TRX1 Date: 05/05/17 Time: 03.15.00 Range: 100 Sec

| ID | CHID | Type | Function Jobname | ASID | Status | Alloc Time% | PCI Load | Operations Store | Rate Block | Rate Refr | -Xfer Read | Rate Write |
|------|------|-------|------------------|------|--------|-------------|----------|------------------|------------|-----------|------------|------------|
| 2601 | 0204 | RoCE2 | VTAM | 002B | Alloc | 100 | 0.050 | 0.060 | 0 | 0 | 0 | 0 |
| 2602 | 0204 | RoCE2 | VTAM | 002B | Alloc | 100 | 0.060 | 0 | 0 | 0 | 0 | 0 |

- The current RMF PCIE Activity reporting already provides statistics and performance measurements for 10GbE RoCE and 10GbE RoCE Express adapter cards.
- RMF is enhanced to recognize the new 10GbE RoCE Express2 card type and display the card in the RMF Postprocessor and Monitor III PCIE Activity reports
- The RMF support for 10GbE RoCE Express2 will be available with z/OS V2.3 RMF. New function APAR OA50762 provides the support for z/OS 2.2.

z14 - Support of Crypto Express6S Card

CRYPTO HARDWARE ACTIVITY

z/OS V2R2 SYSTEM ID S35 DATE 04/26/2017 INTERVAL 29.59.848
 RPT VERSION V2R2 RMF TIME 09.00.00 CYCLE 1.000 SECONDS

----- CRYPTOGRAPHIC CCA COPROCESSOR -----

| TOTAL | | | | KEY-GEN | |
|-------|----|------|-----------|---------|------|
| TYPE | ID | RATE | EXEC TIME | UTIL% | RATE |
| CEX5C | 0 | 1882 | 0.206 | 38.8 | 0.25 |
| | 1 | 1927 | 0.201 | 38.8 | 0.24 |
| CEX6C | 5 | 3479 | 0.111 | 38.6 | 0.46 |
| | 6 | 3235 | 0.120 | 38.7 | 0.45 |
| | 7 | 3504 | 0.110 | 38.7 | 0.43 |


----- CRYPTOGRAPHIC PKCS11 COPROCESSOR -----

| TOTAL | | | | OPERATIONS DETAILS | | | | |
|-------|----|-------|-----------|--------------------|---------------|-------|-----------|-------|
| TYPE | ID | RATE | EXEC TIME | UTIL% | FUNCTION | RATE | EXEC TIME | UTIL% |
| CEX5P | 3 | 549.6 | 0.611 | 33.6 | ASYM FAST | 289.4 | 0.575 | 16.6 |
| | | | | | ASYM GEN | 0.10 | 1.730 | 0.0 |
| | | | | | ASYM SLOW | 219.3 | 0.651 | 14.3 |
| | | | | | SYMM COMPLETE | 4.45 | 0.542 | 0.2 |
| | | | | | SYMM PARTIAL | 36.27 | 0.663 | 2.4 |
| CEX6P | 8 | 659.2 | 0.508 | 33.5 | ASYM FAST | 331.3 | 0.464 | 15.4 |
| | | | | | ASYM GEN | 0.13 | 1.203 | 0.0 |
| | | | | | ASYM SLOW | 260.0 | 0.552 | 14.4 |
| | | | | | SYMM COMPLETE | 5.39 | 0.291 | 0.2 |
| | | | | | SYMM PARTIAL | 62.41 | 0.571 | 3.6 |

----- 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% |
| CEX6A | 4 | 13504 | 0.029 | 38.7 | 1024 | 4361 | 0.010 | 4.4 | 1439 | 0.034 | 4.9 |
| | | | | | 2048 | 1633 | 0.026 | 4.2 | 0.00 | 0.000 | 0.0 |
| | | | | | 4096 | 6072 | 0.042 | 25.2 | 0.00 | 0.000 | 0.0 |

Crypto Express6 Card reported with same set of measurements as Crypto Express5



- The new Crypto Express6S card available on z14 is functionally equivalent to Crypto Express5S
- Crypto Express6S may be configured as:
 - An accelerator (CEX6A), a CCA coprocessor (CEX6C), or an EP-11 coprocessor (CEX6P)
- RMF for z/OS 2.3 and new function APAR OA50693 enhances RMF to recognize and exploit performance data for the new Crypto Express6S (CEX6) card.
- PTFs are available for z/OS 2.1 and z/OS 2.2.
- In detail, RMF collects and reports performance measurements for operations executed on Crypto Express6S CCA coprocessors (CEX6C) , PKCS11 coprocessors (CEX6P) and accelerators (CEX6A).
- The crypto measurements are stored to SMF 70 subtype 2 data sections:
 - CEX6C measurements are stored in the Cryptographic Coprocessor data section,
 - CEX6P measurements are stored in the Cryptographic PKCS11 Coprocessor data section and
 - CEX6A 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.

RMF Enhancements at a Glance


- IBM z14 Support
 - Monitoring of Virtual Flash Memory
 - RMF support for for Crypto Express6S
 - IBM 10GbE RoCE Express2 support



- DS8880 Support
 - RMF Monitoring of SuperPAV capability
 - RMF Enhancements for zHyperWrite
 - Monitor Performance of PCIE Synchronous I/O links

- RMF Enhancements for Coupling Facility

- Monitor Throughput and Delays of Asynchronous CF duplexing for lock structures
- Predict overhead of CF encryption and report encryption state of CF structures



- RMF Enhancements for Spark

- Report on new WLM functionality for Spark

- RMF WLM Reporting Enhancements

- RMF support for Shorter WLM Response Time Goals



- Large Page Enhancements

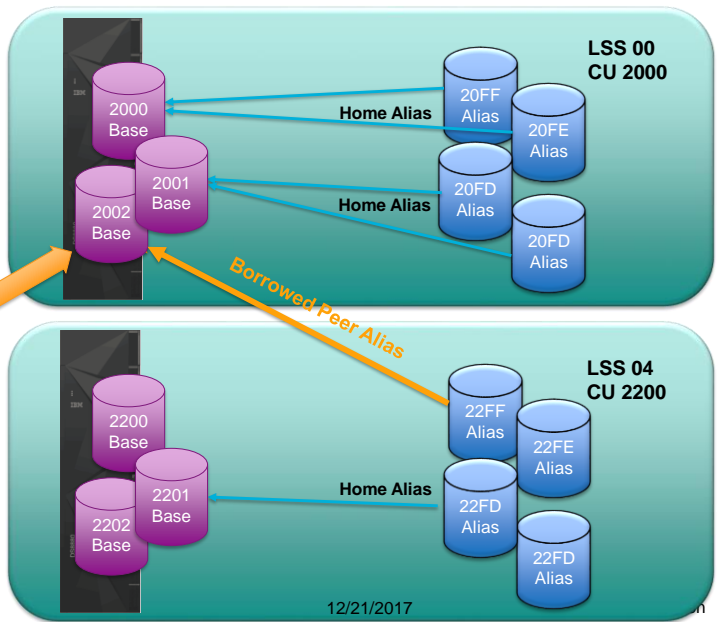
- Statistics for 2 GB Large Pages

- SuperPAV is a new functionality of the IBM DS8880 series that allows to share Aliases among multiple control units. RMF supports the monitoring of the DS8880 SuperPAV capability with new statistics in SMF records and report enhancements in the RMF Postprocessor I/O Queuing Activity report.
- The zHyperWrite technology enables DB2 to perform parallel log writes to PPRC primary and secondary volumes. This improves both DB2 log throughput and transactional response time. RMF is enhanced to monitor I/O to PPRC secondary devices.
- New RMF functionality helps to monitor the performance of PCIE Synchronous I/O links provided by the zHyperLink Express I/O feature.



DS8880 – 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.

DS8880 – 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 TA0 | | | | DATE 03/31/2016 | | | | INTERVAL 09.59.995 | | | | | |
| | | | RPT VERSION V2R2 RMF | | | | TIME 11.40.00 | | | | CYCLE 1.000 SECONDS | | | | | |
| ALIAS MANAGEMENT GROUPS | | | | | | | | | | | | | | | | |
| DATA | AMG | DCM | GROUP | CHAN | CHPID | % DP | % CU | CUB | CMR | CONTENTION | Q | CSS | HPAV | AVG | OPEN | XFER |
| | | MIN | MAX | DEF | TAKEN | BUSY | BUSY | DLY | DLY | RATE | LNGLTH | DLY | WAIT | MAX | 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 | 7.50 7.23 |
| 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 | 7.65 7.38 |
| 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.

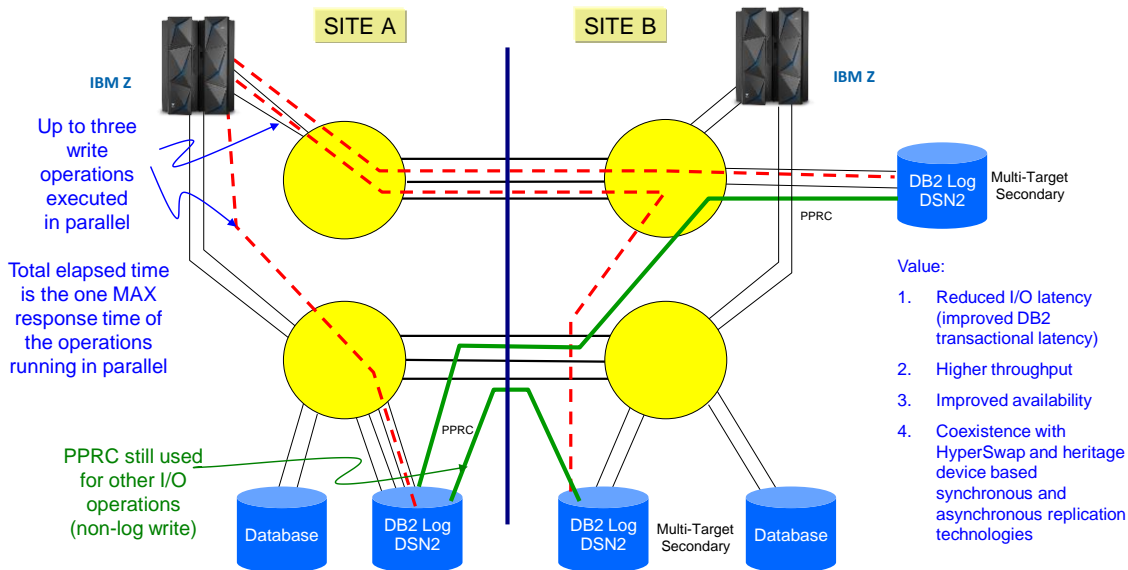
SuperPAV – RMF Postprocessor I/O Queuing Activity Report

| 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 | | LNTH | 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)

- 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.

IBM zHyperWrite - Overview

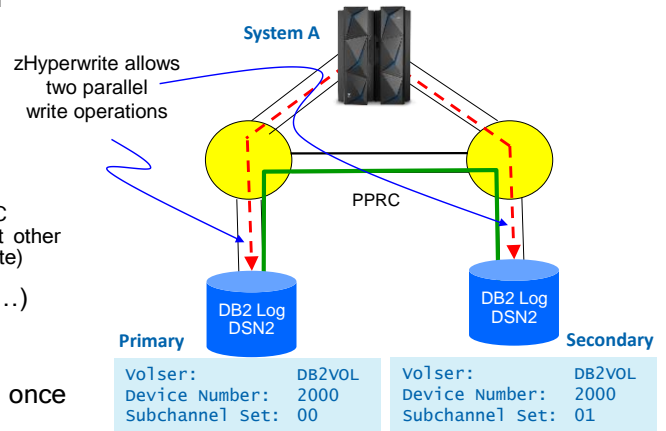


- IBM zHyperWrite storage technology provides I/O performance improvements for critical applications and middleware by providing parallel execution of I/O requests when synchronous replication is used, such as IBM DS8870 Metro Mirror technology.
- DB2 for z/OS is the first middleware that can exploit the zHyperWrite technology.
- The zHyperWrite technology speeds the DB2 log writes by eliminating a majority of the time required for the synchronous replication of data. It enables DB2 to perform parallel log writes to PPRC primary and secondary volumes.
- This improves both DB2 log throughput and transactional response time.

