



## Where is my RMF?

z/VM Performance and Capacity Management from a z/OS Perspective

2003 zSeries EXPO – Session V94

November 2003

William L. Sheldon, Jr., Ph.D.

Bill.Shelden@PerfMan.com

ISM, Inc.

## Trademarks and Registered Trademarks



- PerfMan®, the PerfMan logo and all other ISM Software product are registered trademarks or trademarks of The Information Systems Manager, Inc.
- OS/390®, Parallel Sysplex®, RMF, z/OS®, z/Series®, and z/VM® are trademarks or registered trademark of International Business Machines, Inc.
- All other registered trademarks or trademarks belong to their respective companies.

## Where is my RMF?



### Abstract

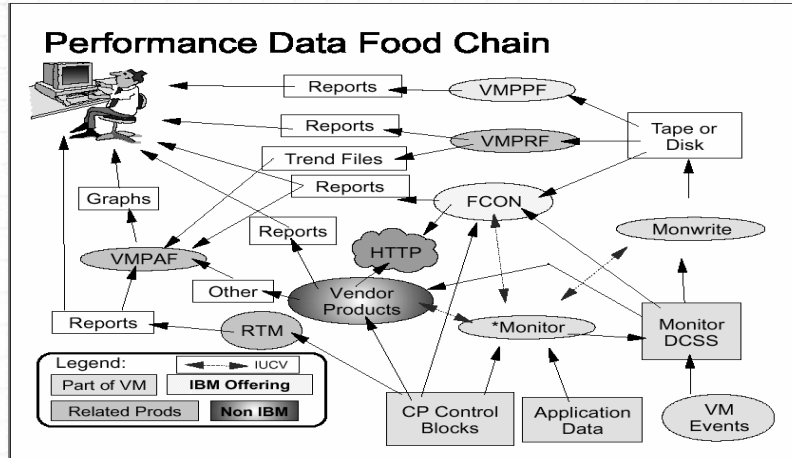
*z/VM, an operating system many z/OS system programmers believed would be replaced by PR/SM is now making a comeback based on its almost infinite ability to host Linux guests. With IBM's new zSeries 990, data centers can now run 1000's of Linux images within a single z990. Well-established z/OS strongholds will likely need to lower their drawbridges and welcome z/VM as a partner in reducing the vast complexity caused by 1000's of independent rack mounted Intel based Linux systems. But what should I expect? How do I measure this? How do I deal with virtual devices? Where is my RMF & SMF data?*

## Agenda



- Introduction
  - z/VM Performance Data from \*MONITOR
  - RMF and \*MONITOR Data
  - Post Processing Data
  - \*MONITOR Data Considerations
  - \*MONITOR Data Domains
- RMF and \*MONITOR Data Presentation
  - CPU Activity
  - DASD I/O
  - Workloads

## VM Performance Data ...from Bill Bitner Presentation



## RMF and \*MONITOR Data



- Resource Measurement Facility (RMF)
  - Write SMF 70-79's to SMF datasets
  - Variable length spanned records
    - Many subtypes
    - Many dependent segments
- \*MONITOR
  - Writes 4K blocks to VM shared segment
  - MONWRITE is a sample program in z/VM that writes \*MONITOR data to disk or tape
  - 4K blocks are control and data blocks
  - Reference:
    - z/VM Performance 4.3.0 (SC24-5999)
    - Appendices B and C

## Post Processing RMF and \*MONITOR Data



- RMF
  - RMF's Post Processor Reports
  - Vendor Products
- \*MONITOR
  - IBM's Performance Reporting Facility (PRF)
    - Functionally replaced by Performance Toolkit for VM in z/VM 4.4.0
  - Vendor Products

## \*MONITOR Data Details



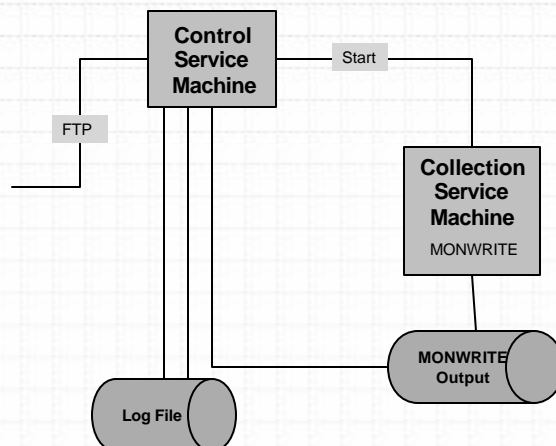
- Data written to shared segment in 4K blocks
- Organized by record type within domain
- Most (not all) metrics are cumulative
- Cumulative metrics will wrap (some quickly)
  - Makes it hard to compute totals
- Time stamps are in TOD format
- Can be retrieved and written to disk by sample program MONWRITE
- Interval and Event data
- References:
  - z/VM Monitor Records 4.1.0 (SC24-6012)
  - z/VM website

