

z/VM System Performance Tool Review

WAVV 2004

Bill Bitner
bitnerb@us.ibm.com

last updated: April 25, 2004

Legal Stuff

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "as is" basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used instead.

Any performance data contained in this document was determined in a controlled environment and, therefore, the results which may be obtained in other operating environments may vary significantly.

Users of this document should verify the applicable data for their specific environments.

It is possible that this material may contain references to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country or not yet announced by IBM. Such references or information should not be construed to mean that IBM intends to announce such IBM products, programming, or services.

Should the speaker start getting too silly, IBM will deny any knowledge of his association with the corporation.

The following are **Trademarks** of the IBM Corporation:

VM/ESA, e-business logo*, HiperSockets, IBM*, IBM logo*,
IBM eServer, RAMAC*, TotalStorage, z/OS, z/VM, zSeries
LINUX is a registered trademark of Linus Torvalds

Performance Product Strategy

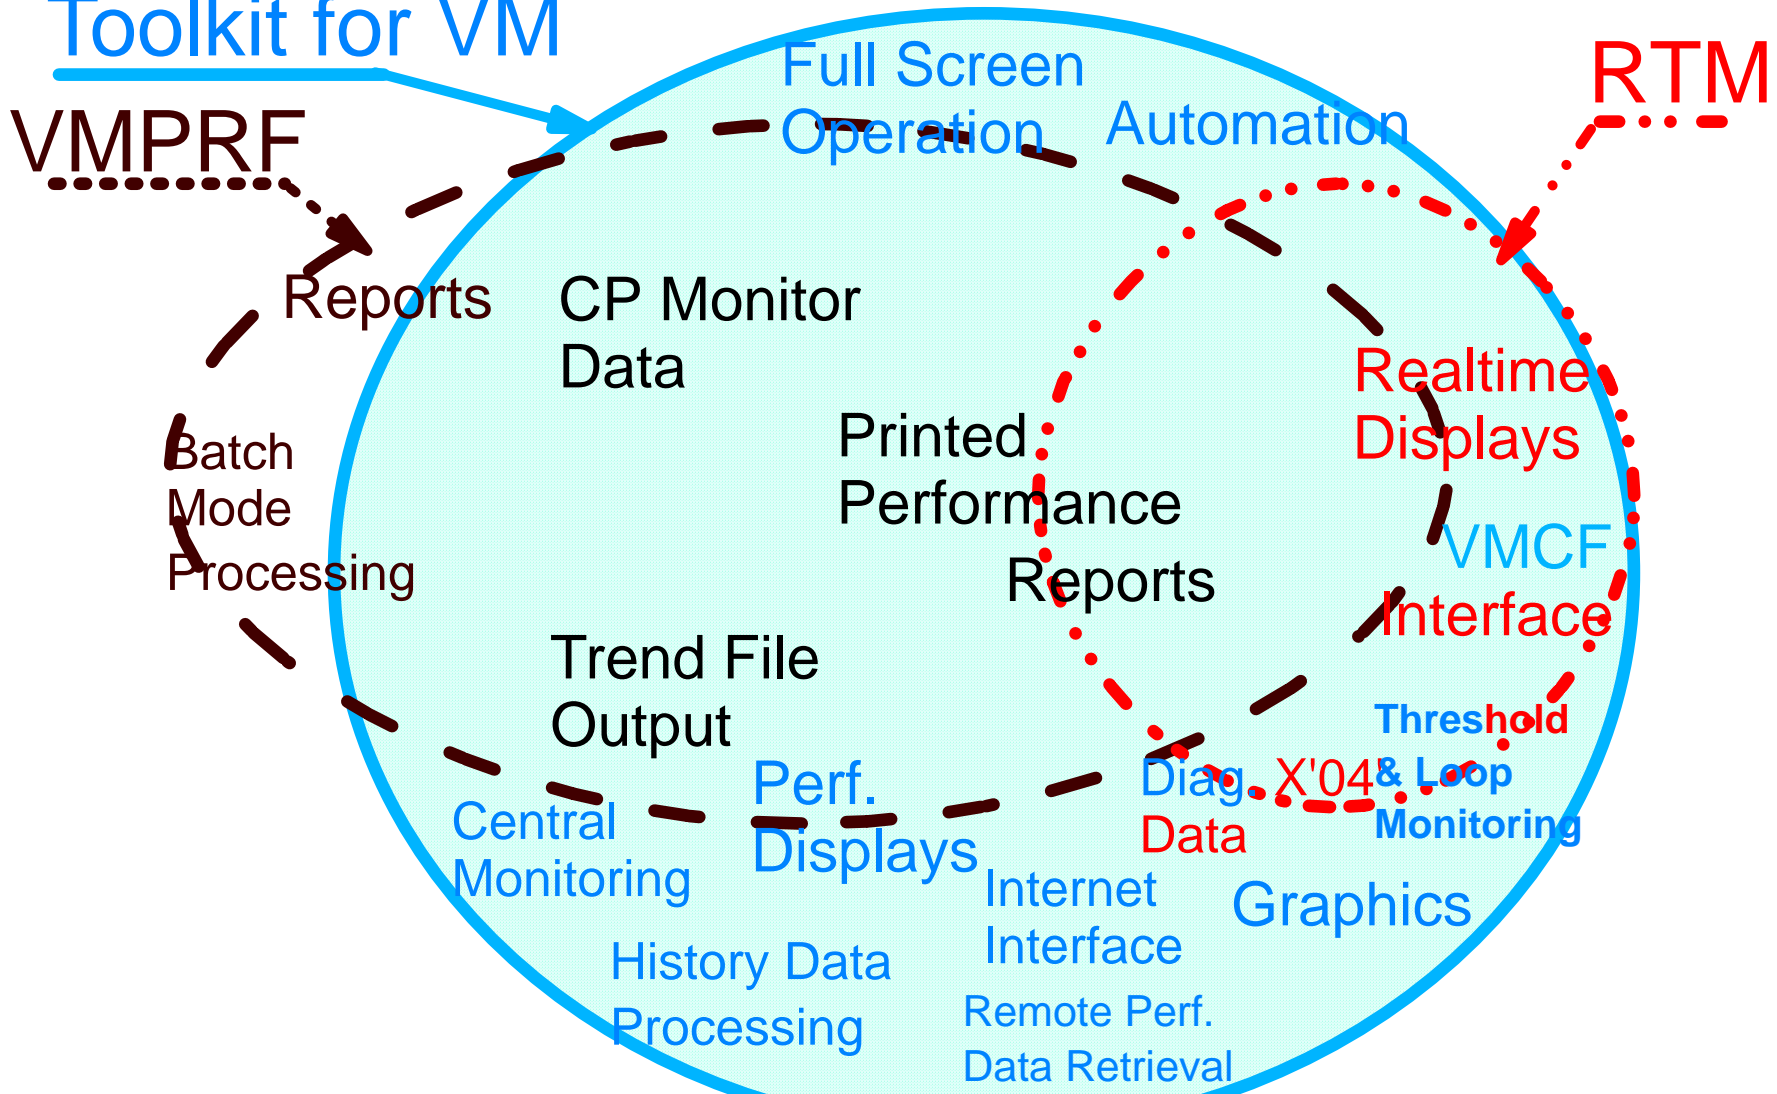
- VMPRF and RTM will be phased out
 - ▶ high development costs
- FCON/ESA phased in as Performance Toolkit for VM
 - ▶ adds significant new function
- Continue to encourage vendor activity
 - ▶ competition breeds excellence
 - ▶ greater percentage of customer needs met

Performance Toolkit Naming

- FCON = Full Screen Operator Console
 - ▶ FCON/XA, FCON/ESA
- FCX = 3 letter module prefix
 - ▶ used in messages, displays, etc.
- Performance Toolkit for VM = full name
- PERFKIT = module that invokes it
- PERFSVM = default userid it runs in

