IBM Advanced Technical Support

A Beginner's Guide to Measuring and Understanding z/VM Guest Performance

Jon vonWolfersdorf
wolff@us.ibm.com

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: CICS*, DB2*, DB2 Universal Database, e-business logo, IBM, IBM Logo*, IMS, Intelligent Miner*, Lotus*, MQSeries*, Multiprise*, MVS, Net.Commerce, NetView*, On demand business logo, OS/390, Parallel Sysplex*, PassPort Advantage*, PR/SM, Tivoli*, Websphere, z/OS*, zSeries*, z/VM*, zSeries Entry License Charge.

The following are trademarks or registered trademarks of other companies

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

LINUX is a registered trademark of Linux Torvalds

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

Intel is a registered trademark of Intel Corporation

* All other products may be trademarks or registered trademarks of their respective companies.

NOTES:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Acknowledgements

- **With special thanks to:**
  - Bill Bitner
  - Richard Lewis
  - Chuck Morse

Agenda

- **Introduction**
- **Resources that affect z/VM guest performance**
  - What they are?
  - How they can be controlled?
  - How to recognize a problem using the Performance Toolkit for VM ™
- **Summary**

Introduction

Performance Problem Analysis Methodology

- **Performance problems result from high usage and/or contention for key system resources**
  - **CPU**
  - **Memory**
  - **I/O**
- **Begin measuring system performance characteristics before problems occur**
  - **Use Performance Toolkit for VM or other monitoring tools**
  - **Understand performance characteristics of well running system**
- **When a problem occurs:**
  - **Identify the resource in contention**
  - **Reduce or eliminate the contention**
    - **Reduce the demand for the resource**
    - **Increase the resource capacity**
    - **Reallocate the usage of the existing resource**