## \*MONITOR data Domains



- 0 = System
- 1 = Monitor
- 2 = Scheduler
- 3 = Storage
- 4 = User
- 5 = Processor
- 6 = I/O
- 7 = Seek
- 10 = Application

## Sample Environment in z/VM to Capture \*MONITOR Data




## Data Presentation CPU Activity




- Measured Dispatch Time
- % Utilization, MIPS, MSUs etc.
- CPU Queuing in the Primary LPAR
- LPAR Configuration Information
  - LPAR names and numbers
  - Weights, Capping, Wait Completion, number of LP's, etc.





## z/OS Dispatch Time RMF Type 70



**z/VM Dispatch Time  
\*MONITOR (D0,R16)**



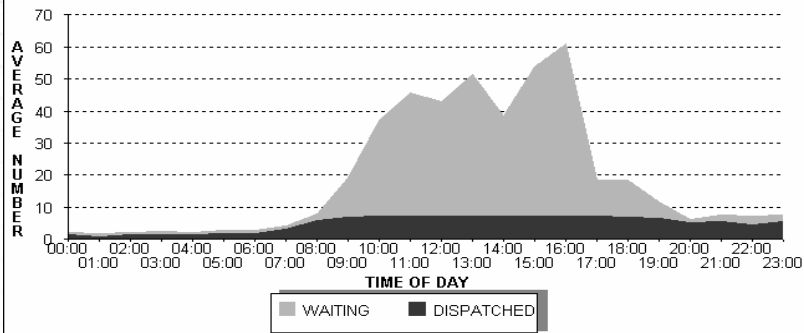
**z/VM (or z/OS) Dispatch Time  
Expressed in Other Terms**



**z/OS CPU Queuing  
RMF Type 70**



***In and Ready Address Spaces  
SYSA***

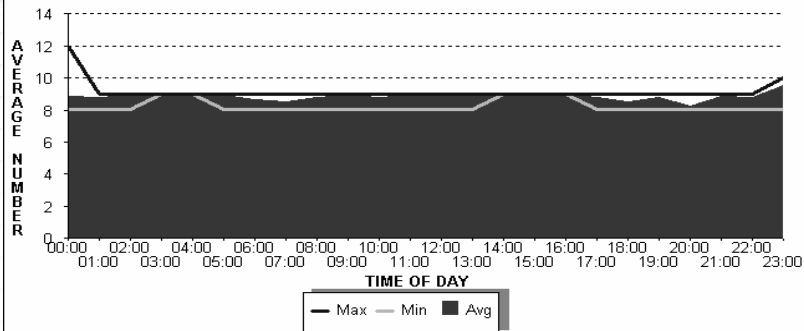


Day of 08/15/03

**z/VM CPU Queuing  
\*MONITOR (D0,R10)**



***Dispatch List  
VMSYSA***



Day of 06/12/03



## z/VM LPAR Information \*MONITOR (D0,R15-16)



PerfMan for z/VM LPAR Information - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Address [C:\ntmlocat\NSMHTML\\_TEMP\zVMLPARS.HTML](C:\ntmlocat\NSMHTML_TEMP\zVMLPARS.HTML) Go Links »

### LPAR Information

In a logically partitioned system the user defines a number of logical partitions (LPARs). Each LPAR, identified by name and number, is assigned a number of logical processors (LCP's) and a weight. Also, for each LPAR a specification of wait completion (Yes or No) and capped (Yes or No) is specified.