4.4.0 Comparison with VMPRF and RTM

Performance Toolkit for VM



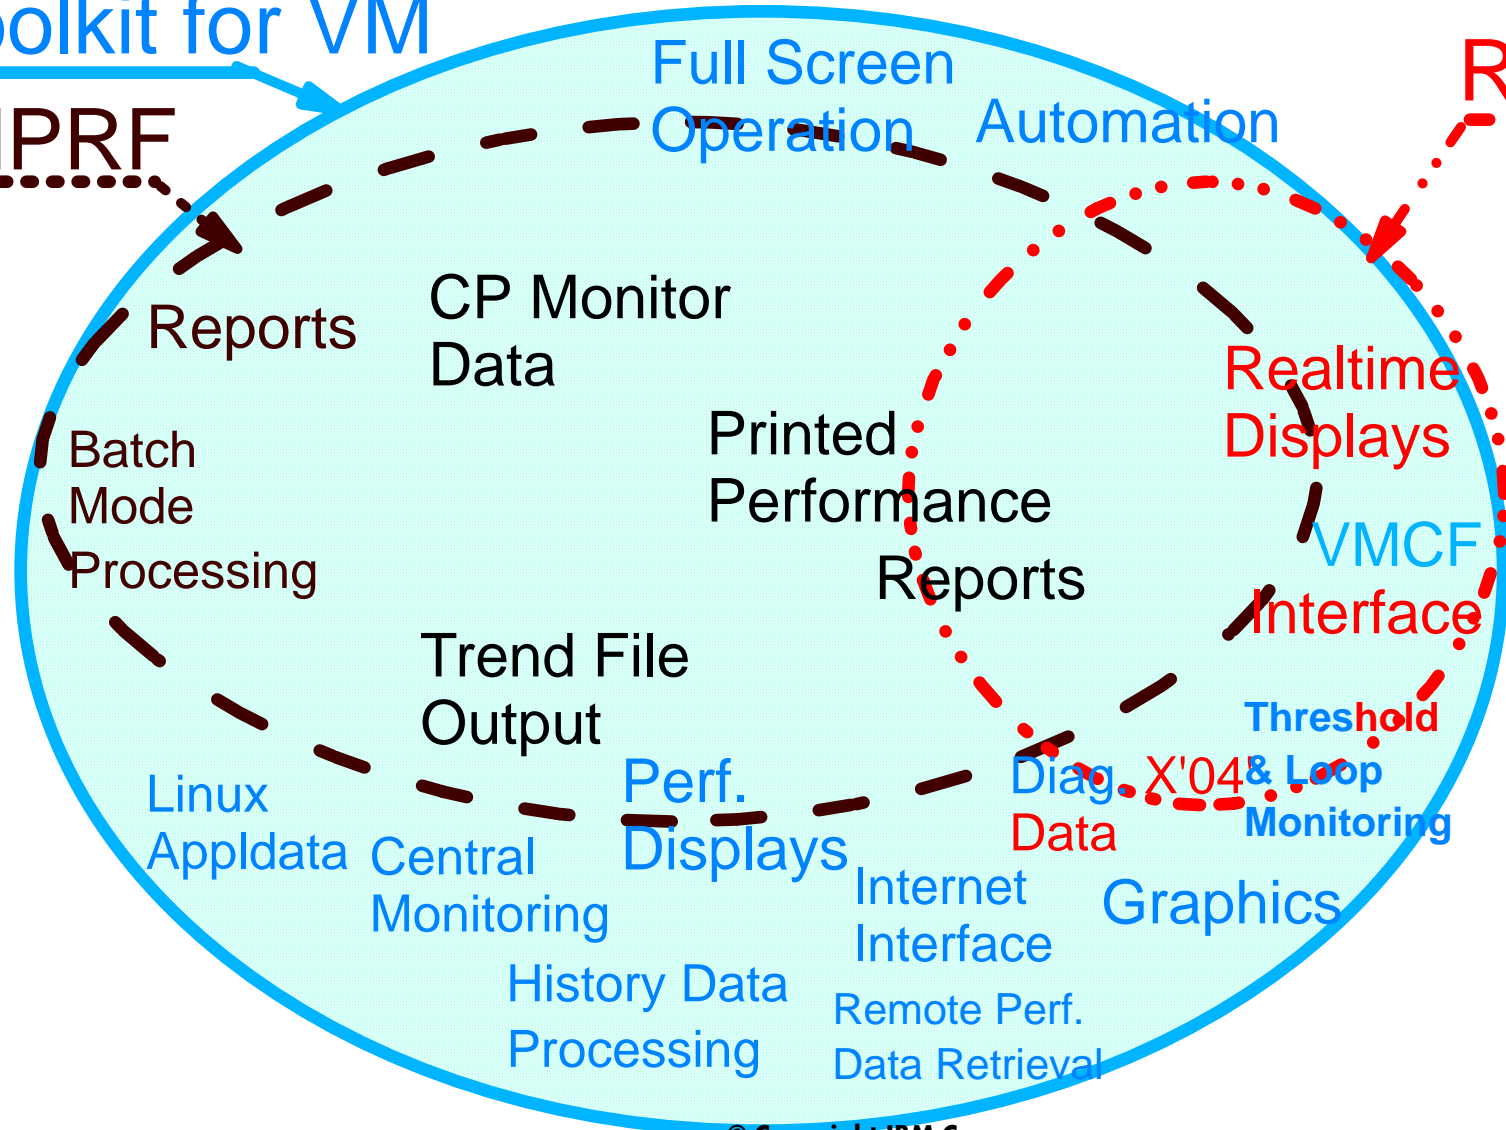
5.1.0 Comparison with VMPRF and RTM

Performance

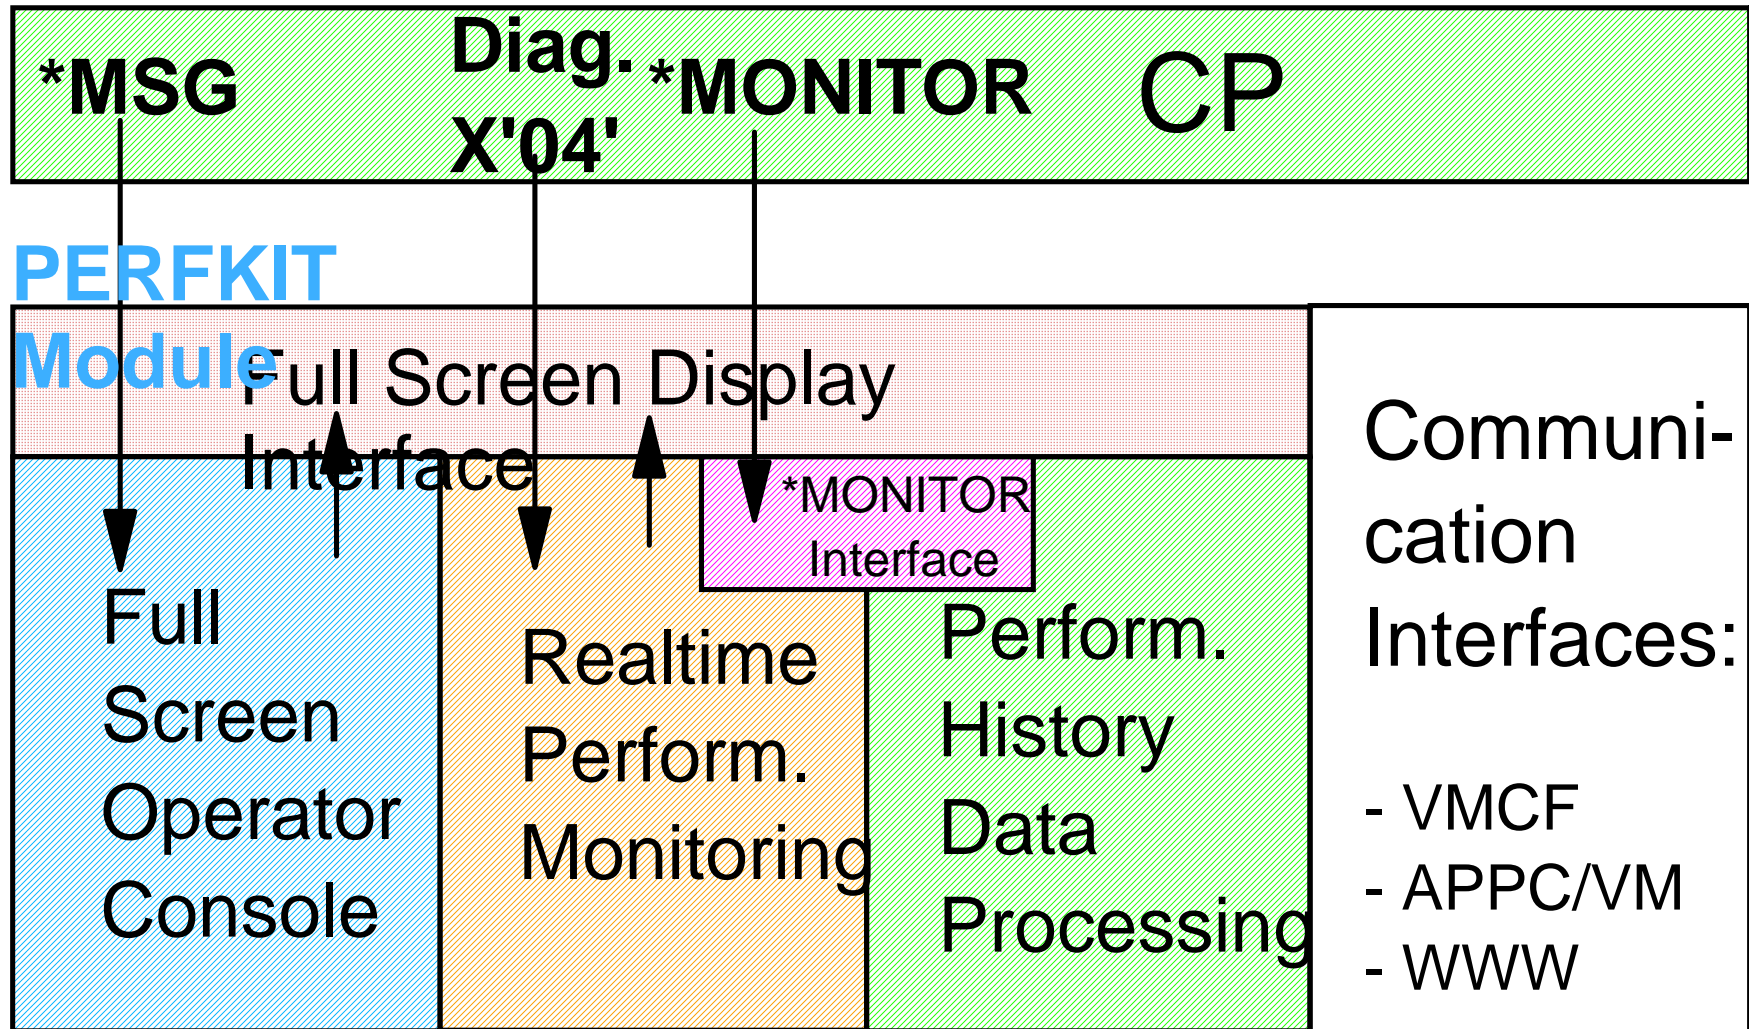
Toolkit for VM

VMPRF

RTM



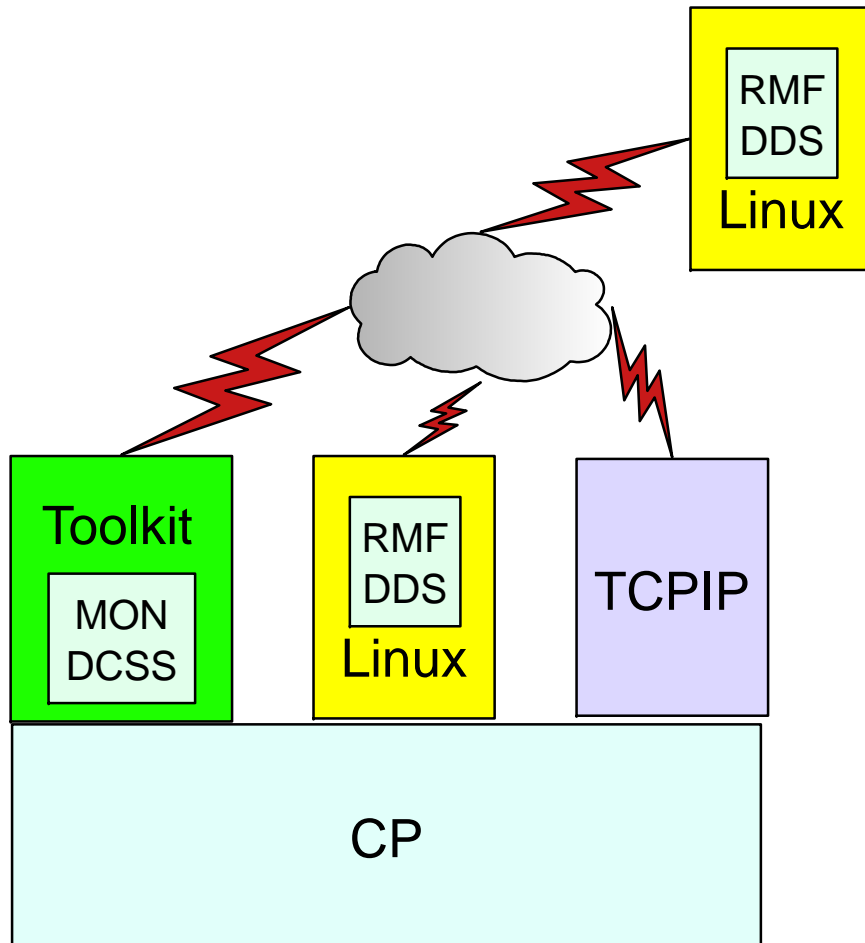
The PERFKIT Module



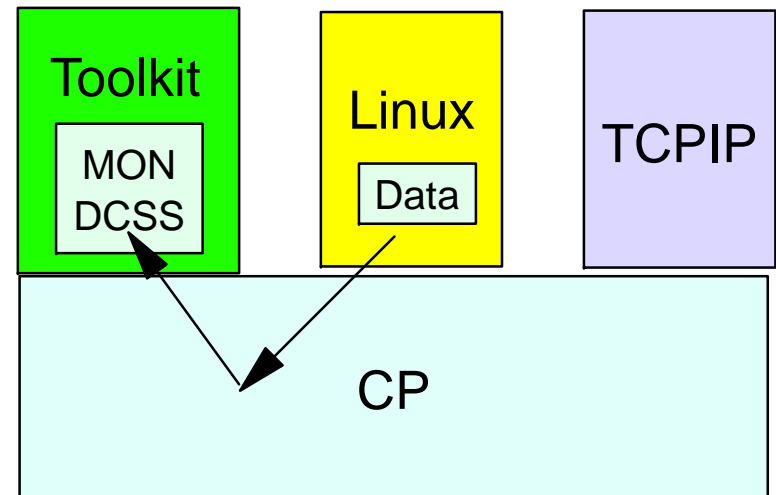
Linux Performance Data

- RMF DDS Interface
 - Originally developed for use with RMF PM
 - Permanent Data Collection in Linux
 - History Data Saved in Linux
 - Selective 'ad hoc' Retrieval via TCP/IP
 - XML data retrieval requests - only data for requested report are retrieved
 - Linux systems not necessarily under same VM
- VM Monitor Application Data (APPLDATA)
 - Linux use of Diagnose x'DC' in January 2004 code drop
 - Allows high level overview of performance
 - Very low overhead for data collecton