RMF and IBM zHyperWrite

Rationale

- In the past, RMF only reported on devices defined in subchannel set 0
- Enhancements:
 - OA42068 (z/OS 2.1) OA40977 (z/OS 1.13 and 1.12): Report on PPRC Primary devices and LCUs, where the devices are defined in a subchannel set other than 0.
 - OA45985 (z/OS 2.1): SMF 74-1 enhanced to provide performance data for PPRC Secondary devices that are defined in a subchannel set other than 0 and are eligible for read/write activity (zHyperwrite)
- Full RMF support (Postprocessor, Monitor III, DDS, ...) was still missing
- IBM zHyperwrite allows scenarios where the same four-digit device number can be active more than once in a z/OS system.
 - Reporting of four-digit device numbers is not sufficient anymore.



➔ **RMF is enhanced to collect and report 5-digit device numbers (including the subchannel set)**

- In the past, RMF only reported on devices defined in subchannel set 0.
- To relief the four digit device number constraint, PPRC secondary devices can be defined in the alternate subchannel set (subchannel set other than 0)
- With various enhancements RMF provides partial support for PPRC primary and secondary devices defined in the alternate subchannel set.
- However, full RMF support to monitor I/O activity of devices defined to an alternate subchannel set was still missing.
- IBM zHyperwrite allows I/O activity from the z/OS system to PPRC primary and secondary devices. zHyperwrite allows environments where PPRC primary and secondary devices have the same device number but different subchannel sets. Since RMF only reports 4 digit device numbers (without subchannel set information), it's not possible to monitor these environments with existing RMF reports.
- With RMF for z/OS 2.3, RMF is enhanced to collect and report the 4 digit device numbers together with the subchannel set id as a 5 digit device number. This allows to report on PPRC devices where PPRC primary and secondary has the same device number.

IBM zHyperWrite – Overview RMF support

- RMF for z/OS 2.3 is enhanced to support IBM zHyperwrite environments and five-digit device numbers across all RMF monitoring and reporting components:
 - RMF Postprocessor Device Activity, Shared Device Activity, Cache Activity, and PAGESP reports
 - Monitor I data gathering option DEVICE(NMBR(...))
 - RMF Postprocessor control statements REPORTS and SYSRPTS
 - Monitor II report commands and data gathering options DEV and DEVV
 - Monitor II DEV, DEVV, LLI, PGSP, and SENQR reports
 - Monitor III DEVR, DEVT, CACHDET, DSNJ, DSNV, and JOB reports
 - Data collection of more than 65535 active devices in SMF 74 subtype 1
 - Subchannel set id is stored in SMF 74 subtypes 5 and 8 as well as in SMF 75 and SMF 79 subtypes 6 and 11
 - Enhanced RMF programming interfaces ERBSMFI and ERB2XDGS
- In general, five-digit device numbers have the format *sdddd* where *s* represents a subchannel set id between 0 and 3 and *dddd* represents the “old” four-digit device number

z/OS V2.3

- z/OS V2.3 RMF will be enhanced to support monitoring and reporting of system environments with more than 65535 DASD devices.
- The new support allows to specify five-digit device numbers in RMF data gathering and reporting options whereby the first digit of the five-digit device number represents the ID of the subchannel set to which the device is physically configured.
- The RMF device and storage subsystem reporting is enhanced to report five-digit device numbers.

RMF Monitor I Cache Subsystem Activity

RMF Postprocessor Cache Subsystem Activity report is enhanced to display 5-digit device numbers in report sections

- Top-20 Device Lists
- Cache Subsystem Device Overview
- Cache Device Activity

Top-20 Device Lists

Device List by DASD I/O Rate

| Volume Serial | Device Number | Subsystem ID | Caching Status | %I/O of Total | I/O Rate | Cache Hit Read Rate |
|---------------|---------------|--------------|----------------|---------------|----------|---------------------|
| D83F02 | 0AF02 | A223 | Active | 1.8 | 83.4 | 0.7 |
| D83F0A | 0AF2F | A223 | Active | 16.8 | 795.4 | 735.3 |
| D83F27 | 0AF27 | A223 | Active | 1.0 | 44.9 | 0.6 |
| D83F04 | 0AF04 | A223 | Active | 1.4 | 68.1 | 0.6 |
| D83F07 | 0AF07 | A223 | Active | 16.5 | 780.9 | 735.2 |
| D83F2C | 0AF2C | A223 | Active | 1.1 | 54.1 | 0.5 |
| D83F2A | 0AF2A | A223 | Active | 1.0 | 47.4 | 0.5 |

Cache Device Activity for Volume Serial D83F00

Storage Subsystem Descriptor

Subsystem : 2107-01 Physical CU-ID : AF18 Subsystem ID : A223 Type-Model : 2107-941
 Volume Serial : D83F00 **Device Number : 0AF00** Extent Pool ID : 0001

Cache Device Status

Caching : Active DASD Fast Write : Active Pinned Data : None Duplex Pair : Not established

Cache Device Activity

| Cache I/O Request Category | Read Count |
|----------------------------|------------|
| Normal | 300 |
| Sequential | 48 |
| CFW Data | 0 |
| Total | 348 |

Cache Subsystem Device Overview

| Volume Serial | Device Number | Extent Pool ID | Caching Status |
|---------------|---------------|----------------|----------------|
| *ALL | | | |
| *CACHE-OFF | | | |
| *CACHE | | | |
| D83F00 | 0AF00 | 0001 | Active |
| D83F01 | 0AF01 | 0001 | Active |
| D83F02 | 0AF02 | 0001 | Active |
| D83F03 | 0AF03 | 0001 | Active |

- The RMF Postprocessor Cache Subsystem Activity report is enhanced to display five-digit device numbers in the following report sections:
 - Top-20 Device List by DASD I/O Rate and Top-20 Device List by total I/O Rate
 - Cache Subsystem Device Overview
 - Cache Device Activity

RMF Monitor I Device Activity

RMF Postprocessor Device Activity report displays 5-digit device numbers

▼ Direct Access Device Activity

Total Samples : 600 IODF Name Suffix : 20 IODF Creation Date : 03/31/2017 IODF Creation

| Storage Group | Device Number | Device Type | Number of Cylinders | Volume Serial Number | Available PAV Devices | HyperPAV | LCU Number | Device Activity Rate | | | | |
|---------------|---------------|-------------|---------------------|----------------------|-----------------------|----------|------------|----------------------|-------|-------|-------|-------|
| EFDMPK | 0AF00 | 33909 | 65520 | D83F00 | 1.0 | Yes | 008C | 0.002 | | | | |
| EFDMPK | 0AF01 | 33909 | 65520 | D83F01 | 1.0 | Yes | 008C | 0.002 | | | | |
| EFDMPK | 0AF02 | 33909 | 65520 | D83F02 | 1.0 | Yes | 008C | 0.002 | 0.128 | 0.000 | 0.000 | 0.000 |
| EFDMPK | 0AF03 | 33909 | 65520 | D83F03 | 1.0 | Yes | 008C | 0.002 | 0.384 | 0.000 | 0.128 | 0.000 |
| EFDMPK | 0AF04 | 33909 | 65520 | D83F04 | 1.0 | Yes | 008C | 0.002 | 0.384 | 0.000 | 0.000 | 0.000 |
| EFDMPK | 0AF05 | 33909 | 65520 | D83F05 | 1.0 | Yes | 008C | 0.002 | 0.256 | 0.000 | 0.000 | 0.000 |
| EFDMPK | 0AF06 | 33909 | 65520 | D83F06 | 1.0 | Yes | 008C | 0.002 | 0.512 | 0.000 | 0.000 | 0.000 |
| EFDMPK | 0AF07 | 33909 | 65520 | D83F07 | 1.0 | Yes | 008C | 0.002 | 0.000 | 0.000 | 0.000 | 0.000 |
| LCU | | | | | | | 008C | 0.013 | 0.320 | 0.000 | 0.016 | 0.000 |

RMF Postprocessor Device Activity reporting:

REPORTS control statement allows to specify a 5-digit device number for **DEVICE** report

Example:
`REPORTS(DEVICE(NMBR(0AF00:0AF07)))`

- Monitor I Device Data gathering:
 - Gatherer option **DEVICE** allows to specify 5-digit device numbers
 - Example: `DEVICE(NMBR(10288,10291))`
requests monitoring of devices 0288 and 0291 configured to subchannel set 1

- The RMF Postprocessor Device Activity and Shared Device Activity reports are enhanced to display five-digit device numbers.
- RMF Monitor I device data gathering:
 - Monitor I data gatherer option `DEVICE|NODEVICE` controls the device data gathering.
 - `DEVICE` option enhanced to specify 5 digit device number for a device range or list of devices that should be monitored.
 - Example: `DEVICE(NMBR(10288,10291))`
-> requests monitoring of devices 0288 and 0291 configured to subchannel set 1
- RMF Postprocessor Device Activity report:
 - `REPORTS` control statement allows to specify a 5-digit device number when using `NMBR` suboption of `DEVICE` report.:
 - The device number can be specified in format `[s]dddd`. The first digit `s` represents an optional subchannel set ID. If the subchannel set ID is omitted data reporting for devices `0dddd`, `1dddd`, `2dddd`, and `3ddd` is requested.
 - Example: `REPORTS(DEVICE(COMM,NMBR(12180,12183:12185,2188)))` generates report data for all communication devices as well as for devices 2180, 2183, 2184, 2185 configured to subchannel set 1 and devices 2188 configured to any subchannel set
- RMF Postprocessor Shared Device Activity report:
 - `SYSRPTS` control statement allows to specify a 5-digit device number when using `NMBR` or `EXNMBR` suboptions of `SDEVICE` report:
 - The five-digit device number is specified in the format `[s]dddd`. The first digit `s` represents an optional subchannel set ID. If the subchannel set ID is omitted device activity reporting is requested (suboption `NMBR`) or suppressed (suboption `EXNMBR`) for devices `0dddd`, `1dddd`, `2dddd`, and `3ddd`.
 - Example: `SYSRPTS(SDEVICE(EXNMBR(02180,02183:02184)))` will exclude devices 2180, 2183, 2184 configured to subchannel set 0 from Shared Device Activity report.
- In addition the RMF Postprocessor Page Data Set Activity report is also enhanced to display 5-digit device numbers.

RMF Monitor II Device Activity

RMF - DEV Device Activity Line 67 of 18596

CPU= 15/ 1 UIC= 65K PR= 0 System= N64 Total

09:17:38 I=76% **DEV** ACTV RESP IOSQ -DELAY- PEND DISC CONN %D %D
 STG GRP VOLSER NUM PAV LCU RATE TIME TIME CMR DB TIME TIME TIME UT RV

| | | | | | | | | | | | | | |
|--------|-------|------|-------|-------|------|------|-----|------|------|------|------|----|---|
| H10900 | 00900 | 1.0H | 000F | 35.95 | 3.24 | .001 | .02 | .00 | .133 | 1.17 | 1.93 | 11 | 0 |
| H10901 | 00901 | 1.0H | 000F | 36.01 | 3.22 | .001 | .02 | .00 | .133 | 1.17 | 1.92 | 11 | 0 |
| H10902 | 00902 | 1.0H | 000F | 35.95 | 3.23 | .000 | .02 | .00 | .133 | 1.17 | 1.92 | 11 | 0 |
| H10903 | 00903 | 1.0H | 000F | 35.97 | 3.23 | .000 | .02 | .00 | .132 | 1.17 | 1.92 | 11 | 0 |
| 10901 | 1.0H | 00CD | 18.48 | 9.84 | .240 | .01 | .00 | .124 | .000 | 9.47 | 18 | 0 | |
| 10902 | 1.0H | 00CD | 18.39 | 9.76 | .124 | .01 | .00 | .124 | .000 | 9.47 | 18 | 0 | |
| 10903 | 1.0H | 00CD | 18.40 | 9.64 | .000 | .01 | .00 | .124 | .000 | 9.47 | 18 | 0 | |

RMF Monitor II - Device Activity Options - Single Device

Volume ==> _____ Specify a volume serial number.
 Device Number ==> 00900 Specify a hexadecimal device number.

zHyperwrite I/O activity to PPRC Secondaries

The DEV Report Option

The DEVV Report Option

RMF - DEVV Device Activity Line 1 of 5

CPU= 15/ 1 UIC= 65K PR= 0 System= N64 Total

TIME I=76% **DEV** ACTV RESP IOSQ -DELAY- PEND DISC CONN %D %D
 VOLSER NUM PAV LCU RATE TIME TIME CMR DB TIME TIME TIME UT RV

| | | | | | | | | | | | | | |
|--------|-------|------|------|-------|------|------|-----|-----|------|------|------|----|---|
| H10900 | 00900 | 1.0H | 000F | 35.95 | 3.24 | .001 | .02 | .00 | .133 | 1.17 | 1.93 | 11 | 0 |
|--------|-------|------|------|-------|------|------|-----|-----|------|------|------|----|---|