LPAR Name	LPAR Number	Primary?	No. of LCP's	Weight	Weight %	Wait Completion?	Capped?
CF01	4	No	1	24	2.4	No	No
CF02	5	No	1	24	2.4	No	No
MVSTEST	6	No	1	24	2.4	No	No
VMSYSA	7	Yes	1	142	14.2	No	Yes
VMSYSB	8	No	1	165	16.5	No	Yes
MVSA	10	No	1	56	5.6	No	No
MVSB	11	No	1	56	5.6	No	No
MVSC	12	No	3	509	50.9	No	No


Done My Computer

## Data Presentation DASD Subsystem




- I/O Rates
- I/O Response Times
- Cache Metrics
- For One or Multiple devices
- Some metrics will reflect z/OS or z/VM architecture differences


**z/OS I/O Rates by Device  
RMF Type 74\_1**




**z/VM I/O Rates by Device  
\*MONITOR (D6,R3)**




**z/VM Avoided I/O  
\*MONITOR (D6,R3)**




**z/OS I/O Response Time  
RMF Type 74\_1**




**z/VM I/O Response Time**  
\*MONITOR (D6,R3)



**z/OS Cache Metrics by Dev**  
RMF Type 74\_5



z/OS Cache Metrics  
RMF Type 74\_5



z/VM Cache Metrics  
\*MONITOR (D6,R4)




- \*MONITOR data contains same cache metrics as RMF
  - Read and Write content
  - Read Hit %'s
  - Write Hit %'s
  - Sequential I/O
  - Inhibit and Bypass cache I/O
- Represents I/O activity from all sharing systems
- Can be presented at Device or Controller level

## Data Presentation Workloads

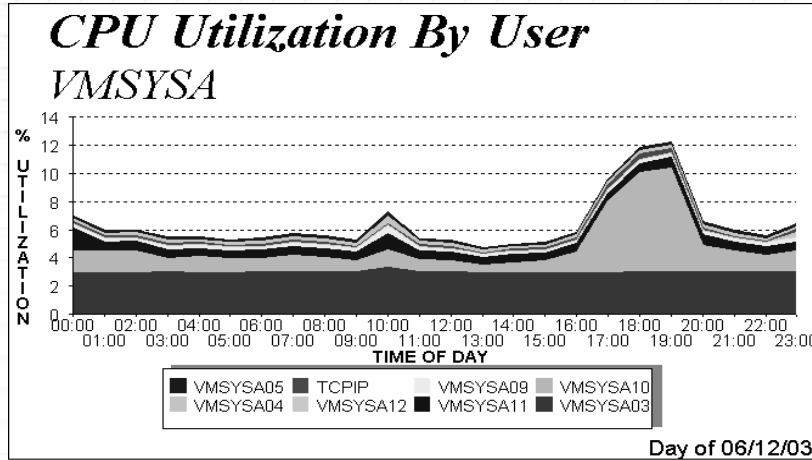


- In z/OS: SC's, SCP's, RC's, RCP's and Workloads
- In z/VM: Users or Virtual Machines
- Resource Usage and Responsiveness
  - CPU Time
  - DASD I/O Rates and Response Times
  - Central and Expanded Storage
- Sampling Profiles
  - Using Samples
    - CPU, DASD and Other
  - Delay Samples
    - CPU, DASD and Many Other