Linux Performance Data Collection



RMF DDS



VM Monitor Appldata

USER userid

User Resource Details

FCX115 CPU 2064 SER 51524 Interval 03:16:14 - 03:16:45 Perf. Monitor

Detailed data for user LXM00001

Total CPU	: 4.5%	Storage def.	: 500MB	Page fault rate:	.0/s
Superv. CPU	: 1.5%	Resident <2GB:	35988	Page read rate	: .0/s
Emulat. CPU	: 3.0%	Resident >2GB:	90717	Page write rate:	.0/s
VF total	:%	Proj. WSET	: 126678	Pgs moved >2GB>:	.0/s
VF overhead	:%	Reserved pgs	: 0	Main > XSTORE	: .0/s
VF emulation:%	Locked pages	: 10	XSTORE > main	: .0/s
VF load rate:/s	XSTORE dedic.:	0MB	XSTORE > DASD	: .0/s
I/O rate	: 5.6/s	XSTORE pages	: 0	SPOOL pg reads	: .0/s
DASD IO rate:	5.6/s	DASD slots	: 0	SPOOL pg writes:	.0/s
UR I/O rate	: .0/s	IUCV X-fer/s	: .0/s	MDC insert rate:	.0/s
Diag. X'98'	: .0/s	Share	: 3000	MDC I/O avoided:	.0/s
*BLOCKIO	: .0/s	Max. share	: ...		

...