© 2017 IBM Corporation

- RMF Monitor II Device Activity reports DEV and DEVV are enhanced to display 5-digit device numbers
- PPRC secondary devices are usually offline. In the Monitor II device reports, these devices are reported without volser.
- The report option panel of DEV report allows to specify a single 5-digit device number and a list or range of 5-digit device numbers
- The report option panel of DEVV report allows to specify a single 5-digit device number
- Monitor II device activity reports are based on the data in SMF 79-9. Since SMF79-9 is limited to a maximum number of 65535 active devices, the reports cannot show more than 65535 devices.
- If RMF Monitor II detects more than 65535 active devices in a IBM zHyperwrite environment
 - RMF Monitor II Device background session writes error message ERB439I
 - RMF Monitor II Device Activity report DEV displays error message ERBA092I
- To avoid this problem, it is recommended to specify a device number range that does not encompass more than 65535 devices (e.g. 00000:0FFFF).
- Further RMF Monitor II reports that are enhanced to display 5-digit device numbers:
 - Library List report (LLI)
 - Page Data Set Activity (PGSP)
 - System Enqueue Reserve (SENQR)

RMF Monitor III Cache Subsystem and Device Activity

RMF V2R3 Cache Detail - SVPLEX1 Line 44 of 496

Samples: 60 Systems: 4 Date: 05/02/17 Time: 09.50.00 Range: 60 Sec
 CDate: 05/02/17 CTime: 09.49.56 CRange: 60 Sec

| Volume /Num | SSID | I/O % | I/O Rate | Hit % | Cache Read | Hit Rate | DFW | CFW | DASD Total | I/O Stage | Seq Rate | Async Rate |
|-------------|-------|-------|----------|-------|------------|----------|------|-----|------------|-----------|----------|------------|
| H10900 | 00900 | 1903 | 6.2 | 1183 | 100 | 0.0 | 1183 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| 20900 | 1983 | 6.2 | 1201 | 100 | 0.0 | 1201 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| 10900 | 1943 | 6.2 | 1201 | 100 | 0.0 | 1201 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |

RMF Monitor III Cache Detail report

RMF Monitor III Device Resource Delays report

zHyperwrite I/O activity to PPRC Secondaries

RMF V2R3 Device Resource Delays Line 21 of 70

Samples: 60 System= N64 Date: 05/02/17 Time: 09.50.00 Range: 60 Sec

| Volume /Num | S/PAV | Act Rate | Resp Time | ACT % | CON % | DSC % | PND % | Reasons | DEV/CU Type | Jobname | Service C Class | USG % | DLY % |
|-------------|-------|----------|-----------|-------|-------|-------|-------|---------|-------------|---------|-----------------|-------|-------|
| H10A01 | S | 17.8 | 8.3 | 15 | 15 | 0 | PND | 0 | 33903 2107 | ZHWWKL1 | B JESLOW5 | 17 | 0 |
| 00A01 | 1.0H | | | | | | | | | | | | |
| H10909 | S | 18.5 | 8.4 | 15 | 15 | 0 | PND | 0 | 33903 2107 | ZHWWKLD | B JESMED4 | 17 | 0 |
| 00909 | 1.0H | | | | | | | | | | | | |

- In RMF Monitor III, the following reports are enhanced to display display 5-digit device numbers:
 - Cache Detail report (CACHEDET)
 - Device Resource Delays report (DEVR)
 - Device Activity Trend report (DEVT)
 - Data Set Delays – Job report (DSNJ)
 - Data Set Delays – Volume report (DSNV)
 - Job Delay report (JOB)

New SMF 74 subtype 1 logical record concept

- One SMF 74 subtype 1 record can only hold up to 65535 ('FFFF'x) devices
 - Two-byte field SMF74DDN describes number of device data sections in SMF record
 - Sufficient for reporting of four-digit device numbers
 - One SMF 74 subtype 1 record written per device class
- With zHyperwrite, the same four-digit device number can be active more than once in a z/OS system, e.g. devices 02080 and 12080
 - ➔ More than 65535 devices can be active on a z/OS system
- RMF now supports multiple SMF 74 subtype 1 logical records per device class
- New logical SMF record sequence number in SMF 74-1 Device Control Data Section

SMF record type 74 subtype 1 – Device Control Data Section

| Offsets | Name | Len | Format | Description |
|---------|----------|-----|--------|--|
| 11 B | SMF74SMF | 1 | Binary | Logical SMF record flag Bit Meaning when Set 0 There are more logical SMF records for this device class 1-7 Reserved. |
| 27 1B | SMF74LSN | 1 | Binary | Logical SMF record sequence number within a device class. This number can be used by SMF record assembly programs to recognize conditions where logical SMF records of a device class are not sorted in chronological order. |

- One SMF 74-1 record can only hold data for up to 65535 devices.
- To support more than 65535 active devices, the RMF Monitor I data gatherer was changed to create multiple SMF 74-1 records per device type (e.g. device type DASD). Each logical SMF 74-1 record has a logical SMF record sequence number that can be used by programs to determine the chronological order of logical SMF records belonging to same device class.


Exploitation Considerations

- Application programs exploiting SMF 74 subtype 1 data in IBM zHyperwrite environments must handle the condition that performance data for more than 65535 devices are reported
- If such a condition is indicated in the Device Control data section (SMF74SMF/SMF74LSN) of SMF 74 subtype 1, additional SMF record(s) must be processed by the application program if all devices of a device class are of interest to the program

IBM zHyperLink Express - Overview

What is IBM zHyperLink™?

- zHyperLink Express is a direct connect short distance IBM Z I/O feature designed to work in conjunction with a FICON or High Performance FICON SAN infrastructure
- IBM zHyperLink™ dramatically reduces latency by interconnecting the z14 CPC directly to the I/O Bay of the DS8880
- zHyperLink is FAST enough the CPU can just wait for the data
 - No Un-dispatch of the running task
 - No CPU Queueing Delays to resume it
 - No host CPU cache disruption
 - Very small I/O service time
- Operating System and Middleware (e.g. DB2) are changed to keep running over an I/O



- IBM zHyperLink Express is a short distance, IBM Z I/O adapter which works in conjunction with a FICON SAN infrastructure. The feature is a 2 port PCIe adapter which resides in the z14 PCIe I/O drawer and supports direct connectivity to the I/O Bay of the DS8880 at distances up to 150m.
- IBM zHyperLink offers extremely low latency connectivity to FICON storage systems. Working in conjunction with existing FICON SAN infrastructure, zHyperLink fosters a new I/O paradigm for IBM mainframes.
- zHyperLink improves application response time, cutting I/O sensitive workload response time by up to 50% without requiring application changes.
- zHyperLink is intended to speed DB2 for z/OS transaction processing and later improve DB2 active Log throughput.

RMF Postprocessor PCIE Activity report

- RMF **Postprocessor PCIE Activity** report enhanced to provide performance metrics for synchronous I/O over IBM zHyperLink

OA50755

Number of megabytes per second transferred to and from the PCIE function.

IBM zHyperLink

General PCIE Activity

| Function ID | Function CHID | Function Name | Function Status | Owner Job Name | Owner Address Space ID | Function Allocation Time | PCI Load Operations Rate | PCI Store Operations Rate | PCI Store Block Operations Rate | Refresh PCI Translations Operations Rate | Data Transfer Rate |
|-------------|---------------|----------------|-----------------|----------------|------------------------|--------------------------|--------------------------|---------------------------|---------------------------------|--|--------------------|
| 0300 | 01C8 | 8GB zHyperLink | Allocated | IOSAS | 0019 | 600 | | | | | 3.24 |
| 0301 | 01C8 | 8GB zHyperLink | Allocated | IOSAS | 0019 | 600 | | | | | 3.24 |
| 0302 | 01C8 | 8GB zHyperLink | Allocated | IOSAS | 0019 | 600 | | | | | 3.24 |
| 0303 | 01C8 | 8GB zHyperLink | Allocated | IOSAS | 0019 | 600 | | | | | 3.24 |
| 0304 | 01C8 | 8GB zHyperLink | Allocated | IOSAS | 0019 | 600 | | | | | 3.24 |

- RMF is enhanced to monitor and report on performance of :
 - PCIE Synchronous I/O links, defined by a zHyperlink port of a server, a fibre optic cable and the target port on a connected storage controller
 - DASD devices that actively perform synchronous I/O reads and writes
- The RMF support for zHyperLink will be available with z/OS V2.3 RMF. New function APAR OA50755 provides the support for z/OS 2.1 and z/OS 2.2.
- RMF collects and reports new zHyperlink synchronous I/O statistics in following reports:
 - RMF Postprocessor PCIE Activity report
 - RMF Postprocessor Device Activity report
 - RMF Monitor III PCIE report.
- The synchronous I/O statistics are stored in SMF 74-1 and SMF 74-9
- New/modified OVW Conditions are provided for synchronous I/O metrics based on the data in SMF 74-1 and 74-9 records.
- The **RMF Postprocessor PCIE Activity** report is enhanced to provide performance metrics for synchronous I/O over IBM zHyperLink.
- In the General PCIE Activity section of the report zHyperLinks that were active on the system are displayed as PCIe function name '8GB zHyperLink'
- Together with the new support for zHyperLinks, the Postprocessor PCIE Activity report is restructured:
 - Function specific metrics are moved from General PCIE activity section to the Hardware Accelerator sections and the new ROCE and ISM report sections.
 - The Read and write transfer rates are combined in a single 'Data Transfer Rate' value.

RMF Postprocessor PCIE Activity report ...

Connectivity data and request performance measurements on LPAR and CPC level

▼ Synchronous I/O Link Activity

| Function ID | Function CHID | Port ID | Serial Number | Type-Model | Total Request Rate | Total Request Rate (CPC) | Successful Request % | Successful Request % (CPC) | Read Transfer Rate | Read Transfer Rate (CPC) | Read Transfer Ratio | Read Transfer Ratio (CPC) | Write Transfer Rate | Write Transfer Rate (CPC) | Write Transfer Ratio | Write Transfer Ratio (CPC) | Time Busy % | Time Busy % (CPC) |
|-------------|---------------|---------|---------------|------------|--------------------|--------------------------|----------------------|----------------------------|--------------------|--------------------------|---------------------|---------------------------|---------------------|---------------------------|----------------------|----------------------------|-------------|-------------------|
| 0300 | 01C8 | 1 | 0000000YT111 | 002107-981 | 791 | 6277 | 100 | 100 | 3.24 | 25.7 | 0.004 | 0.004 | 0 | 0 | 0 | 0 | 1.41 | 11.4 |
| 0301 | 01C8 | 1 | 0000000YT111 | 002107-981 | 791 | 6277 | 100 | 100 | 3.24 | 25.7 | 0.004 | 0.004 | 0 | 0 | 0 | 0 | 1.41 | 11.4 |
| 0302 | 01C8 | 1 | 0000000YT111 | 002107-981 | 791 | 6277 | 100 | 100 | 3.24 | 25.7 | 0.004 | 0.004 | 0 | 0 | 0 | 0 | 1.41 | 11.4 |
| 0303 | 01C8 | 1 | 0000000YT111 | 002107-981 | 791 | 6277 | 100 | 100 | 3.24 | 25.7 | 0.004 | 0.004 | 0 | 0 | 0 | 0 | 1.41 | 11.4 |
| 0304 | 01C8 | 2 | 0000000YT111 | 002107-981 | 791 | 6277 | 100 | 100 | 3.24 | 25.7 | 0.004 | 0.004 | 0 | 0 | 0 | 0 | 1.57 | 12.6 |

▼ Synchronous I/O Response Time Distribution

| Function ID | % Read < 20usec | % Read < 30usec | % Read < 40usec | % Read < 50usec | % Read < 60usec | % Read < 70usec | % Read < 80usec | % Read < 90usec | % Read < 100usec | % Read >=100usec |
|-------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|------------------|------------------|
| 0300 | 82.5 | 17.5 | 0.026 | 0.003 | 0.004 | 0.003 | 0.002 | <.001 | 0 | <.001 |
| 0301 | 82.4 | 17.5 | 0.026 | 0.009 | 0.004 | 0.004 | 0.003 | <.001 | 0 | <.001 |
| 0302 | 82.4 | 17.6 | 0.031 | 0.006 | 0.006 | 0.005 | 0.003 | <.001 | <.001 | <.001 |
| 0303 | 82.3 | 17.7 | 0.023 | 0.007 | 0.003 | 0.003 | 0.001 | <.001 | 0 | 0.001 |
| 0304 | 22.3 | 77.6 | 0.082 | 0.008 | 0.004 | 0.004 | 0.003 | <.001 | 0 | 0.001 |

I/O Response Time Distribution for instructions executed synchronously on zHyperLinks.