Performance Toolkit for VM

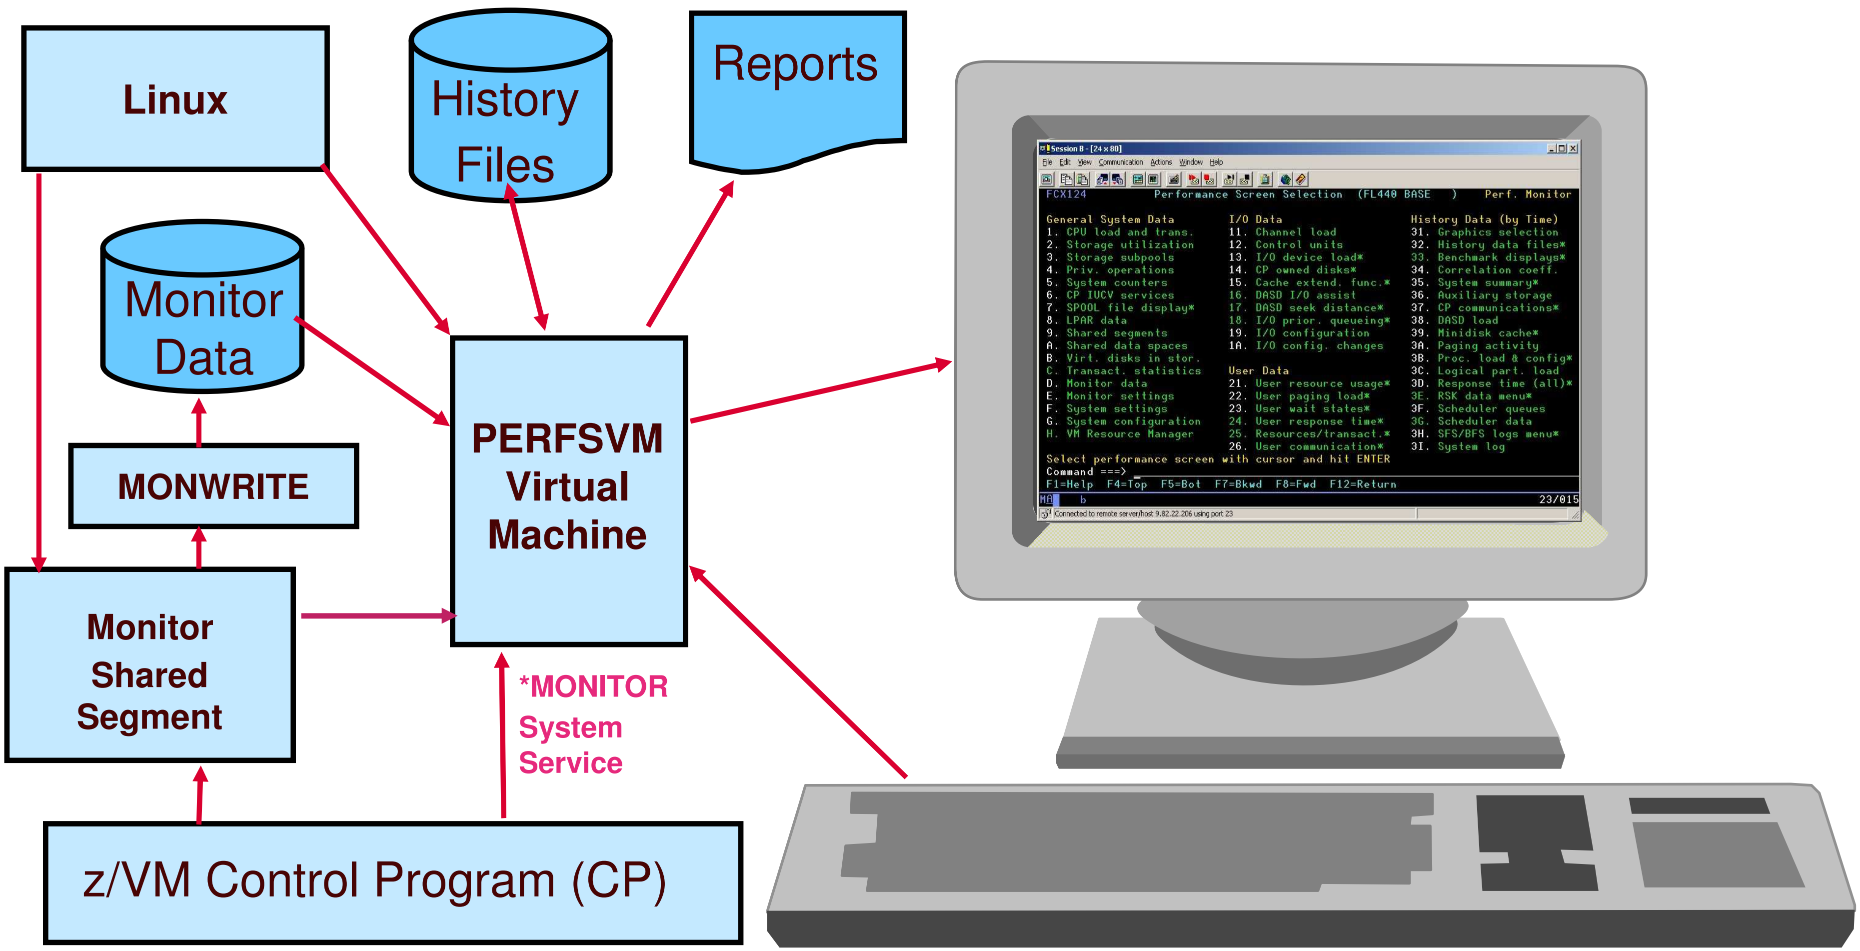

z/VM Internals - Definitions

- **Virtual Machine Definition Block (VMDBK)**
  - One per virtual processor defined to a virtual machine
  - CP representation of a virtual machine
- **Scheduler lists**
  - Dormant list
  - Eligible list
  - Dispatch list
- **Start Interpretive Execution (SIE)**
  - Method z/VM uses to dispatch a virtual processor on a real processor
- **Working Set Size (WSS)**
  - Scheduler estimate of memory pages needed to run a virtual machine

Resources That Affect Performance

- **Processor**
- Memory
- I/O

Physical Processors

- **z/VM5.3 supports up to 32 physical processors**
- **Physical processors may be dedicated to virtual machines**
  - Not commonly done
  - A dedicated processor not available to run other virtual machines
- **z/VM dispatches virtual processors on physical processor using the Start Interpretive Execution (SIE) instruction**
  - Intercept occurs based upon criteria specified by CP
  - SET SRM DSPSLICE – amount of time a virtual CPU can use a physical CPU before being interrupted

Physical Processors…

- **Utilization of physical processors is controlled by:**
  - Number of VMDBKs in the dispatch list
- **Dispatch list is controlled by:**
  - SET DSPBUF command
  - Other SRM commands such as STORBUF & LDUBUF
- **Utilization reported by:**
  - CP INDICATE LOAD command – reports utilization based upon a 4 minute smoothed average
  - Performance Toolkit for VM and CP Monitor also report utilization without the smoothed average

Physical Processors…

- **Utilization is divided into:**
  - Time where CP is running on real processor (%CP)
    - Performing tasks not charged to a specific user (%SYS)
      - Scheduling, and other housekeeping
    - Performing tasks on behalf of virtual processors (%CP - %SYS)
      - Instruction simulation, etc.
  - Time where a virtual processor is running on a real processor (%EMU)
    - Actual user instructions