Proc.	%CPU	%CP	%EM	%VECT	%VOHD	%VEMU	VLD/S	IO/S	Status
00	1.2	.4	.8	5.5	ESA,P03,IOWT
01	1.2	.4	.80	ESA,P03,PSWT
02	.8	.2	.51	ESA,P03,IOWT
03	1.3	.4	.91	ESA,P03,IOWT

...

LXCPU userid

LINUX CPU Utilization Details

FCX230 CPU 9672 SER 15585 Interval 01:33:00 - 01:34:00 Perf. Monitor

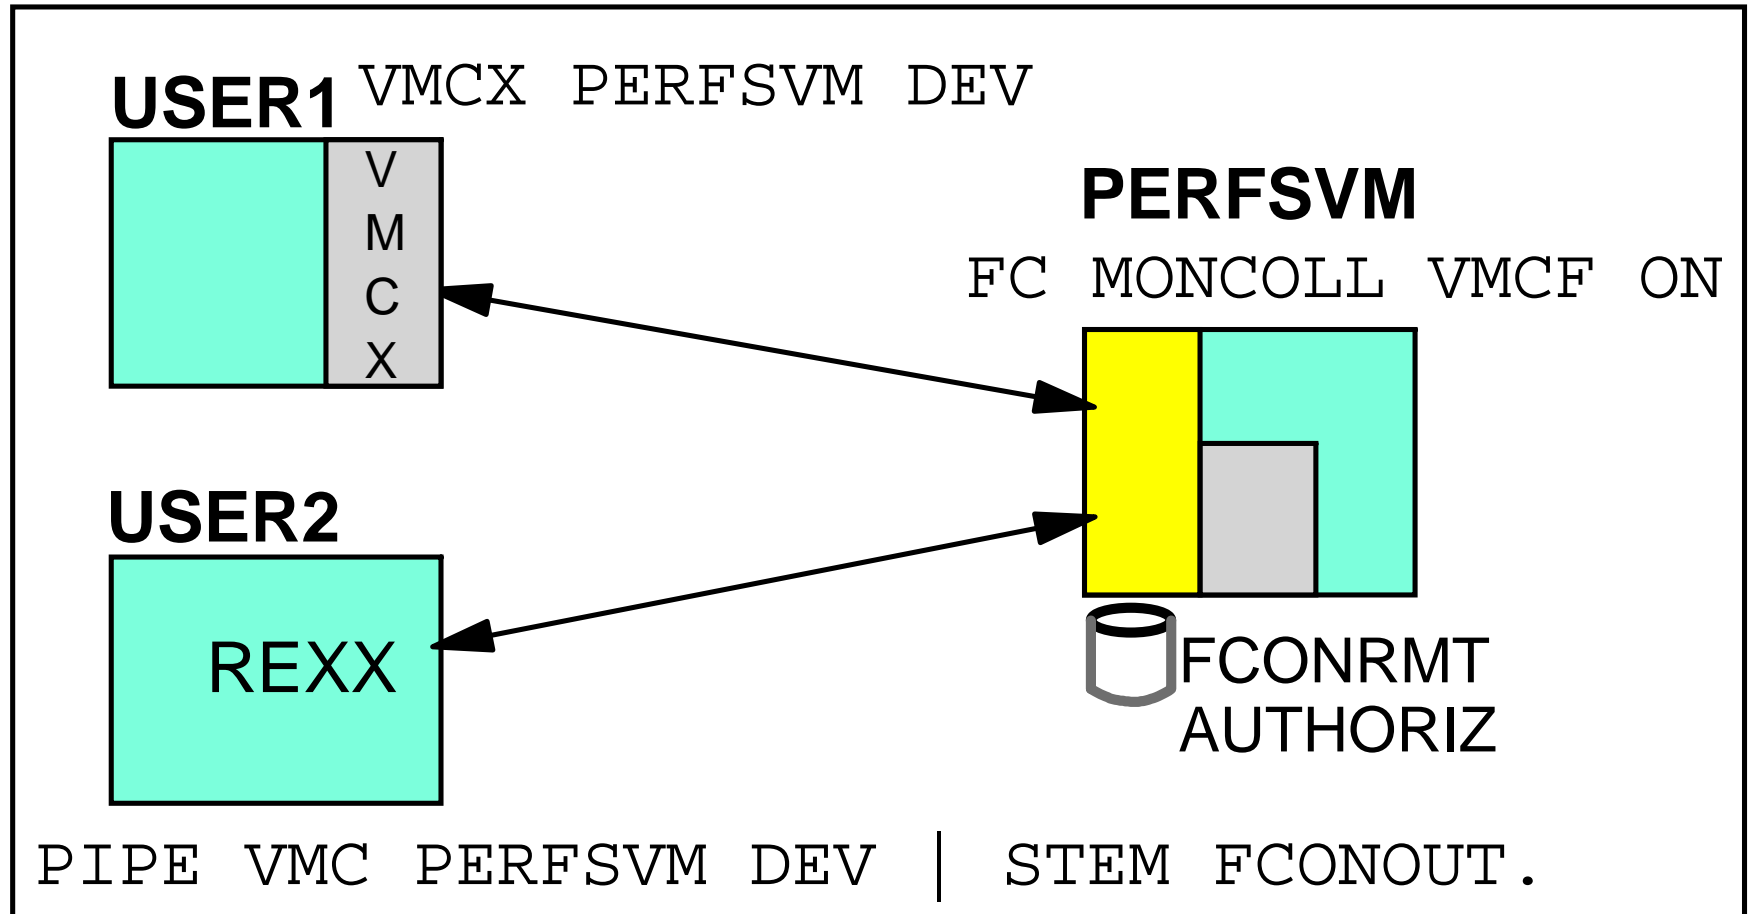
Linux CPU Utilization for System LXM00001

Processor	<--- Percent CPU Utilization --->					<-Accumulated (s)->		
	Total	User	Kernel	Nice	Idle	TotTm	UserTm	KernTm
>>Mean>>	0.70	0.06	0.64	0	99.29	---	---	---
cpu0	0.78	0.06	0.71	0	99.21	---	---	---
cpu1	0.73	0.18	0.54	0	99.26	---	---	---
cpu2	0.48	0	0.48	0	99.51	---	---	---
cpu3	0.86	0.01	0.84	0	99.13	---	---	---
Process Name								
syslogd.293	0.78	0	0.78	0	---	3657	36.08	3621
nmbd.499	0.46	0.03	0.43	0	---	9166	1649	7517
apachegat.29502	0.3	...	0.3	...	---	353.9	1.48	352.5
gengat.29511	0.3	...	0.3	...	---	498.0	1.37	496.7
procgat.29517	0.23	...	0.23	...	---	370.7	13.47	357.2
httpd.464	0.18	0	0.18	0	---	953.6	18.3	935.3
kupdate.3	0.16	0	0.16	0	---	603.9	...	603.9
filegat.29508	0.06	...	0.06	...	---	89.36	0.21	89.15
...								