- There are two new report sections to report PCIe function specific measurements for zHyperLinks:
 - Synchronous I/O Link Activity report section
 - Synchronous I/O Response Time Distribution report section
- The Synchronous I/O Link Activity section has either system scope (metrics showing values per function) or CEC scope (metrics showing values with a CPC wide view). For IBM zHyperLinks the section displays:
 - Interconnection metrics. These are metrics like port ID, serial number, type and model of the storage controller the synchronous I/O link is connected to
 - Transfer metrics,. For example: data transfer rate
 - Metrics on requests processed, like percentage of successful requests
 - Utilization metrics, like time busy percentage.
- Values on CPC level are only reported if Global Performance Reporting is enabled in the LPAR image profile of the Hardware Management Console (HMC).
- The Synchronous I/O Response Time Distribution section provides a response time distribution overview for read and write instructions executed on the allocated synchronous I/O links.
- The response time of synchronous I/Os are sampled in 10 buckets:

Leftmost bucket: Percentage of I/Os with a response time less than n microseconds.
 Rightmost bucket: Percentage of I/Os with a response time greater or equal n microseconds.
 Other buckets: Percentage of I/Os with a response time less than n microseconds and greater or equal to the prior bucket time limit.

Example:
 % Read <30msec = 15.3 means that 15.3 percent of the read I/Os had a response time of more than or equal to 20 microseconds but less than 30 microseconds.

RMF Postprocessor Device Activity report

| DIRECT ACCESS DEVICE ACTIVITY | | | | | | | | | | |
|-------------------------------|---------|-------------|----------------------|---------------|------------|-----------------|---------------|-----------|------|---------|
| z/OS V2R3 | | | SYSTEM ID TA3 | | | DATE 05/04/2017 | | | | |
| | | | RPT VERSION V2R3 RMF | | | TIME 01.50.00 | | | | |
| TOTAL SAMPLES = | 600 | IODF = | 67 | CR-DATE: | 04/26/2017 | CR-TIME: | 15.02 | | | |
| STORAGE GROUP | DEV NUM | DEVICE TYPE | NUMBER OF CYL | VOLUME SERIAL | PAV | LCU | ACTIVITY RATE | RESP TIME | IOSQ | AVG DLY |
| CDRIVER | 05003 | 3390A | 65667 | CD5003 | 1.0H | 0061 | 0.002 | .128 | .000 | .000 |
| CDRIVER | 05004 | 3390A | 65667 | CD5004 | 1.0H | 0061 | 0.002 | .128 | .000 | .000 |
| CDRIVER | 05005 | 3390A | 65667 | CD5005 | 1.0H | 0061 | 0.002S | .128 | .000 | .000 |
| CDRIVER | 05006 | 3390A | 65667 | CD5006 | 1.0H | 0061 | 0.002S | .128 | .000 | .000 |
| CDRIVER | 05007 | 3390A | 65667 | CD5007 | 1.0H | 0061 | 0.002S | .000 | .000 | .000 |
| CDRIVER | 05008 | 3390A | 65667 | CD5008 | 1.0H | 0061 | 0.002S | .128 | .000 | .000 |

DASD devices which actively performed synchronous I/O requests are marked by character 'S'



SYNCHRONOUS I/O DEVICE ACTIVITY report section with detailed synchronous I/O performance statistics for marked DASD devices

SYNCHRONOUS I/O DEVICE ACTIVITY

| z/OS V2R3 | | | SYSTEM ID TA3 | | | DATE 05/04/2017 | | | | | INTERVAL 09.59.999 | | | | | | | |
|--------------------------|---------|-------------|----------------------|----------|------------|-----------------|----------|---------------|-------|-------|---------------------|---------------|-------------|-----------|-----------------|-------------|------|--|
| | | | RPT VERSION V2R3 RMF | | | TIME 01.50.00 | | | | | CYCLE 1.000 SECONDS | | | | | | | |
| TOTAL SAMPLES = | 600 | IODF = | 67 | CR-DATE: | 04/26/2017 | CR-TIME: | 15.02.37 | ACT: ACTIVATE | | | | | | | | | | |
| - DEVICE ACTIVITY RATE - | | | | | | | | | | | -- AVG RESP TIME -- | | | | AVG SYNCH I/O % | | | |
| STORAGE GROUP | DEV NUM | DEVICE TYPE | VOLUME SERIAL | LCU | READ | WRITE | I/O | ASYNCH | SYNCH | I/O | ASYNCH | TRANSFER RATE | REQ SUCCESS | LINK BUSY | CACHE MISS | --REJECTS-- | | |
| CDRIVER | 05005 | 3390A | CD5005 | 0061 | 3148.02 | 0.000 | 0.002 | 0.020 | 0.000 | 0.128 | 12.89 | 0.000 | 100.0 | 0.00 | 0.00 | 0.00 | 0.00 | |
| CDRIVER | 05006 | 3390A | CD5006 | 0061 | 3148.63 | 0.000 | 0.002 | 0.020 | 0.000 | 0.128 | 12.90 | 0.000 | 100.0 | 0.00 | 0.00 | 0.00 | 0.00 | |
| CDRIVER | 05007 | 3390A | CD5007 | 0061 | 3148.21 | 0.000 | 0.002 | 0.020 | 0.000 | 0.000 | 12.90 | 0.000 | 100.0 | 0.00 | 0.00 | 0.00 | 0.00 | |
| CDRIVER | 05008 | 3390A | CD5008 | 0061 | 3149.34 | 0.000 | 0.002 | 0.020 | 0.000 | 0.128 | 12.90 | 0.000 | 100.0 | 0.00 | 0.00 | 0.00 | 0.00 | |
| | | | | LCU | 12594.2 | 0.000 | 0.007 | 0.020 | 0.000 | 0.096 | 51.59 | 0.000 | 100.0 | 0.00 | 0.00 | 0.00 | 0.00 | |

- RMF **Postprocessor Device Activity** report is enhanced to provide performance metrics for synchronous I/O over IBM zHyperLink
- In the device activity report section, devices with synchronous I/O requests are marked by character 'S' following the device activity rate.
- Detailed synchronous I/O performance statistics for the marked DASD devices are displayed in the new **SYNCHRONOUS I/O DEVICE ACTIVITY** section of the report.
 - The Synchronous I/O Device Activity section is only part of the report if at least one DASD device actively performed synchronous I/O requests.
 - For easy comparison this report lists the device's activity rate and average response time for asynchronous I/O requests adjacent to the appropriate measurements for synchronous I/O read and write requests and provides additional rates and percentages related to synchronous I/O processing.
- A device with synchronous I/O activity may be mapped back to the synchronous I/O link by which it is reached by looking up the serial number and node descriptor information of the device's storage controller in the RMF Cache Subsystem Device Overview report. Serial number and type-model of the storage controller can be looked up in the Synchronous I/O Link Activity section of the RMF Postprocessing PCIE Activity report to identify the appropriate synchronous I/O link.

RMF Monitor III PCIE Activity report

- **Monitor III PCIE Activity** report enhanced to report IBM zHyperLink activities

RMF V2R3 PCIE Activity Line

Samples: 60 System: TA5 Date: 05/04/17 Time: 02.06.00 Range:

| ID | CHID | Type | Function Jobname | ASID | Status | Alloc Time% | PCI Load | Operations Store | Rate Block | Refr |
|------|------|-------|------------------|------|--------|-------------|----------|------------------|------------|------|
| 0078 | 01FC | ZEDC | FPGHWAM | 0013 | Alloc | 100 | 0 | 0 | 0 | 0 |
| 007F | 01FC | ZEDC | FPGHWAM | 0013 | Alloc | 100 | 0 | 0 | 0 | 0 |
| 0500 | 01C8 | zHypL | IOSAS | 001A | Alloc | 100 | | | | |
| 0501 | 01C8 | zHypL | IOSAS | 001A | Alloc | 100 | | | | |
| 0502 | 01C8 | zHypL | IOSAS | 001A | Alloc | 100 | | | | |
| 0503 | 01C8 | zHypL | IOSAS | 001A | Alloc | 100 | | | | |

IBM zHyperLink

RMF Synchronous I/O Link Activity

Function ID : 0500 Alloc Time % : 100
 Allocated : 22.56.47 on 05/03/17

Synchronous I/O Link
 Port ID : 1
 Type-Model : 002107-981
 Serial Number : 0000000YT111

| Adapter | This Function | Link (CEC) |
|---------------|---------------|------------|
| Time Busy % : | 1.38 | 11.0 |

Request
 Rate : 746 6022
 Success % : 100 100

Transfer
 Read Rate : 3.06 24.7
 Read Ratio : 0.004 0.004
 Write Rate : 0 0
 Write Ratio : 0 0

Cursor sensitive fields Function ID, CHID and Type display one of the following pop-ups:

- RMF Hardware Accelerator and Compression Activity
- RMF RoCE Activity
- RMF Internal Shared Memory Activity
- RMF Synchronous I/O Link Activity

- In RMF Monitor III PCIE Activity report, IBM zHyperlinks are reported as PCIe function type 'zHypL'.
- Together with the zHyperlink support, the Monitor III PCIE Activity report is enhanced by additional pop-up panels displaying function type specific PCIE connection and performance measurements.
- If cursor is placed on one of the cursor sensitive fields *Function ID*, *Function CHID* or *Function Type*, one of the following pop-up panels are displayed:
 - RMF Hardware Accelerator and Compression Activity
 - RMF RoCE Activity
 - RMF Internal Shared Memory Activity
 - RMF Synchronous I/O Link Activity
- The details for zHyperLinks are displayed in pop-up panel 'RMF Synchronous I/O Link Activity'

RMF Enhancements at a Glance


- IBM z14 Support
 - Monitoring of Virtual Flash Memory
 - RMF support for for Crypto Express6S
 - IBM 10GbE RoCE Express2 support



- DS8880 Support
 - RMF Monitoring of SuperPAV capability
 - RMF Enhancements for zHyperWrite
 - Monitor Performance of PCIe Synchronous I/O links

RMF Enhancements for Coupling Facility

- Monitor Throughput and Delays of Asynchronous CF duplexing for lock structures
- Predict overhead of CF encryption and report encryption state of CF structures



- RMF Enhancements for Spark
 - Report on new WLM functionality for Spark


- RMF WLM Reporting Enhancements
 - RMF support for Shorter WLM Response Time Goals



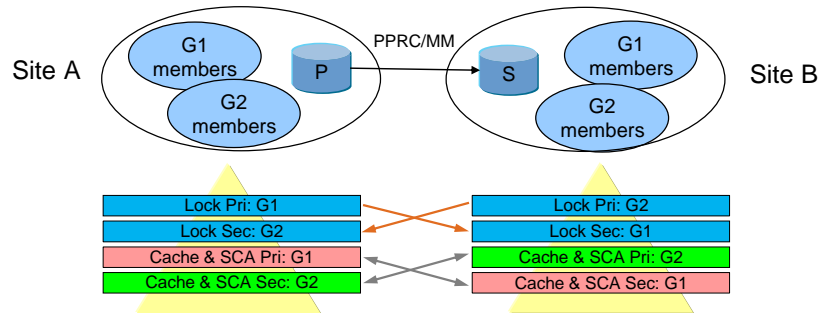
- Large Page Enhancements
 - Statistics for 2 GB Large Pages

- Asynchronous CF duplexing for lock structures is a solution for continuous availability. RMF provides new statistics that helps to monitor throughput and delays of asynchronous CF duplexing for lock structures.
- z/OS V2.3 provides support for end-to-end encryption for both CF data in flight and data at rest in CF structures. RMF is enhanced to report the encryption state per CF structure and collects new CF structure statistics that can be used to predict the overhead of CF encryption and decryption.