Virtual Processors

- **Virtual machines may be defined with up to 64 virtual processors**
  - Recommend not defining more virtual processors than physical
- **Virtual processors may be dedicated or shared**
  - Recommend not mixing shared and dedicated processors
  - Dedicated processors are rarely necessary
- **SHARE setting represents portion of system resources virtual processors should receive**
  - Absolute or Relative
  - Target minimum and maximum values
  - Maximum values can be either hard or soft limits
  - SHARE value divided between virtual processors

Monitoring CPU - Examples

1. CPU load and trans.

```
FCX100      CPU 7060  SER 10431  Interval 13:47:19 - 13:48:19    ZVMV4R40

CPU Load                                                 Vector Facility    Status or
PROC   %CPU   %CP %EMU   %WT %SYS   %SP %SIC %LOGLD    %VTOT %VEMU  REST   ded. User
P00     100     0  100     0    0     0   10    100      not installed     Master
P01     100     0  100     0    0     0    0    100      not installed     Alternate

Total SSCH/RSCH      2/s     Page rate         .0/s     Priv. instruct.    21/s
Virtual I/O rate     1/s     XSTORE paging     .0/s     Diagnose instr.    19/s
Total rel. SHARE    3200     Tot. abs SHARE      0%
```

Both CPUs at 100%

CP utilization is 0%

Virtual machine utilization at 100% - next find which virtual machine is using the CPU.

Monitoring CPU - Examples

21. User resource usage

FCX112 CPU 7060 SER 10431 Interval 14:08:19 - 14:09:19 ZVMV4R40

	<----- CPU Load ----->				<------ Virtual IO/s ------>						
		<-Seconds->		T/V							
Userid	%CPU	TCPU	VCPU	Ratio	Total	DASD	Avoid	Diag98	UR	Pg/s	User Status
>System<	14.2	8.538	8.531	1.0	.0	.0	.0	.0	.0	.0	---,---,---
WOJLINUX	99.7	59.79	59.75	1.0	.1	.1	.0	.0	.0	.0	ESA,CL0,DIS
WOJLIN02	99.5	59.72	59.68	1.0	.0	.0	.0	.0	.0	.0	ESA,CL0,DIS
VMRTM	.03	.018	.005	3.6	.0	.0	.0	.0	.0	.0	ESA,---,DOR
PERFSVM	.01	.005	.004	1.3	.1	.1	.0	.0	.0	.0	ESA,---,DOR
DISKACNT	0	0	0	...	0	0	0	0	0	0	ESA,---,DOR
EREP	0	0	0	...	0	0	0	0	0	0	ESA,---,DOR
FTPSERVE	.00	.000	.000	...	.0	.0	.0	.0	.0	.0	XC, ---,DOR
MAINT	.00	.000	.000	...	.0	.0	.0	.0	.0	.0	ESA,---,DOR
OPERATOR	0	0	0	...	0	0	0	0	0	0	ESA,---,DOR
OPERSYMP	0	0	0	...	0	0	0	0	0	0	ESA,---,DOR
TCPIP	.00	.001	.000	...	.1	.0	.0	.1	.0	.0	ESA,CL0,DIS
VMSERVR	0	0	0	...	0	0	0	0	0	0	ESA,---,DOR
VMSERVS	0	0	0	...	0	0	0	0	0	0	XC, ---,DOR
VMSERVU	0	0	0	...	0	0	0	0	0	0	XC, ---,DOR

From this report we can see that the Linux guests were using all of the CPU – the z/VM scheduler distributed processor resource fairly evenly across the virtual machines that needed it (each has SHARE Relative 100).

Monitoring CPU - Examples

3B. Proc. load & config
(PROCSUM)

```
FCX239        CPU 7060  SER 10431  Interval 13:30:19 - 14:25:19    ZVMV4R40

              <------- CPU --------> >Spin Lock Activity --------->
                    <--Ratio-->      >l -----> <--- Scheduler --->     <--SIE/sec-->
Interval       Pct          Cap-  On- >ge    Pct  Locks Average    Pct     In-  Inter-
End Time      Busy    T/V   ture line >ec   Spin   /sec    usec   Spin  struct    cept
>>Mean>>      99.9   1.00 .9979  2.0 >52   .000     .0   8.145   .000   424.3    21.8
14:04:19      99.9   1.00 .9979  2.0 469   .000     .0     ...      0   422.6    20.5
14:05:19      99.9   1.00 .9976  2.0 ...      0     .0     ...      0   423.5    21.3
14:06:19      99.9   1.00 .9981  2.0 ...      0     .0     ...      0   423.1    20.7
14:07:19      99.9   1.00 .9981  2.0 ...      0     .0     ...      0   422.0    19.9
14:08:19      99.9   1.00 .9981  2.0 ...      0     .0     ...      0   422.7    20.5
14:09:19      99.9   1.00 .9975  2.0 026   .000     .0     ...      0   422.0    20.2
14:10:19      99.9   1.00 .9981  2.0 483   .000     .0     ...      0   427.8    24.5
14:11:19      99.9   1.00 .9981  2.0 211   .000     .0   3.063   .000   424.6    21.9
```

T/V ratio shows low CP overhead (0%), further evidence of this is the ratio of SIE intercepts/second to SIE Instructions/second. Intercepts less than half the number of SIE Instructions is consistent with low CP overhead.