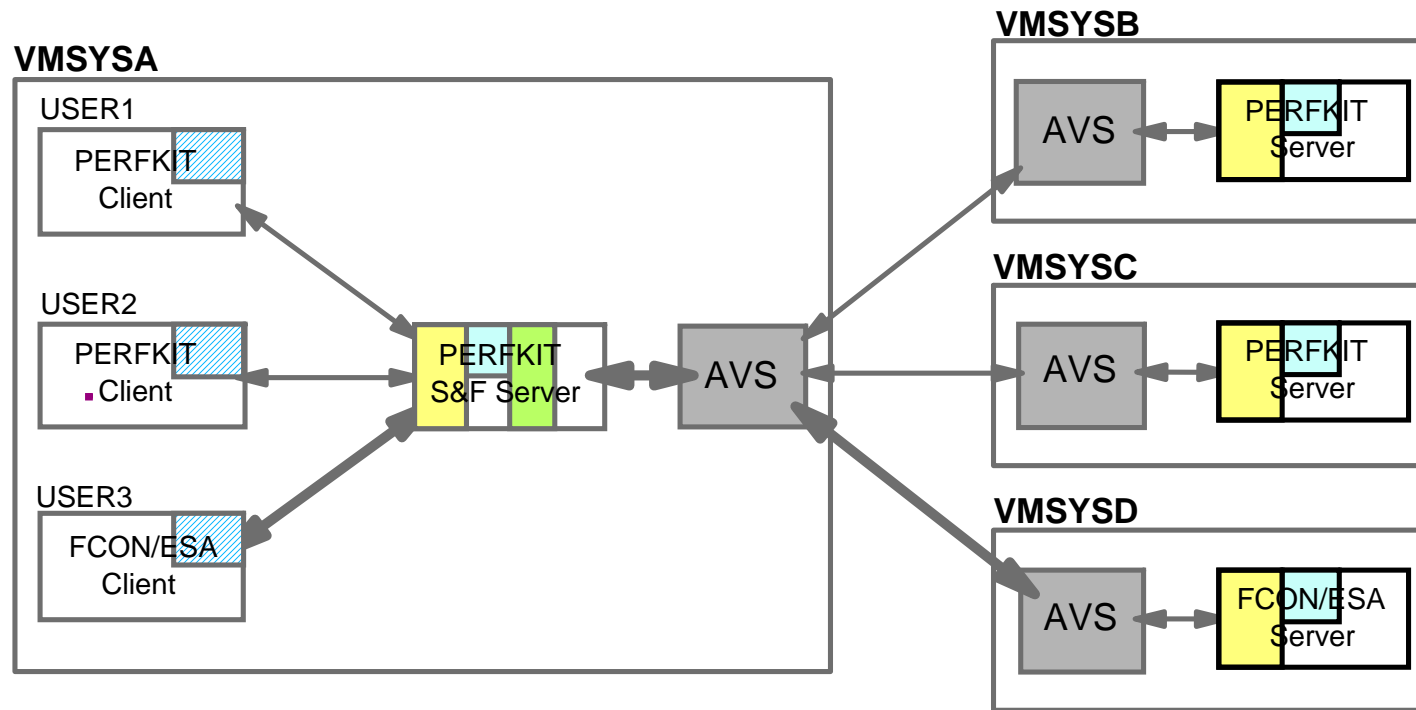
Multiple (Remote) Access to Performance Data

- Local VMCF Interface
 - Patterned after RTM VMC interface
 - Useful when used with CMS Pipelines
- APPC/VM Local & Remote
- WWW Interface for Standard Web Browsers
 - Only serves Performance Toolkit data

Local VMCF Interface

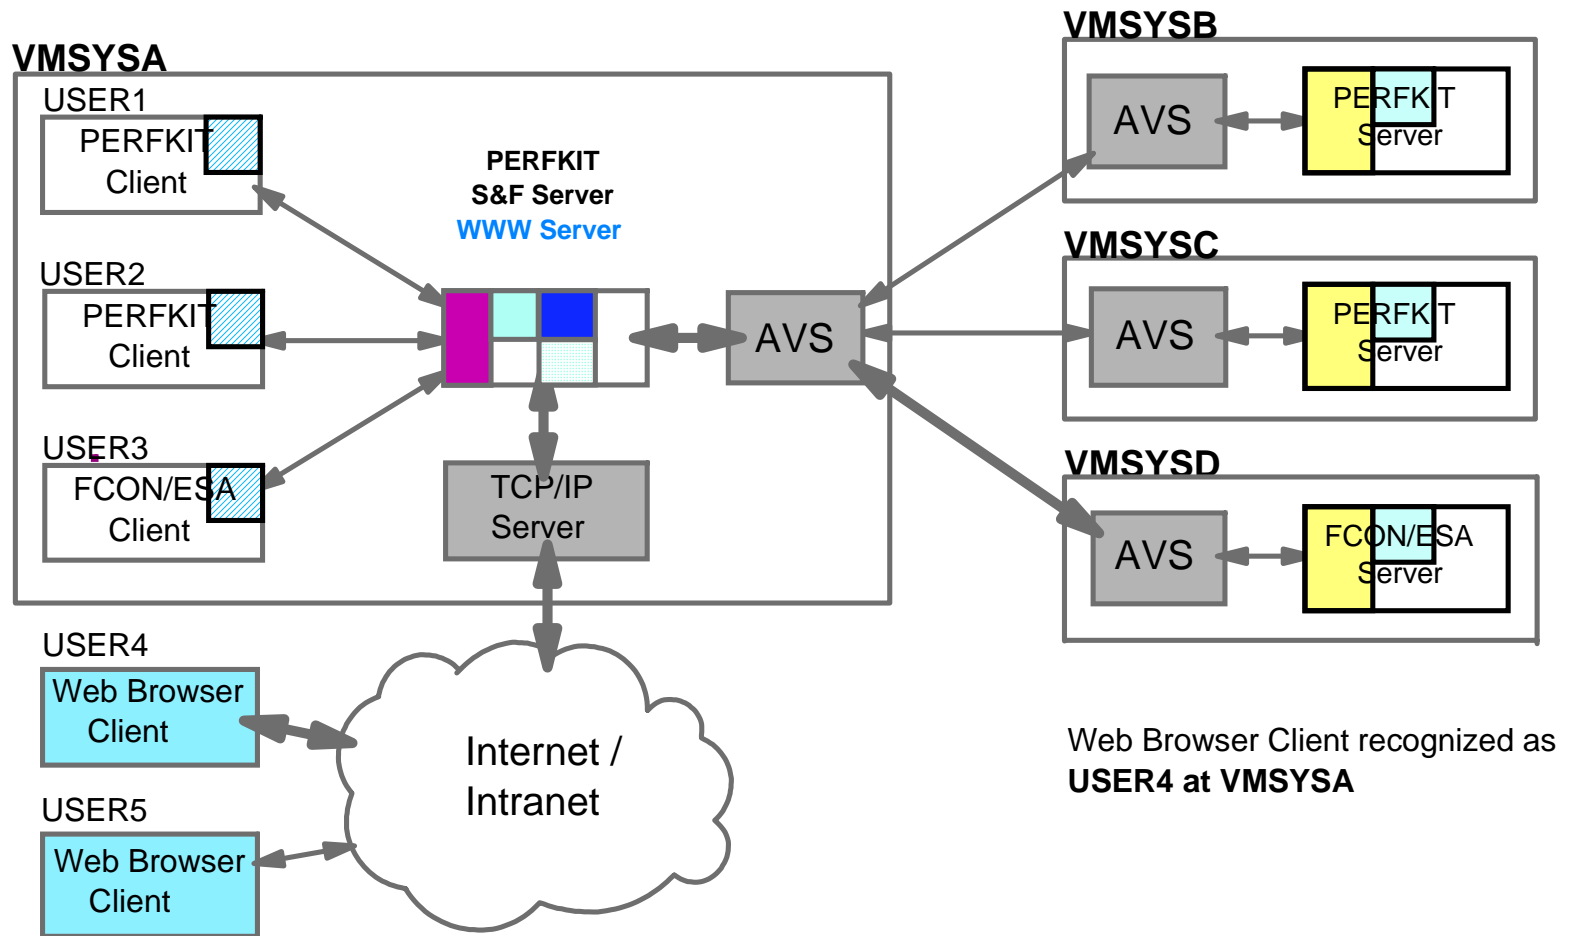


APPC/VM: Store & Forward Logic



Client recognized as
USER3 at VMSYSA

WWW Server Interface



Web Browser Client recognized as **USER4** at **VMSYSA**

Initial Perf. Data Selection Menu

Initial Performance Data Selection Menu - Netscape

File Edit View Go Communicator Help

Location: 8DE/a?v=%2F00A23B18%2F88DE%2F&form=02&s=Menu

FCON IBM ESA v.3.2

Data Retrieval Session with System VM4ALL
Initial Performance Data Selection Menu

Command Refresh Systems Forw Help Auto-Refresh

Performance Screen Selection

General System Data	I/O Data	History Data (by Time)
1. CPU load and trans.	11. Channel load	31. Graphics selection
2. Storage utilization	12. Control units	32. History data files*
3. Storage subpools	13. I/O device load*	33. Benchmark displays*
4. Priv. operations	14. CP owned disks*	34. Correlation coeff.
5. System counters	15. CU-cached disks*	35. System overview
6. CP IUCV services	16. Cache extend. func.*	36. Auxiliary storage
7. SPOOL file display*	17. DASD I/O assist	37. CP communications*
8. LPAR data	18. DASD seek distance*	38. DASD load
9. Shared segments	19. DASD seek locations*	39. Minidisk cache*
A. Shared data spaces	1A. I/O configuration	3A. Paging activity
B. Virt. disks in stor.	1B. I/O config. changes	3B. Processor load
C. Transact. statistics	User Data	3C. Logical part. load
D. Monitor data	21. User resource usage*	3D. Response time (all)*
		3E. RSK data menu*