Asynchronous System Managed CF Duplexing for Lock Structures





- **Secondary structure updates are performed asynchronously with respect to primary updates**
 - Designed to drive out cross-site latencies that exist today when replicating CF data across distance
 - Designed to avoid the need for synchronous speed-of-light communication delays during the processing of every duplexed update operation
- **Value**
 - Improves performance with cross-site duplexing of lock structures at distance
 - Maintains robust failure recovery capability through the redundancy of duplexing
 - Reduces z/OS, CF, and link utilization overhead costs associated with synchronous duplexing of lock structures
- **Requirements**
 - IBM z13 GA2 and z13s, CFCC Level 21, z/OS V2.2 with PTFs for APAR OA47796 and OA49148
 - First Exploiter: IRLM/DB2 Lock Structure; IRLM version 2.3 with PTF for PI68378 DB2 V12 with PTF for PI66689



- Asynchronous CF duplexing for lock structures is designed to be a continuously available solution that:
 - Makes duplexing Coupling Facility (CF) lock structures more practical, even at extended distances
 - Provides performance advantages for duplexing lock structures
- Asynchronous CF duplexing is designed to improve the performance of the duplexing of lock structures by reducing the overhead costs associated with the synchronous duplexing of lock structures while maintaining the robust failure recovery capability inherent in the redundancy of duplexing.
- With asynchronous CF duplexing, secondary CF lock structure updates are performed asynchronously with respect to primary updates in order to drive out cross-CF latencies that exist today. This is intended to avoid the need for synchronous communication delays during the processing of every duplexed update operation.
- When the coupling facility (CF) is set up in duplex mode, DB2 for z/OS data sharing performance can improve if asynchronous CF duplexing is in effect.
- With asynchronous CF duplexing, multi-site data sharing groups that implement duplexing of all DB2 CF structures can have continuous availability without significant performance impact after site failures.

RMF Postprocessor CF Activity report



COUPLING FACILITY NAME = CF03

COUPLING FACILITY USAGE SUMMARY

GENERAL STRUCTURE SUMMARY

| TYPE | STRUCTURE NAME | STATUS CHG | ALLOC SIZE | % OF CF STOR | # REQ | % OF ALL REQ | % OF CF UTIL | AVG REQ/SEC | LST/DIR ENTRIES TOT/CUR | DATA ELEMENTS TOT/CUR | LOCK ENTRIES TOT/CUR | DIR REC/DIR REC XI'S |
|------|-----------------|------------------|------------|--------------|-------|--------------|--------------|-------------|-------------------------|-----------------------|----------------------|----------------------|
| LOCK | ISGLOCK | ACTIVE | 47M | | | 7.4 | 0.7 | 5766.7 | 0 | 0 | 4194K | N/A |
| | THRLCKDB2_ASYNC | ACTIVE SEC A | 220M | | | 0.0 | 0.7 | 0.00 | 0 | 0 | 2048 | N/A |
| | THRLCKHT_ASYNC | ACTIVE SEC A | 220M | | | 0.0 | 0.2 | 0.00 | 715K | 835 | | |
| | THRLCKIMS_ASYNC | ACTIVE PRIM A | 220M | | | 0.0 | 1.5 | 0.00 | 803K | 0 | | |

Secondary of asynchronously duplexed structure

Primary of asynchronously duplexed structure

New section with details for asynchronously duplexed secondary structures


ASYNCHRONOUS CF DUPLEXING SUMMARY

| TYPE | STRUCTURE NAME | TOTAL | ASYNC DUPLEX CF OPERATIONS | | | | TOTAL | #SUSPEND | ASYNC DUPLEX SYNC_UP REQUESTS | |
|------|-----------------|-------|----------------------------|-------|---------|------------------|-------|----------|-------------------------------|---------|
| | | | --TRANSMIT TIME-- | AVG | STD_DEV | --SERVICE TIME-- | | | AVG | STD_DEV |
| LOCK | THRLCKDB2_ASYNC | 1917K | 226.6 | 640.0 | 278.4 | 1154.7 | 0 | 0 | 0.0 | 0.0 |
| | THRLCKHT_ASYNC | 1132K | 210.4 | 539.8 | 217.9 | 612.5 | 0 | 0 | 0.0 | 0.0 |

- With APAR OA49148, RMF provides the support for asynchronous CF duplexing of lock structures
- The support is available for z/OS V2.2 and provides enhancements for following RMF reports:
 - Postprocessor CF Activity report
 - Monitor III CFACT report

Summary of RMF reporting enhancements for asynchronous CF duplexing of lock structures

- The status field of asynchronous duplexed structures in the Postprocessor CF Activity report and the Monitor III CFACT report indicates whether the structure is the asynchronous primary or asynchronous secondary instance of the structure during a rebuild.
- The CF Usage Summary section of RMF Postprocessor CF Activity report is enhanced with a new **Asynchronous CF Duplexing Summary** subsection. For asynchronous duplexed secondary structures it lists:
 - the number of asynchronous duplex CF operations and their consumed transmit and service times (average and standard deviation)
 - the number of issued and suspended sync-up requests together with suspend times (average and standard deviation)



RMF Postprocessor CF Activity report...

COUPLING FACILITY ACTIVITY

z/OS V2R2

SYSPLEX ENGTST2
RPT VERSION V2R2 RMF

DATE 11/22/2016
TIME 03.00.00

INTERVAL 030.00.000
CYCLE 01.000 SECONDS


COUPLING FACILITY NAME = CF03

CF TO CF ACTIVITY

| PEER CF | -RECEIVER- | | --SENDER-- | | # REQ | REQUESTS | | | # REQ | % OF REQ | DELAYED REQUESTS | | | |
|---------|------------|-----|------------|-----|-------|----------|------------------------|---------|-------|----------|------------------|---------------|---------|------|
| | TYPE | USE | TYPE | USE | | AVG/SEC | -SERVICE TIME(MIC)-AVG | STD_DEV | | | /DEL | AVG TIME(MIC) | STD_DEV | /ALL |
| CF01 | CIB | 2 | CIB | 2 | SYNC | 32997K | 18332 | 127.7 | 346.9 | 40K | 0.1 | 1657 | 2513 | 2.0 |
| | CS5 | 2 | CS5 | 2 | | | | | | | | | | |

CHANNEL PATH DETAILS

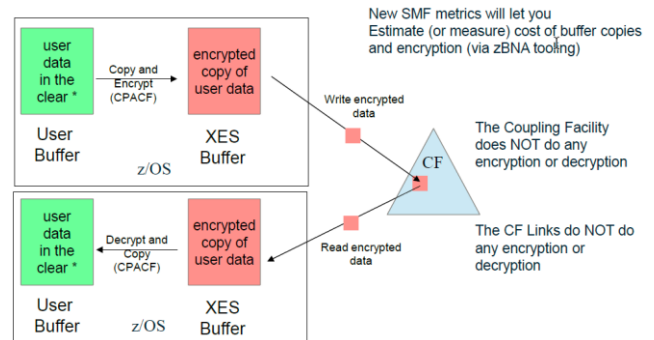
| PEER CF | ID | TYPE | R/S | OPERATION | MODE | DEGRADED | DISTANCE |
|---------|----|------|-----|-----------|----------------|----------|----------|
| CF01 | 2C | CS5 | R | 8X | GEN3 PCIE-O SR | N | <1 |
| | 2C | CS5 | S | 8X | GEN3 PCIE-O SR | N | <1 |
| | 2D | CS5 | R | 8X | GEN3 PCIE-O SR | N | <1 |
| | 2D | CS5 | S | 8X | GEN3 PCIE-O SR | N | <1 |
| | 2E | CIB | R | 12X | IFB3 HCA3-O | N | <1 |
| | 2E | CIB | S | 12X | IFB3 HCA3-O | N | <1 |
| | 2F | CIB | R | 12X | IFB3 HCA3-O | N | <1 |
| | 2F | CIB | S | 12X | IFB3 HCA3-O | N | <1 |



- The CF to CF report section of RMF Postprocessor CF Activity report is updated to show channel path detail data for a remote CF's sender channel paths.


Coupling Facility Encryption

- Today, customer data that flows through the Coupling Facility (CF) and the CF Link infrastructure is vulnerable to potential for exposure because the data is not encrypted
- With z/OS 2.3 for Coupling Facility Encryption:
 - Via CFRM policy you can direct that CF structures be transparently encrypted, with no middleware or application changes needed
 - For those structures, host -based CP Assist For Cryptographic Functions (CPACF) protected key AES 256-CBC encryption will be used to encrypt (decrypt) data and adjunct data that is written to (read from) the designated structure
- z14 and CF Encryption:
 - z14 CF is not required, but is recommended in order to simplify some sysplex recovery/reconciliation scenarios involving encrypted CF structures and "lost" CFRM datasets
 - z14 at the z/OS side is not required, but is recommended for the improved AES CBC encrypt/decrypt performance that z14 provides



- z/OS V2.3 provides support for end-to-end encryption for both CF data in flight and data at rest in CF structures (as part of a more pervasive encryption solution).
- Host-based CPACF encryption is used for high performance and low latency.
- Support provided can only be enabled when **all** systems are z/OS 2.3 or higher
 - Support for CF encryption metrics in z/OS and RMF is rolled back to 2.2
 - Toleration support provided on z/OS 2.2 gracefully prevents connections to encrypted CF structures, which cannot be accessed and used on a downlevel z/OS 2.2 or lower system
- IBM z14 CFLEVEL 22 CF images are not required, but are recommended in order to simplify some sysplex recovery and reconciliation scenarios involving encrypted CF structures. (Note that the CF image itself never decrypts, nor encrypts, any data).
- IBM z14 z/OS images are not required, but are recommended for the improved AES CBC encrypt/decrypt performance that z14 provides.

RMF Enhancements for CF Encryption



- RMF V2.2 and V2.3 stores new metrics on Coupling Facility read and write statistics for every CF structure in SMF type 74.4 record
- RMF V2.3 enhances Postprocessor CF Activity report and Monitor III CFACT report to display the encryption state of each CF structure.

| COUPLING FACILITY ACTIVITY | | | | | | | | | | | | | |
|---|----------------|----------------------|-----|-----------------|--------------|------------------------|--------------|--------------|--------------|-------------------------|-----------------------|----------------------|---------------|
| z/OS V2R3 | | SYSDPLEX SYSDDPLEX | | DATE 02/03/2017 | | INTERVAL 015.00.000 | | PAGE 1 | | | | | |
| | | RPT VERSION V2R3 RMF | | TIME 09.59.35 | | CYCLING 01.000 SECONDS | | | | | | | |
| ----- | | | | | | | | | | | | | |
| COUPLING FACILITY NAME - CP02 | | | | | | | | | | | | | |
| TOTAL SAMPLES (AVG) - 900 (MAX) - 900 (MIN) - 900 | | | | | | | | | | | | | |
| ----- | | | | | | | | | | | | | |
| COUPLING FACILITY USAGE SUMMARY | | | | | | | | | | | | | |
| ----- | | | | | | | | | | | | | |
| GENERAL STRUCTURE SUMMARY | | | | | | | | | | | | | |
| ----- | | | | | | | | | | | | | |
| TYPE | STRUCTURE NAME | STATUS CHG | ENC | ALLOC SIZE | % OF CF STOR | # REQ | % OF ALL REQ | % OF CF UTIL | AVG REQ/ SEC | LST/DIR ENTRIES TOT/CUR | DATA ELEMENTS TOT/CUR | LOCK ENTRIES TOT/CUR | DIR REC/ XI'S |
| LIST | IXCPLX_PATH3 | ACTIVE | YES | 12M | 0.8 | 2202 | 7.5 | 4.8 | 2.45 | 987 1 | 950 16 | N/A | N/A |
| | IXCPLX_PATH4 | ACTIVE | NO | 12M | 0.8 | 25352 | 86.1 | 17.2 | 28.17 | 987 1 | 950 21 | N/A | N/A |
| | IXCVLF | ACTIVE | YES | 12M | 0.8 | 1876 | 6.4 | 1.7 | 2.08 | 987 1 | 950 16 | N/A | N/A |
| LOCK | IGNLOCK00 | ACTIVE SEC | N/A | 79M | 5.1 | 0 | 0.0 | 66.8 | 0.00 | 184K 0 | 0 | 8389K 0 | N/A |
| STRUCTURE TOTALS | | | | 115M | 7.4 | 29430 | 100 | 90.5 | 32.70 | | | | |

- With z/OS V2.3 RMF and new function APAR OA52003, available for z/OS 2.2, RMF provides the following enhancements for CF encryption:
 - RMF is enhanced to gather new Coupling Facility structure related metrics and store the metrics in SMF74.4 records.
 - These new SMF 74.4 data can be used by zBNA tooling to predict the overhead of CF encryption and decryption, prior to actually turning on CF encryption.
 - In addition, z/OS V2.3 RMF is enhanced to report the encryption state per CF structure in the RMF Postprocessor CF Activity report and RMF Monitor III CFACT report.