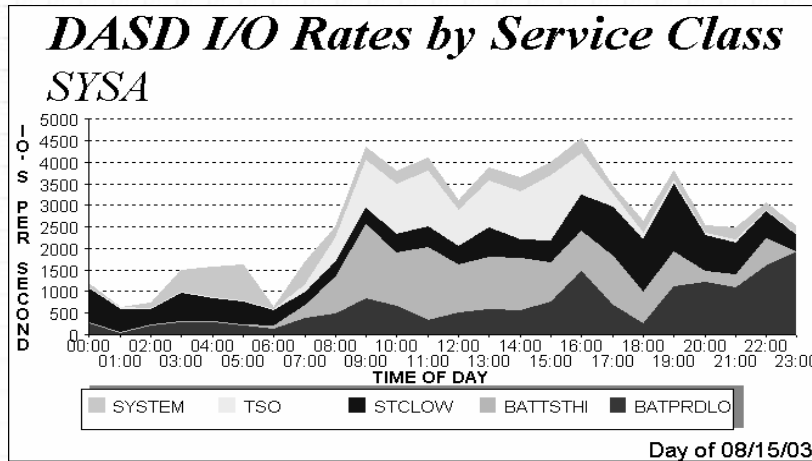
## z/OS CPU Usage by SC RMF Type 72



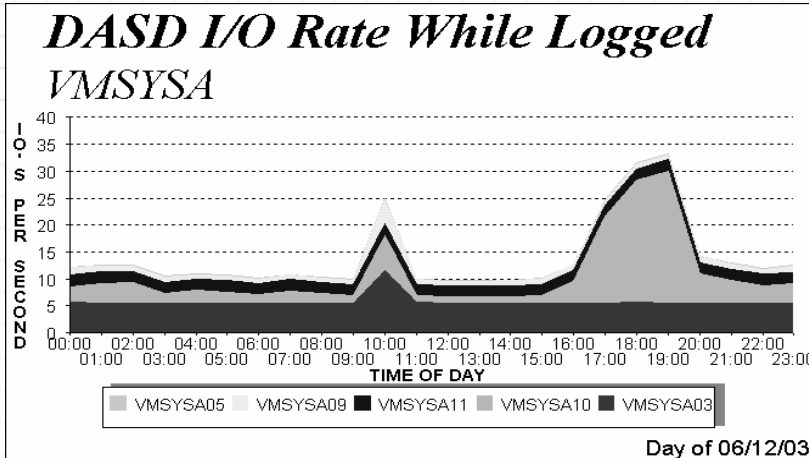
**z/VM CPU Usage by User  
\*MONITOR (D4,R3)**



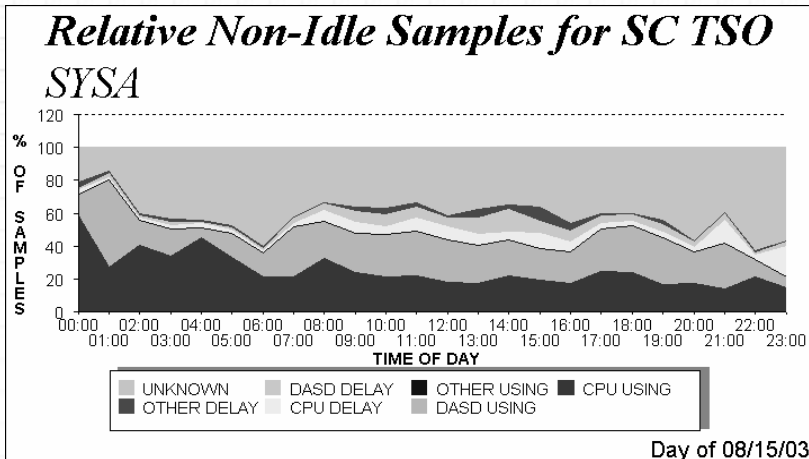
**z/OS DASD I/O by SC  
RMF Type 72**



**z/VM DASD I/O by User  
\*MONITOR (D4,R3)**



**z/OS WLM Sampling  
RMF Type 72**

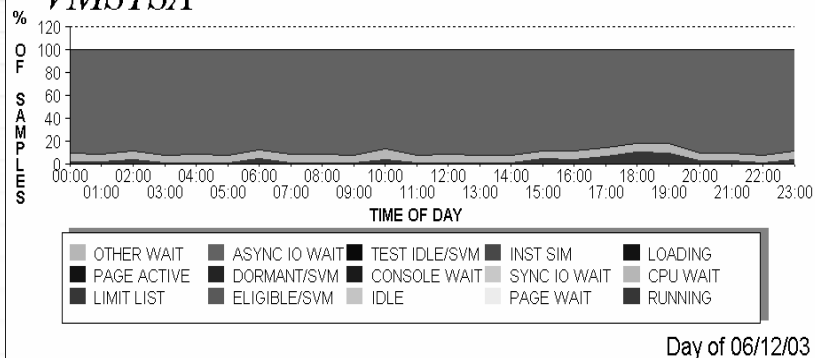




## z/VM Sampling \*MONITOR ( D4,R4)



### True Non-Dormant Time for VMSYSA10 VMSYSA



## Summary



- \*MONITOR is a rich performance data source for z/VM and it's part of z/VM
- \*MONITOR data is different in structure than RMF but record content is similar
- Presentation of Performance data can be very similar for z/OS and z/VM
- References:
  - z/VM Monitor Records 4.1.0 (SC24-6012)
  - z/VM Performance 4.3.0 (SC24-5999)