Example for Performance Data Display

Hyperlink selection of:

Sort sequence

Context help

Device details

General I/O Device Load and Performance - Netscape

File Edit View Go Communicator Help

FCON IBM
ESA
V.3.2

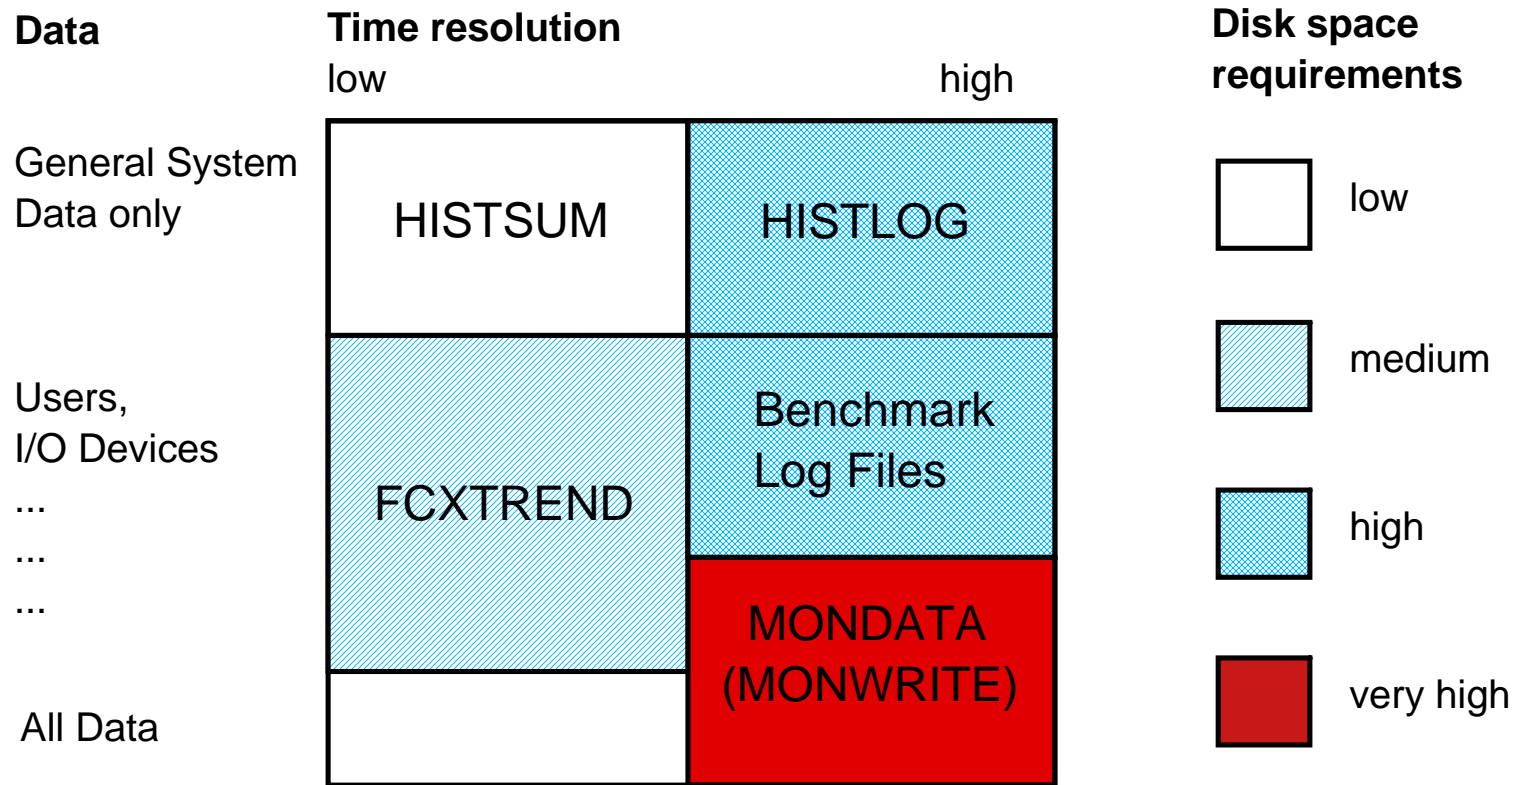
Data Retrieval Session with System VM4ALL
Select a device for I/O device details

Command Refresh Systems Menu Forw

Interval 15:36:35-15:41:37, on 2000/04/06 (Select [average](#) for

<-- Device Descr. -->	Mdisk	Pa-	<-Rate/s->	<----- Time (msec			
<u>Addr</u> <u>Type</u> <u>Label/ID</u>	<u>Links</u>	<u>ths</u>	<u>I/O</u> <u>Avoid</u>	<u>Pend</u>	<u>Disc</u>	<u>Conn</u>	<u>Serv</u>
>> All DASD <<0 .0	.2	.9	3.5	4.6
<u>0200</u> 3380-K XDSK01	0	8	.0 .0	.3	.1	.4	.8
<u>0201</u> 3380-K 1IBK31	8	8	.1 .0	.4	.5	2.7	3.6
<u>0202</u> 3380-K 1IBK32	9	8	.0 .3	.1	.1	.4	.6
<u>0203</u> 3380-K 1IBK33	18	8	.1 .0	.4	.3	3.1	3.8
<u>0204</u> 3380-K 1IBK34	9	8	.0 .0	.3	.1	.4	.8
<u>0205</u> 3380-K 1IBK35	10	8	.1 .0	.3	.1	3.5	3.9
<u>0206</u> 3380-K 1IBK36	11	8	.1 .0	.2	.7	3.8	4.7
<u>0207</u> 3380-K 1IBK37	7	8	.1 .0	.4	.6	3.5	4.5

History Data Files



Graphics

- Simple Plots with Commands **PLOT...**
 - No additional graphics SW required
- GDDM Line Graphics with Commands **GRAPH...**
 - Requires GDDM on the system where graphics are to be shown
- Line Graphics with Java Applet via WWW Interface
 - Based on graphics capability of WS and Web Browser's Java support
 - No additional graphics SW required



Data Retrieval Session with System VM4ALL
Graphics Selection Menu

Command

Refresh

Systems

Menu

Return

Help

Auto-Refresh

General Specifications

Output format

Line graphics

Validate

Submit

Data origin

File VM4ALL FCXTREND A

Graphics type

Summary graphics (coarse time scale)