RMF Enhancements at a Glance

- IBM z14 Support
 - Monitoring of Virtual Flash Memory
 - RMF support for for Crypto Express6S
 - IBM 10GbE RoCE Express2 support



- DS8880 Support
 - RMF Monitoring of SuperPAV capability
 - RMF Enhancements for zHyperWrite
 - Monitor Performance of PCIe Synchronous I/O links

- RMF Enhancements for Coupling Facility
 - Monitor Throughput and Delays of Asynchronous CF duplexing for lock structures
 - Predict overhead of CF encryption and report encryption state of CF structures



- RMF Enhancements for Spark
 - Report on new WLM functionality for Spark

- RMF WLM Reporting Enhancements
 - RMF support for Shorter WLM Response Time Goals



- Large Page Enhancements
 - Statistics for 2 GB Large Pages

- z/OS Platform for Apache Spark is the IBM solution to enable Apache Spark natively on z/OS. RMF reports on new WLM functionality that allows a better management of Spark on z/OS.

WLM & SRM Enhancements for Spark



- IBM is enabling the Apache Spark platform on z/OS. The z/OS Platform for Apache Spark offers an enterprise-grade solution to enable Apache Spark natively on z/OS.
- Spark is an in-memory compute engine for analytics, operating on data in memory.
- Problem:
 - Since Spark can be very processor and memory intensive customers need a means to limit Spark's consumption of both resources.
 - Spark is primarily running on specialty engines but WLM's resource groups cannot be used to limit processor access to zAAPs and zIIPs.

WLM APAR OA50845
RSM APAR OA51171


- Solution:
 - WLM provides new service class attribute that allows to prohibit specialty-engines eligible work to overflow to regular CPs for needs help processing.
 - Limit Spark's processor consumption to specialty-engines
 - WLM provides new data field for resource groups that allows to specify an upper limit for real memory consumption.
 - Limit Spark's memory consumption

Honor Priority = DEFAULT|NO

Memory Limit = xxx GB

- As part of the overall z/OS Spark and Cloud Solution, the z/OS Platform for Apache Spark is the IBM solution to enable Apache Spark natively on z/OS.
- Spark is an in-memory compute engine for analytics that operates on data in memory.
- Problems that need to be addressed:
 - Since Spark can be very processor and memory intensive customers need a means to limit Spark's consumption of both resources.
 - Spark is primarily running on specialty engines but WLM's resource groups cannot be used to limit processor access to zAAPs and zIIPs.
- Solution:
 - As a first step to limit Spark's processor consumption, WLM implements a new attribute on service class level that allows to prohibit specialty-engines eligible work to overflow to regular CPs for needs help processing. Customers need to classify Spark workloads into such service classes to exploit this function.
 - New attribute for service classes: Honor Priority = DEFAULT|NO
 - DEFAULT keeps OPT=IFA/SUP HonorPriority in effect.
 - NO ignores OPT=IFA/SUP HonorPriority
 - To allow limitation of Spark's memory consumption, WLM implements a new data field for resource groups that allows to specify an upper limit for real memory consumption. Customers need to classify Spark workloads into service classes that are connected to such resource groups to exploit this function.
 - New field for resource groups: Memory Limit
 - To be specified in multiples of Giga Bytes
 - The support is provided by WLM APAR OA50845 and RSM APAR OA51171

RMF Enhancements for Spark



- Service Class (Period) header section of RMF **Postprocessor Workload Activity** report enhanced to provide 'HONOR PRIORITY' setting as defined in WLM service class definition

Prohibit offloading to regular CPs

| | | | |
|----------------------------|--------------|----------------------------|-----------------------|
| REPORT BY: POLICY=STANDARD | WORKLOAD=STC | SERVICE CLASS=DB2P | RESOURCE_GROUP=RGPIX1 |
| | | CRITICAL =NONE | HONOR_PRIORITY=NO |
| | | DESCRIPTION =DB2 subsystem | |

| -TRANSACTIONS- | TRANS-TIME | HHH.MM.SS.TTT | --DASD | I/O-- | ---SERVICE--- | SERVICE TIME | ---APPL %--- | --PROMOTED-- | ---STORAGE--- | |
|----------------|------------|---------------|--------|-------|---------------|--------------|--------------|--------------|---------------|-----------------|
| AVG 1.00 | ACTUAL | 0 | SSCHRT | 0.0 | IOC | 0 | CPU 42.328 | CP 4.70 | BLK 0.000 | AVG 1498.00 |
| MPL 1.00 | EXECUTION | 0 | RESP | 0.0 | CPU | 33694K | SRB 0.005 | AAPCP 0.00 | ENQ 0.000 | TOTAL 1492.94 |
| ENDED 0 | QUEUED | 0 | CONN | 0.0 | MSO | 1233K | RCT 0.000 | IIPCP 0.00 | CRM 0.000 | SHARED 1.00 |
| END/S 0.00 | R/S AFFIN | 0 | DISC | 0.0 | SRB | 4204 | IIT 0.000 | | LCK 0.002 | |
| #SWAPS 3 | INELIGIBLE | 0 | Q+PEND | 0.0 | TOT | 34931K | HST 0.000 | AAP N/A | SUP 0.000 | -PAGE-IN RATES- |
| EXCTD 0 | CONVERSION | 0 | IOSQ | 0.0 | /SEC | 38812 | AAP N/A | IIP 0.00 | | SINGLE 0.0 |
| AVG ENC 0.00 | STD DEV | 0 | | | ABSRPTN | 39K | IIP 0.000 | | | BLOCK 0.0 |
| REM ENC 0.00 | | | | | TRX SERV | 39K | | | | SHARED 0.0 |
| MS ENC 0.00 | | | | | | | | | | HSP 0.0 |

| | | | | |
|---------------------|---------|--------------------|--------------------|--------------|
| TRANSACTION APPL% : | TOTAL : | CP 4.70 | AAP/IIP ON CP 0.00 | AAP/IIP 0.00 |
| MOBILE : | CP 0.00 | AAP/IIP ON CP 0.00 | AAP/IIP 0.00 | |

- With RMF for z/OS 2.3 and new function APAR OA50760, RMF is enhanced to report the value of the new "Honor Priority" attribute at service class, address space and enclave level, and the value of the new memory limit at resource group level. PTFs are available for z/OS 2.1 and z/OS 2.2.
- The Service Class (Period) header section in the RMF Postprocessor Workload Activity report is enhanced to provide the "HONOR PRIORITY" setting as defined in the WLM service class definition.
 - The only possible value is "NO". It means that CP help processing (honor priority processing) is suppressed by WLM
 - Otherwise, the "HONOR PRIORITY" field does not appear, i.e. OPT settings IFAHONORPRIORITY and IIPHONORPRIORITY are in effect


RMF Enhancements at a Glance

- IBM z14 Support
 - Monitoring of Virtual Flash Memory
 - RMF support for for Crypto Express6S
 - IBM 10GbE RoCE Express2 support



- DS8880 Support
 - RMF Monitoring of SuperPAV capability
 - RMF Enhancements for zHyperWrite
 - Monitor Performance of PCIe Synchronous I/O links

- RMF Enhancements for Coupling Facility
 - Monitor Throughput and Delays of Asynchronous CF duplexing for lock structures
 - Predict overhead of CF encryption and report encryption state of CF structures



- RMF Enhancements for Spark
 - Report on new WLM functionality for Spark

- RMF WLM Reporting Enhancements
 - RMF support for Shorter WLM Response Time Goals

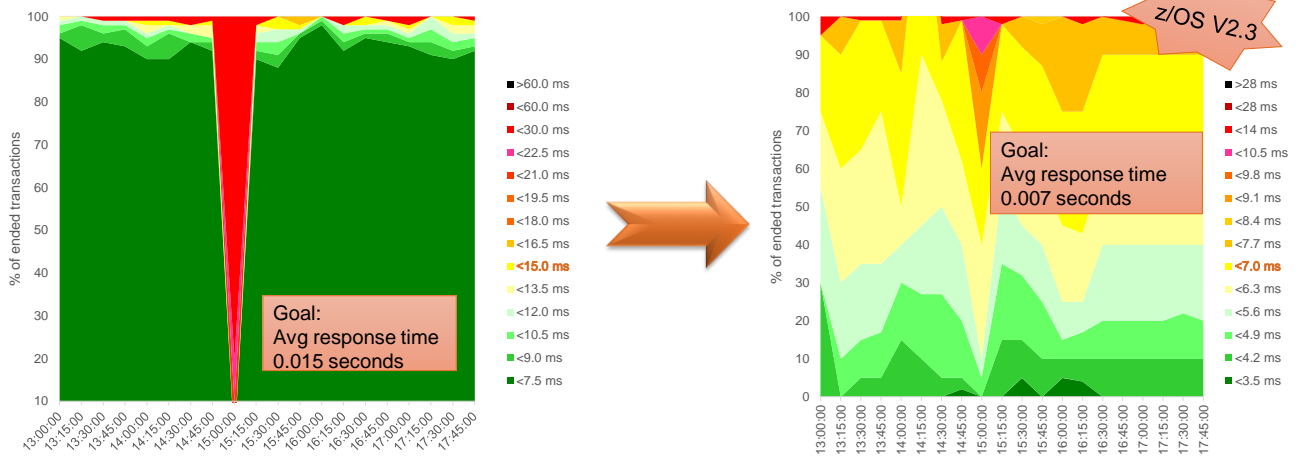


- Large Page Enhancements
 - Statistics for 2 GB Large Pages

- With z/OS V2R3, WLM enhances the definition of response time goals for service classes. The current lower bound of 15 milliseconds for a response time goal is replaced by one millisecond. RMF is enhanced to report on shorter response times in various RMF reports.

RMF Support for WLM Shorter Response Time Goals

- WLM enhances the definition of response time goals for service classes
- For service class periods with an *average response time goal* or a *response time goal with percentile*, the lowest goal which can be specified changes from 15 milliseconds to one millisecond



44 | The Latest and Greatest

12/21/2017

© 2017 IBM Corporation

- Currently, the shortest average or percentile response time that can be specified in the WLM service definition is 15 milliseconds. This makes it difficult to manage such transactional work effectively where transactions complete well below 15 milliseconds (e.g. in DB2 DDF and CICS environments).
- The figure on the left illustrates an exemplary response time distribution for a service class with a 15 millisecond response time goal and very fast running transactions.
- During the presented time period, approximately 95% of the transactions complete with an average response time below 7.5 milliseconds (dark green area).
- Thus, the service class considerably overachieves its goal allowing it to run at a low dispatching priority.
- This makes the service class susceptible to workload spikes where the average response time of all transactions rises far above the goal (red area with the outlier around 15:00).
- With z/OS V2R3, WLM replaces the current lower bound of 15 milliseconds for a response time goal by one millisecond. This allows to specify meaningful goal values for very fast running transactions.
- The figure on the right illustrates a response time distribution if a more meaningful goal of 7 milliseconds will be used.
- The response time distribution is better balanced. The workload always runs at reasonable dispatch priority preventing bad response times even with workload spikes.
- With z/OS V2R3, RMF gathers these more granular response time and distribution information from WLM by calling the IWMRCOLL service.
- RMF adapts the workload reporting in various Monitor III reports, Postprocessor Workload Activity report (WLMGL) and SMF type 72-3 records to properly display response times of transactions which run faster than one millisecond on the average. The Distributed Data Server is extended to deliver new performance metrics for the shorter response time values.