Selected period

All data

Selected days

All days

Selected hours

All hours

Variable Selection

X-Variable

IO/S

SSCH/RSCH Rate

SYSTEM

Truncate at

Y-Variables

1

IO/S

SSCH/RSCH Rate

SYSTEM

2

TOTAL

Total I/O Rate

USER

TCPIP

3

Trend Record Selection

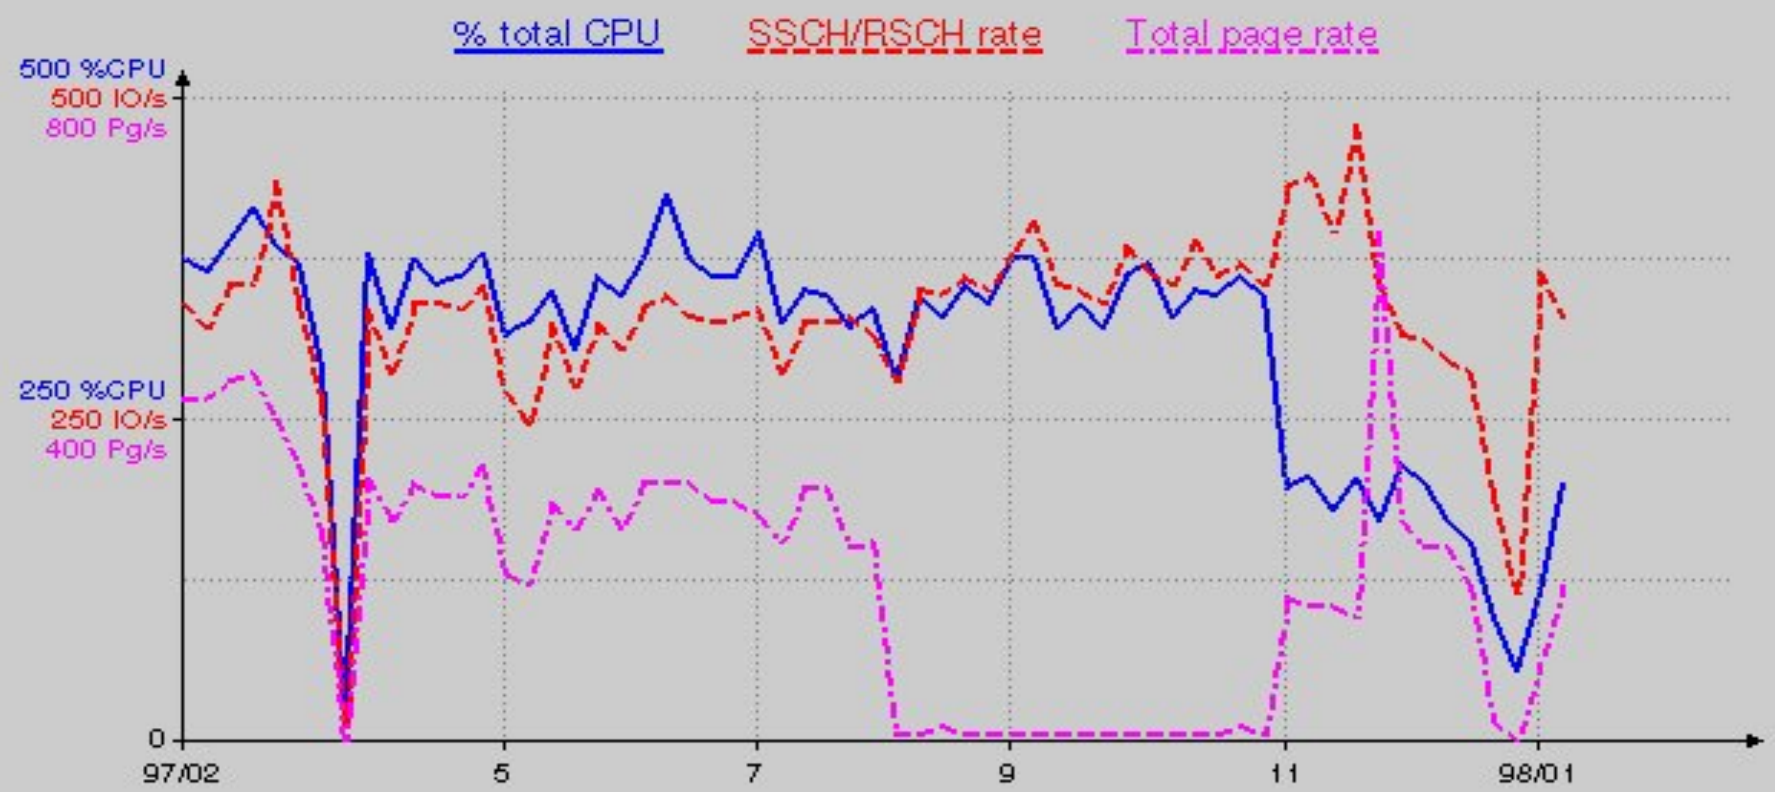
Display

Record Details



Data Retrieval Session with System CHVM1
Graphical Representation of Performance Data

Command Refresh Systems Menu Return Help

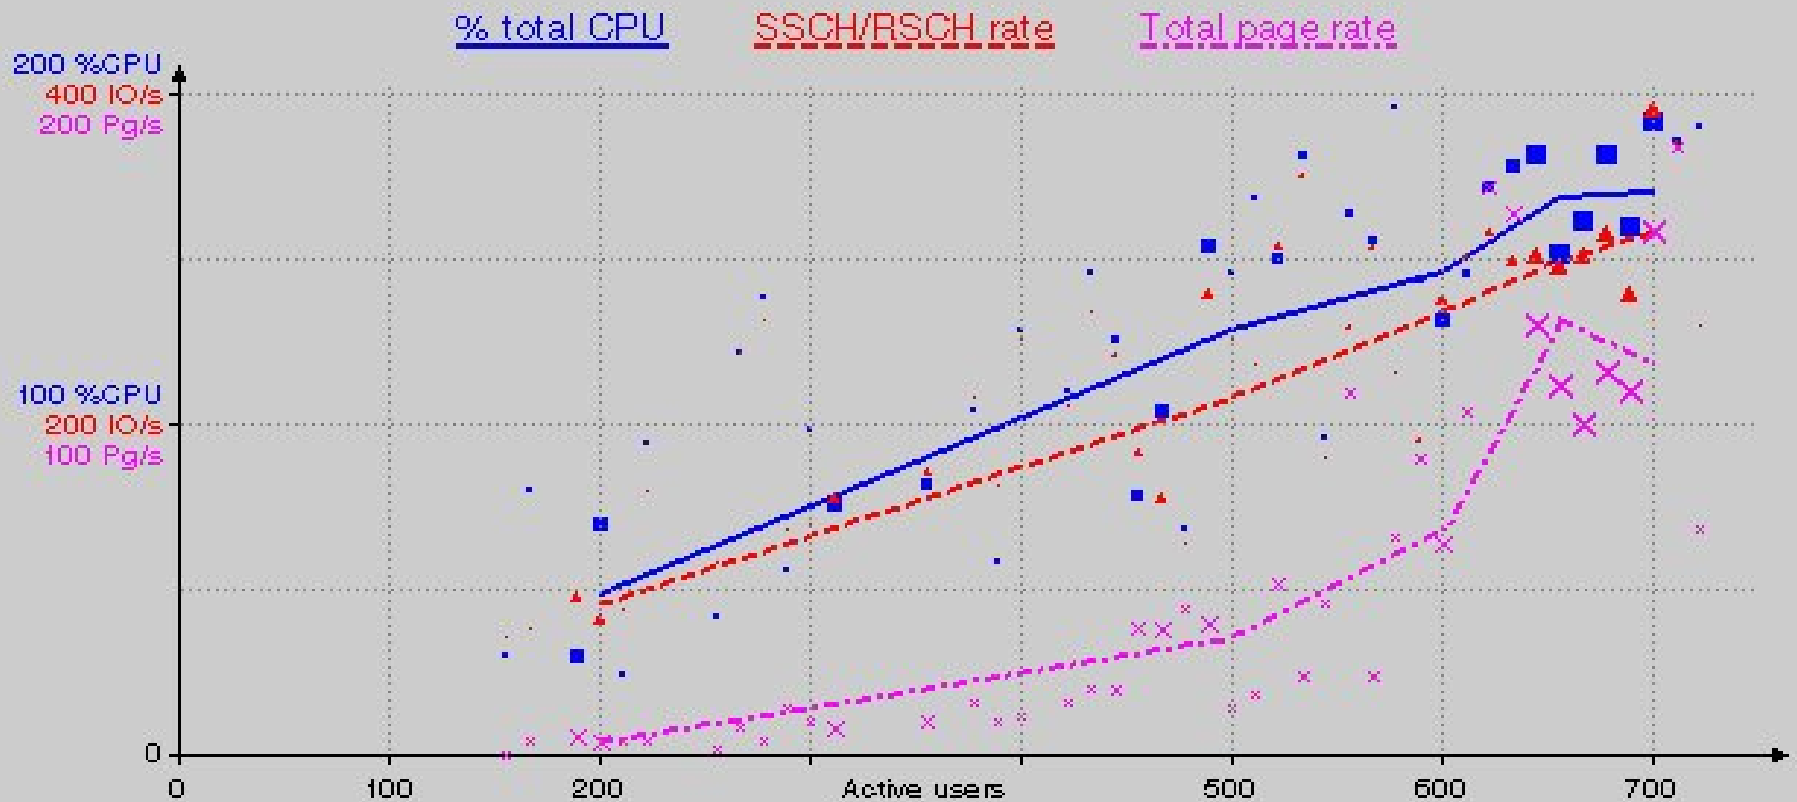


Source data: ACUM HISTSUM A
Sel. days: Monday-Friday
Sel. hours: 08 09 10 11 13 14 15 16



Data Retrieval Session with System CHVM1
Graphical Representation of Performance Data

Command Refresh Systems Menu Return Help



Source data: 19980113 HISTLOG1 *
Period: 07:00h - 19:00h



Data Retrieval Session with System VM4ALL

Graphical Representation of Performance Data

Command

Refresh

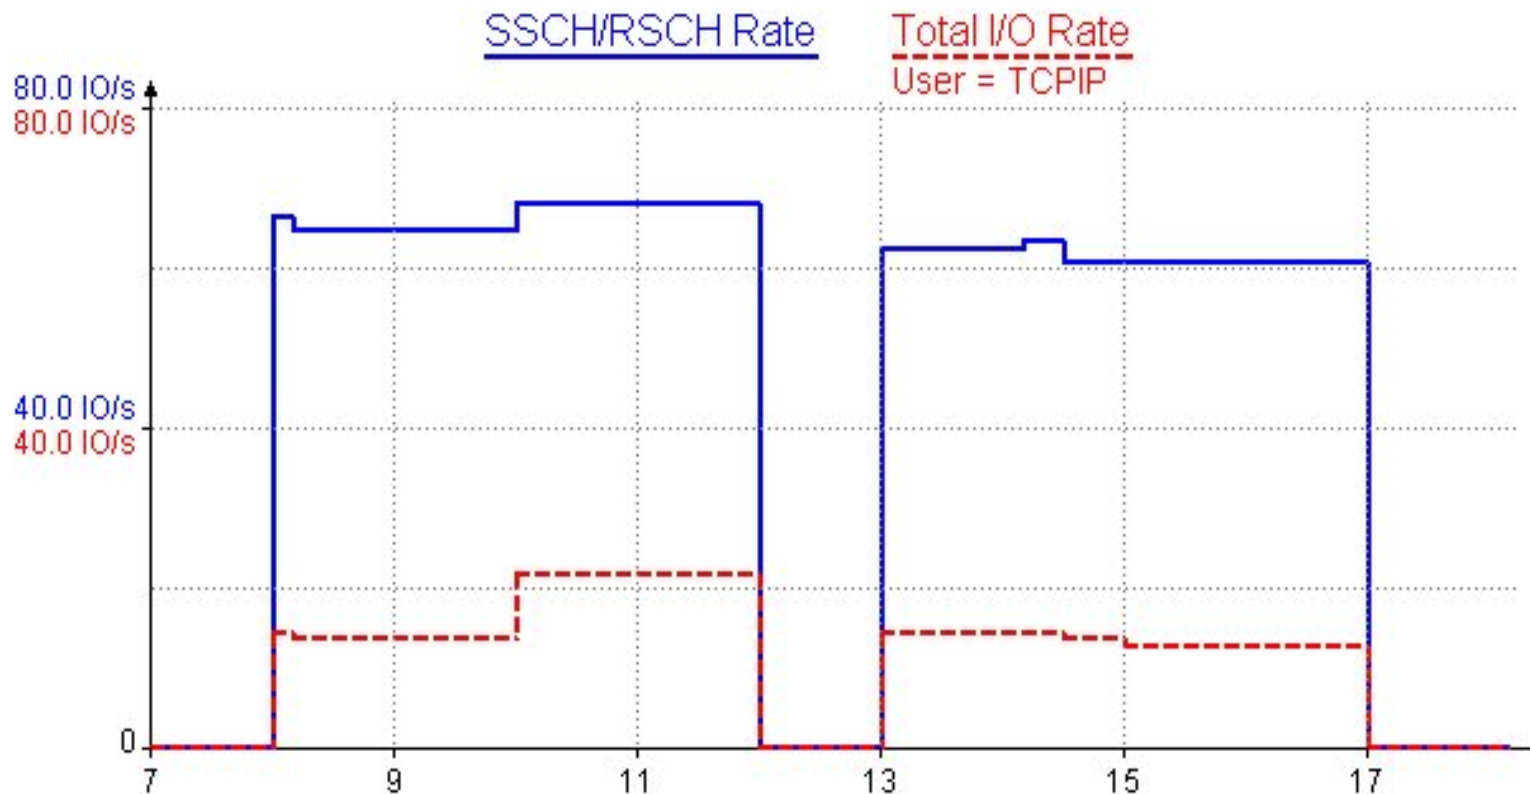
Systems

Menu

Return

Help

Auto-Refresh



Source data: VM4ALL FCXTREND A Sel. days: All
Period: All

IDLEUSER (RTM Compatibility)

FCX238	CPU 2064	SER 51524	Status	06:58:11	Perf. Monitor				
-----	·	-----	·	-----	·				
-----	Min	-----	Min	-----	Min				
Userid	Idle	Userid	Idle	Userid	Idle				
AVS	30	BUCKETS	189	DATAMOVE	2	DGTSRV01	189	DGTSRV02	189
DGTSRV03	189	DIRMAINT	1	FTPSERVB	189	GCS	189	HAMILTJL	189
HARDWARE	189	HOLDER	189	IMAPAUTH	189	IMAP3SRV	8	ISPVM	189
JSWITZER	189	K4SERV	8	LININS	189	LNXMLSTR	189	ML1	189
MULTISRV	8	OPERACCT	6	OPERSYMP	189	OSADMIN1	189	PERFSVM	189
PERFTEST	189	PORTMAP	8	RAICHER	189	RECOVERY	189	RMSMASTR	5
RTMTST4	189	RXAGENT1	8	SFCM1	189	SFSESA	189	SMSMASTR	189
SMSSRV01	189	SNMPD	8	SNMPQE	1	SQLMACH	2	TCPMAINB	12
TCPMAINC	189	TFTPD	8	TOOLS	189	TPOPER	189	VMNFS	3
...									
...									

Selectable with IDLEUSER command

PROCSUM Log (VMPRF Compatibility)

```

FCX239      CPU 2064  SER 51524  Interval 03:49:11 - 07:14:11      Perf. Monitor

      <----- CPU ----->      <----- Spin Lock Activity ----->
      <---Ratio--->                <----- Total ----->      <--- Scheduler --->      <-
Interval    Pct          Cap-  On-  Locks Average    Pct  Locks Average    Pct  st
End Time    Busy    T/V  ture  line  /sec   usec  Spin  /sec   usec  Spin
>>Mean>>    1.0  1.66 .7789  9.0   74.8   1.825 .002  58.4   2.029 .001
05:30:11    1.0  1.71 .7760  9.0   65.9   1.941 .001  50.5   2.208 .001
05:35:11    1.0  1.71 .7672  9.0   69.1   1.881 .001  52.7   2.151 .001
05:40:11    1.0  1.69 .7717  9.0   71.7   1.611 .001  55.2   1.798 .001
05:45:11    1.0  1.69 .7794  9.0   72.5   1.646 .001  56.0   1.834 .001
05:50:11    1.0  1.67 .7748  9.0   71.2   1.578 .001  55.1   1.762 .001
05:55:11    1.0  1.69 .7708  9.0   72.4   1.642 .001  56.9   1.801 .001
06:00:11    1.0  1.68 .7804  9.0   70.0   1.558 .001  55.0   1.714 .001
...
...

```

Similar to SYSTEM_SUMMARY2_BY_TIME
 Selectable with PROCSUM command

Performance Toolkit for VM FL440

- z/VM Version 4.4 provides monitor data for
 - Virtual Switches activity
 - VM Resource Manager operation
 - Processor type information
- New Displays
 - VSWITCH display
 - VMRM display with VM Resource Manager data
- Modified
 - Processor type added on LPAR and LPARLOG displays

VMRM (VM Resource Manager Data)

```
FCX241      Data for 2003/04/07  Interval 10:48:14 - 10:49:14  Perf. Monitor
-----      .
VM Resource Manager      .
Server      Workload      Impor  <-- DASD -->  <-- CPU --->  Active
tance      D-Goal D-Act  C-Goal C-Act  Samples
IRDSVM     WORK1      10     100   ...     100   ...     1
IRDSVM     WORK2      5      50    ...     50    ...     1
IRDSVM     WORK3      1      1     ...     1     ...     1
IRDSVM     WORK4     10     100  100     100   72     1
IRDSVM     WORK5      5      50   100     50    77     1
IRDSVM     WORK6      1      1    100     1     63     1
IRDSVM     WORK7     10     100  100     100   63     1
IRDSVM     WORK8      5      50   100     50     3     1
IRDSVM     WORK9      1      1     ...     1     ...     1
...
...
```

Selectable with VMRM command

References

- General information
 - ▶ <http://www.vm.ibm.com/related/perfkit/>
- Performance Toolkit Book
 - ▶ <http://publibz.boulder.ibm.com/epubs/pdf/hcsl7a00.pdf>
- Comparison to VMPRF
 - ▶ <http://www.vm.ibm.com/related/perfkit/pkitprf.html>
- Comparison to RTM
 - ▶ <http://www.vm.ibm.com/related/perfkit/pkitrtm.html>