Shorter RT Goals – RMF Workload Activity Report

New formatting of response times

| WORKLOAD ACTIVITY | | | | | | | | | | | | | |
|---|--------|----------------------|----------|----------------------|------|------------------------|----------|----------------|------|--------------|----------|--------|-----|
| z/OS V | | SYSPLX UTCLXCB | | DATE 12/19/2016 | | INTERVAL 15.00.003 | | MODE = GOAL | | PAGE 1 | | | |
| | | RPT VERSION V2R3 RMF | | TIME 02.45.00 | | | | | | | | | |
| POLICY ACTIVATION DATE/TIME 12/18/2016 09.00.11 | | | | | | | | | | | | | |
| ----- SERVICE CLASS PERIODS | | | | | | | | | | | | | |
| POLICY=BASEPOL | | WORKLOAD=STC_WLD | | SERVICE CLASS=STCLOW | | RESOURCE GROUP=-NONE | | PERIOD=1 | | IMPORTANCE=3 | | | |
| | | CRITICAL | | -NONE | | | | | | | | | |
| ----- TRANSACTIONS----- | | | | | | | | | | | | | |
| TRANS-TIME | | HHH.MM.SS.FFFFFFFF | | TRANS-APPL% | | CP-IIPCP/AAPCP-IIP/AAP | | ---ENCLAVES--- | | | | | |
| AVG | 46.51 | ACTUAL | 5.637787 | TOTAL | 1.04 | 0.00 | 0.00 | AVG ENC | 0.00 | | | | |
| MPL | 46.49 | EXECUTION | 4.957964 | MOBILE | 0.00 | 0.00 | 0.00 | REM ENC | 0.00 | | | | |
| ENDED | 11 | QUEUED | 679822 | CATEGORYA | 0.00 | 0.00 | 0.00 | MS ENC | 0.00 | | | | |
| END/S | 0.01 | R/S AFFIN | 0 | CATEGORYB | 0.00 | 0.00 | 0.00 | | | | | | |
| #SWAPS | 971 | INELIGIBLE | 0 | | | | | | | | | | |
| EXCTD | 0 | CONVERSION | 190650 | | | | | | | | | | |
| | | STD DEV | 2.506325 | | | | | | | | | | |
| -----SERVICE----- | | | | | | | | | | | | | |
| IOC | 23349 | CPU | 9.021 | CP | 1.06 | BLK | 0.000 | SSCHRT | 2.8 | AVG | 4118.93 | SINGLE | 0.0 |
| CPU | 566054 | SRB | 0.383 | IIPCP | 0.00 | ENQ | 0.000 | RESP | 0.1 | TOTAL | 191482.5 | BLOCK | 0.3 |
| MSO | 0 | RCT | 0.107 | IIP | 0.00 | CRM | 0.000 | CONN | 0.0 | SHARED | 0.00 | SHARED | 0.0 |
| SRB | 24025 | IIT | 0.012 | AAPCP | 0.00 | LCK | 708670.7 | DISC | 0.0 | | | HSP | 0.0 |
| TOT | 613428 | HST | 0.000 | AAP | N/A | SUP | 0.000 | Q+PEND | 0.0 | | | | |
| /SEC | 682 | IIP | 0.000 | | | | | IOSQ | 0.0 | | | | |
| ABSRPTN | 15 | AAP | N/A | | | | | | | | | | |
| TRX SERV | 15 | | | | | | | | | | | | |

Added APPL% values for transactions with reporting attributes CategoryA/B

New layout: Re-arranged resource consumption section for service and report classes and workloads

Summary of changes:

- The Postprocessor Workload Activity (WLMGL) report is adapted to present the response time goal, transaction times, and response time distributions with higher precision to reflect actual response times lower than one millisecond.
- The representation of response times in the format HH.MM.SS.TTT is changed to HHH.MM.SS.FFFFFFFF in every place where an actual or average response time has to be reported.
- TTT stands for thousandths of seconds, FFFFFFFF stands for fractional representation of seconds, thus microseconds can be displayed.
- The format for response time goals itself will remain unchanged because the smallest possible response time goal of one millisecond can be represented as 000.00.00.001 (HHH.MM.SS.TTT).
- The data in the resource consumption section is re-arranged to allow for a better usage of the available space

Shorter RT Goals – RMF Mon III SYSSUM

```

RMF V2R3 Sysplex Summary - UTCPLXCB Line 19 of 34
WLM Samples: 479 Systems: 8 Date: 02/09/17 Time: 12.00.00 Range: 120 Sec
>>>>>>XXXXXXXXXXXXXXXXXXXX<<<<<<<<
Service Definition: SYSTEST Installed at: 11/15/16, 08.18.14
Active Policy: BASEPOL Activated at: 02/09/17, 09.00.02
----- Goals versus Actuals -----
Name T I Exec Vel --- Response Time --- Perf Trans
Goal Act ---Goal--- --Actual-- Indx Ended
Rate
----- Avg. Resp. Time-----
WAIT EXECUT ACTUAL
Time Time Time
STCHI S 2 35 69 0.50 0.000
STCLOW S 3 25 45 0.56 0.192
SYSTEM W 81 0.192 812.3 489.5 1165
SYSSTC S N/A 64 N/A
SYSTEM S N/A 84 0.000 0.000 0.000
TSO_WLD W 22 5.175 3.502 10395 10399
TSOHIGH S 67 0.092 197.7 324.8 522.5
1 1 60 400 90% 67% 4.00 0.075 0.000 284.5 284.5
2 2 100 3000 AVG 1593 AVG 0.53 0.017 1087 506.1 1593
TSOLOW S 19 5.083 0.000 10577 10577
1 3 27 500 85% 21% **** 3.583 0.000 8396 8396
2 3 7.1 5000 AVG 15786 AVG 3.16 1.500 0.000 15786 15786
CONSOLES R 40 N/A 0.000 0.000 0.000
IXGLOGR R 82 N/A 0.000 0.000 0.000
TCPIP R 20 N/A 0.000 0.000 0.000
XCFAS R 32 N/A 0.000 0.000 0.000
    
```

Response time goal, actuals and average response times are displayed in milliseconds

- Monitor III SYSINFO, SYSSUM, SYSWKM, SYSRTD, and GROUP reports are adapted to display response time goals and actual response times with higher precision.
- The columns reporting on goal and average response times now represent milliseconds, the values are scaled with factors K (10³), M (10⁶) or G (10⁹) if they do not fit into the defined column width.
- Examples:
 - A response time goal of five milliseconds (defined as 00.005 seconds in the WLM policy definition) is reported as 5 (milliseconds).
 - A goal of one second is reported as 1000 (milliseconds).
 - A goal of one minute is reported as 60000 (milliseconds).
 - A goal of one hour is reported as 3600K (milliseconds) which is 3600 * 10³.

Shorter RT Goals – RMF Mon III SYSRTD

Upper part of the report has been redesigned to contain a table with the complete information of all response time distribution buckets for the selected service class period

Average response times in the scrollable lower part of the report are displayed in milliseconds

```

RMF V2R3 Response Time - UTCPLXCB Line 1 of 9
Command ==> _ Scroll ==> CSR
WLM Samples: 239 Systems: 8 Date: 01/27/17 Time: 02.08.00 Range: 60 Sec

Class: CICSLOW Period: 1 Goal: 500.0 ms for 80%

Trx # 50% 60% 70% 80% 90% 100% 110% 120% 130% 140% 150% 200% 400% >400%
Trx % 55K 24 14 25 13 12 12 14 12 13 8 48 121 11
Trx % 99.4 0.04 0.03 0.05 0.02 0.02 0.02 0.03 0.02 0.02 0.01 0.09 0.22 0.02

--Avg. Resp. Time-- Trx --Subsystem Data-- --Exec Data--
WAIT EXECUT ACTUAL Rate Actv Ready Delay Ex Vel Delay
*ALL 0.000 11.01 10.70 926.0 17 1 62 N/A 7
CB8A all 0.000 11.74 0.491 0.100 0 0 0 N/A 0
CB8B all 0.000 13.52 0.930 0.483 10K 0 95K N/A 1
CB8C all 0.000 18.80 0.894 0.117 48K 0 272K N/A 4
CB8D all 0.000 15.41 0.402 0.150 6912 0 366K N/A 0
CB8E all 0.000 0.796 0.796 0.117 0 0 0 N/A 0
CB86 all 0.000 33.87 38.90 59.33 59 0 53 N/A 10
CB88 all 0.000 7.447 8.775 865.5 3 2 48 N/A 38
CB89 all 0.000 16.49 0.490 0.150 0 0 414K N/A 0
    
```

- The SYSRTD report shows how the response time for a specific service or report class is distributed. Two levels of detail are shown:
 - A table shows the distribution of response time for all systems in a sysplex which have data available in the selected period.
 - A table shows how each system contributed to the overall response time.

How to read the response time distribution table

- To provide a picture of how a performance group was performing, response time distributions are provided for both service classes and report classes. These distributions consist of 14 buckets of information.
- The header contains the value of the particular bucket, which is a percentage of the specified goal. One bucket always maps exactly to the specified goal, with a value of 100%.
- In each bucket is the number of transactions that completed in the amount of time that is represented by that bucket. This number is reported in the first row of the response time distribution table, the second row of the table shows the percentage of transactions that are contained in each bucket.
- In the example, each of the 14 buckets represents a percentage of the specified 0.5-second goal. For instance, bucket 2 represents all transactions that completed in 50% to 60% of the goal, or 250-300 milliseconds, while bucket 10 contains the number of transactions that completed in 130% to 140% of the goal, or 650-700 milliseconds. Notice that bucket 6 falls exactly on the goal (100% of goal, or 0.5 seconds). This bucket captures all those transactions that complete in 400-500 milliseconds.
- The two end buckets (buckets 1 and 14) have special meaning. Bucket 1 (labelled 50%) contains the total number of transactions that completed in up to 50% of the goal. Bucket 14 (labelled >400%) contains the number of transactions that completed in greater than 4 times the goal.

Shorter RT Goals – RMF DDS Metrics

New metrics are available for the SYSPLEX resource:

| | | | |
|---|-------------|--------|---|
| active time (ms) by WLM service class | Explanation | 8D5E20 | } - by WLM service class - by WLM service class period - by WLM report class - by WLM report class period - by WLM workload |
| active time by WLM service class | Explanation | 8D0E60 | |
| execution velocity by WLM service class | Explanation | 8D0F50 | |
| queue time (ms) by WLM service class | Explanation | 8D5E80 | |
| queue time by WLM service class | Explanation | 8D10C0 | |
| response time (ms) by WLM service class | Explanation | 8D5F20 | |
| response time by WLM service class | Explanation | 8D1160 | |
| response time goal (ms) by WLM service class period | Explanation | 8D5F80 | } - by WLM service class period |
| response time goal by WLM service class period | Explanation | 8D11C0 | |
| response time (ms) by WLM service class | Explanation | 8D5F10 | } - by WLM service class - by WLM service class period - by WLM report class - by WLM report class period - by WLM workload |
| response time by WLM service class | Explanation | 8D1150 | |

- All new fields from SYSSUM and SYSINFO reports are available in the RMF Distributed Data Server when requesting the full Monitor III reports with HTTP request:

<http://<host>:8803/gpm/rmf3.xml?report=SYSSUM&resource=„,SYSPLEX”>

http://<host>:8803/gpm/rmf3.xml?report=SYSINFO&resource=„,SYSN,MVS_IMAGE”

- New metrics are available for the SYSPLEX resource:
 - The following 3 metrics are available by WLM service class, by WLM service class period, by WLM report class, by WLM report class period, by WLM workload:
 - active time (ms)
 - queue time (ms)
 - execution time (ms)
 - The following metric is available by WLM service class period:
 - response time goal (ms)
- New metrics are available for the MVS_IMAGE resource:
 - The following metric are available by WLM service class, by WLM service class period, by WLM report class, by WLM report class period, by WLM workload:
 - execution time (ms)


RMF Enhancements at a Glance

- IBM z14 Support
 - Monitoring of Virtual Flash Memory
 - RMF support for for Crypto Express6S
 - IBM 10GbE RoCE Express2 support



- DS8880 Support
 - RMF Monitoring of SuperPAV capability
 - RMF Enhancements for zHyperWrite
 - Monitor Performance of PCIe Synchronous I/O links

- RMF Enhancements for Coupling Facility
 - Monitor Throughput and Delays of Asynchronous CF duplexing for lock structures
 - Predict overhead of CF encryption and report encryption state of CF structures



- RMF Enhancements for Spark
 - Report on new WLM functionality for Spark

- RMF WLM Reporting Enhancements
 - RMF support for Shorter WLM Response Time Goals



- Large Page Enhancements
 - Statistics for 2 GB Large Pages

- RMF provides new support for monitoring and reporting of 2 GB memory frames.



RMF support for 2GB Large Pages

Large Frame Areas

• 1 MB LFArea Large Frame Area

- Fixed storage each frame is 1MB
- Defined by LFAREA (1M =) (IEASYSxx)
- Included in Available Frame Count when INCLUDE1MAFC=YES (Default with z/OS 2.2)

• PLArea –Pageable Large Area

- Pageable storage each frame is 1MB
- System defined size
- Allocated on SCM capable machines
- Can overflow in the LFAREA

• 2GB LFArea Large Frame Area

- Fixed storage each frame is 2GB
- Defined by LFAREA (2G =) (IEASYSxx)
- Not included in Available Frame Count – used only for 2G requests
- Reserved for specific 2GB memory objects

Increased efficiency for DB2 buffer pools and Java heap

2GB Large Frame Area:

- Size defined by LFAREA parameter in IEASYSxx PARMLIB member.
- Amount of 2GB pages can be specified as target and minimum values or percentages of online real storage

2GB Memory Object:

- Obtained by use of IARV64 macro
- ```
IARV64 REQUEST=GETSTOR
 ,UNITS=1,
 ,UNITSIZE=2G
 ,PAGEFRAMESIZE=2G
 ,TYPE=FIXED
```

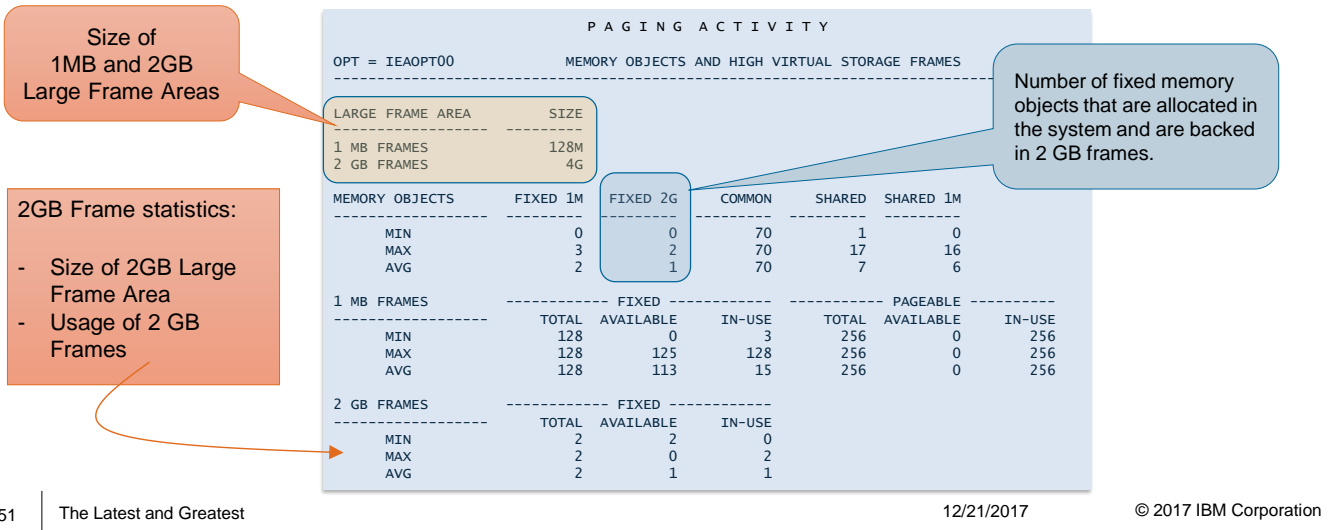
- With the introduction of the enhanced-dynamic address translation facility (EDAT) on z/architecture HW, IBM introduced support for 1 MB pages. Dynamic address translation for a 1 MB page reference does not require a page table. The segment table entry that would normally point to a page table instead points to the 1 MB page directly. The translation lookaside buffer (TLB) coverage is also effectively increased, because one entry for a fixed 1 MB page provides a greater coverage area than one entry for a 4 KB page.
- z/OS supports fixed 1 MB pages in the large frame area (LFAREA)
- FLASH storage capability and the enhanced-DAT facility 2 (EDAT-2) for a 2 GB page size was introduced with IBM zEnterprise EC12 . Beginning with z/OS 2.1 these features are exploited with two new page types:
  - Pageable 1M pages and
  - Fixed 2 GB pages.
- z/OS uses FLASH storage as Storage Class Memory (SMC) for paging the pageable 1M pages.. The fixed 2 GB pages will reside in the LFAREA with the fixed 1 MB pages.
- Fixed 2 GB pages provide an even greater DAT performance benefit because the translation does not require a segment table or its page tables. Instead, the region third table entry points directly to the fixed 2 GB page. Again, TLB coverage is increased, because one entry for a fixed 2 GB page provides an even greater coverage area than one entry for a 4 KB page.
- In order to have 2GB large-page frames available in z/OS, a 2GB large frame area must be configured after the initial program load (IPL). This is achieved by sub-parameter keywords of the LFAREA keyword in the active IEASYSxx member of SYS1.PARMLIB. The system programmer can specify the amount of real storage to be used for 1MB and 2GB large pages using the LFAREA keyword with the following syntax:

```
LFAREA=([1M=(target[%] [,minimum[%]])] [,2G=(target[%] [,minimum[%]])] [,prompt | noprompt])
```

where target is the number of pages of online real storage to be used, or target% is the percentage of online storage at IPL to be used for 1MB or 2GB pages. The default for the 1MB and 2GB large frame area is 0. When the 1M= and 2G= keywords are used, the maximum that can be specified is 80% \* (IPL-time online storage – 4G).

## 2GB Pages - RMF Postprocessor Paging Activity report

- New system wide statistics about 2 GB Memory Objects and Frames in MEMORY OBJECTS AND HIGH VIRTUAL STORAGE FRAMES section



- With APAR OA48913, RMF provides the support for monitoring and reporting of 2 GB memory frames.
- The support is available for z/OS V2.2 and provides following enhancements:
  - System-wide usage data of 2GB Large Pages:
    - SMF 71 records
    - Postprocessor Paging Activity Report (and Overview Conditions)
    - RMF Monitor III STORM Report, System Summary section
  - Address space related usage data of 2 GB Large Pages:
    - SMF 78-2 records
    - RMF Postprocessor Virtual Storage Activity Report (VSTOR)

### RMF Postprocessor Paging Activity report (PAGING):

- The MEMORY OBJECTS AND HIGH VIRTUAL STORAGE FRAMES section displays the size of the 2GB Large Frame Area as well as new statistics about 2 GB memory objects and frames.:
  - LARGE FRAME AREA
    - 1 MB FRAMES      Size of the 1 MB portion of the Large Frame Area  
(This field replaces LFAREA SIZE in the section header.)
    - 2 GB FRAMES      Size of the 2 GB portion of the Large Frame Area.
  - MEMORY OBJECTS
    - FIXED 2G      Number of fixed memory objects that are allocated in the system and are backed in 2 GB frames.
  - 2 GB FRAMES
    - FIXED - TOTAL      Total number of 2 GB frames that can be used by fixed 2 GB memory objects. This value is equal to the size of the 2 GB portion of the Large Frame Area in units of 2 GB.
    - FIXED – AVAILABLE      Number of 2 GB frames in the Large Frame Area that are not in-use.
    - FIXED - IN-USE      Number of 2 GB frames in the Large Frame Area that are in-use.

# 2GB Pages - RMF Monitor III STORM report

**System Summary:**

Average number of fixed memory objects that can be backed in 1 MB frames / 2 GB frames

**At a Glance:**  
Size and Usage of Large Frame Areas

| RMF V2R2 Storage Memory Objects |            |                    |       |                |               |                |        |               |                 | Line 1        |       |                |  |
|---------------------------------|------------|--------------------|-------|----------------|---------------|----------------|--------|---------------|-----------------|---------------|-------|----------------|--|
| Samples: 60                     |            | System: SYSE       |       | Date: 02/10/17 |               | Time: 12.00.00 |        | Range: 60     |                 | Sec           |       |                |  |
| ----MemObj----                  |            |                    |       | ---Frames---   |               |                |        | --1MB Fixed-- |                 | --2GB Fixed-- |       | -1MB Pageable- |  |
| Fixed 1M                        | 3          | Shared             | 1938  | Total          | 512           | Total          | 6      | Initial       | 3221            |               |       |                |  |
| Fixed 2G                        | 1          | Common             | 45977 | Common         | 61            | %Used          | 66.7   | Dynamic       | 0               |               |       |                |  |
| Shared                          | 12         | %Used              | 10.5  | %Used          | 11.9          | %Used          | 76.0   |               |                 |               |       |                |  |
| Common                          | 108        |                    |       |                |               |                |        |               |                 |               |       |                |  |
| -----                           |            |                    |       |                |               |                |        |               |                 |               |       |                |  |
| Service                         |            | --Memory Objects-- |       |                | --1M Frames-- |                |        | 2G-Fr         | -----Bytes----- |               |       |                |  |
| Jobname                         | C Class    | ASID               | Total | Comm           | Shr           | Fixed          | Pgable | Fixed         | Total           | Comm          | Shr   |                |  |
| PGBLE1MC                        | B BATCHLOW | 0057               | 11    | 3              | 0             | 0              | 0      | 0             | 20.0M           | 7168K         | 0     |                |  |
| STOR642G                        | B BATCHLOW | 0056               | 10    | 1              | 0             | 0              | 0      | 4             | 8206M           | 1024K         | 0     |                |  |
| PGBLE1MP                        | B BATCHLOW | 0058               | 10    | 1              | 0             | 0              | 5      | 0             | 19.0M           | 1024K         | 0     |                |  |
| COMP641M                        | B BATCHLOW | 0061               | 10    | 2              | 0             | 30             | 0      | 0             | 44.0M           | 31.0M         | 0     |                |  |
| PGBLE1MS                        | B BATCHLOW | 0059               | 9     | 0              | 1             | 0              | 0      | 0             | 2061M           | 0             | 2048M |                |  |
| COMS641M                        | B BATCHLOW | 0062               | 8     | 0              | 0             | 0              | 0      | 0             | 13.0M           | 0             | 0     |                |  |

**Address space detail:**  
Average number of 2 GB frames in the Large Frame Area

**New and changed System Summary fields in Monitor III Storage Memory Objects (STORM) report:**

|           |          |                                                                                                           |
|-----------|----------|-----------------------------------------------------------------------------------------------------------|
| MemObj    | Fixed 1M | Average number of fixed memory objects that are allocated in the system and can be backed in 1 MB frames. |
|           | Fixed 2G | Average number of fixed memory objects that are allocated in the system and are backed in 2 GB frames.    |
| 2GB Fixed | Total    | Total number of 2 GB frames that can be used by fixed memory objects.                                     |
|           | %Used    | Percentage of 2 GB frames that are used by fixed memory objects in the Large Frame Area.                  |

**New and changed Address Space related fields:**

|       |       |                                                                                    |
|-------|-------|------------------------------------------------------------------------------------|
| 2G-Fr | Fixed | Average number of 2 GB frames in the Large Frame Area owned by this address space. |
|-------|-------|------------------------------------------------------------------------------------|

## Function Reference

| Function                                                  | Availability                           |
|-----------------------------------------------------------|----------------------------------------|
| RMF support for Virtual Flash Memory                      | APAR OA50761 (z/OS 2.2 only)           |
| RMF support for 10GbE RoCE Express2 card                  | APAR OA50762                           |
| RMF support for Crypto Express6S (CEX6) card              | APAR OA50693                           |
| RMF support for zHyperLink                                | APAR OA50755                           |
| RMF z14 Toleration support                                | APAR OA51913                           |
| RMF support for SuperPAV                                  | APAR OA49415                           |
| RMF support for IBM zHyperWrite                           | z/OS 2.3 (Coexistence support OA48870) |
| RMF Asynchronous CF duplexing for lock structures support | APAR OA49148                           |
| RMF Enhancements for CF Encryption                        | APAR OA52003                           |
| RMF Enhancements for Spark                                | APAR OA50760                           |
| RMF Support for WLM Shorter Response Time Goals           | z/OS 2.3                               |
| RMF support for 2GB Large Pages                           | APAR OA48913                           |



RMF z14  
Support



## 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/servers/resourcelink/svc00100.nsf/pages/zOSV2R3RmfPublications?OpenDocument](http://www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zOSV2R3RmfPublications?OpenDocument)