Monitoring CPU - Examples

1. CPU load and trans.

```
FCX100      CPU 7060  SER 10431  Interval 15:14:19 - 15:15:19    ZVMV4R40

CPU Load                                                Vector Facility    Status or
PROC   %CPU   %CP %EMU   %WT %SYS   %SP %SIC %LOGLD   %VTOT %VEMU  REST    ded. User
P00       4     2    2    96    1     0   99      4     not installed    Master
P01       3     2    1    97    1     0   99      3     not installed    Alternate

Total SSCH/RSCH    205/s     Page rate         .0/s     Priv. instruct.   1035/s
Virtual I/O rate   205/s     XSTORE paging     .0/s     Diagnose instr.     22/s
Total rel. SHARE     200     Tot. abs SHARE      0%
```

This example shows low CPU utilization, but fairly high I/O activity (virtual I/O rate is 205 per second). This results in a high number of privileged instructions per second and potential higher CP overhead (T/V ratio).

Monitoring CPU - Examples

21. User resource usage

FCX112 CPU 7060 SER 10431 Interval 15:19:19 - 15:20:19 ZVMV4R40

	<----- CPU Load	----->			<------ Virtual	IO/s	------>				
		<-Seconds->		T/V							
Userid	%CPU	TCPU	VCPU	Ratio	Total	DASD	Avoid	Diag98	UR	Pg/s	User Status
>System<	.36	.218	.115	1.9	14.6	14.4	.0	.1	.0	.0	---,---,---
WOJLINUX	4.43	2.655	1.386	1.9	202	202	.0	.0	.0	.0	ESA,CL0,DIS
WOJLIN02	.58	.345	.204	1.7	.2	.2	.0	.0	.0	.0	ESA,CL0,DIS
VMRTM	.03	.019	.005	3.8	.0	.0	.0	.0	.0	.0	ESA,---,DOR
MAINT	.02	.011	.003	3.7	.6	.0	.0	.0	.0	.0	ESA,---,DOR
PERFSVM	.02	.010	.008	1.3	.2	.2	.0	.0	.0	.0	ESA,---,DOR
TCPIP	.02	.013	.007	1.9	1.5	.0	.0	1.5	.0	.0	ESA,CL0,DIS
DISKACNT	0	0	0	...	0	0	0	0	0	0	ESA,---,DOR
EREP	0	0	0	...	0	0	0	0	0	0	ESA,---,DOR
FTPSERVE	.00	.000	.000	...	.0	.0	.0	.0	.0	.0	XC, ---,DOR
OPERATOR	0	0	0	...	0	0	0	0	0	0	ESA,---,DOR
OPERSYMP	0	0	0	...	0	0	0	0	0	0	ESA,---,DOR
VMSERVR	0	0	0	...	0	0	0	0	0	0	ESA,---,DOR
VMSERVS	0	0	0	...	0	0	0	0	0	0	XC, ---,DOR
VMSERVU	0	0	0	...	0	0	0	0	0	0	XC, ---,DOR

The virtual machine doing most of the I/O (WOJLINUX) shows high CP overhead (T/V ratio –> 2.655 / 1.386 = 1.92 92% overhead). Lots of instruction simulation!

Monitoring CPU - Examples

3B. Proc. load & config
(PROCSUM)

```
FCX239       CPU 7060   SER 10431   Interval 13:30:19 - 16:27:19    ZVMV4R40

           <------- CPU --------> >Spin Lock Activity --------->
                 <--Ratio-->      >l -----> <--- Scheduler ---> <--SIE/sec-->
Interval    Pct         Cap-  On- >ge   Pct  Locks Average   Pct    In-  Inter-
End Time   Busy   T/V  ture line >ec  Spin   /sec    usec  Spin struct   cept
>>Mean>>   37.0  1.01 .9924  2.0 >56  .000    1.8   2.324  .000  530.5  373.9
15:16:19    2.6  2.65 .7482  2.0 636  .001    9.6   1.895  .001   1525   1512
15:17:19    1.8  2.71 .7232  2.0 459  .001    7.4   1.599  .001   1042   1032
15:18:19    1.9  2.36 .7428  2.0 418  .001    7.0   1.604  .001  977.7  967.6
15:19:19    1.7  2.92 .6992  2.0 583  .001    7.1   1.952  .001  955.3  941.8
15:20:19    3.3  2.45 .7723  2.0 551  .001   10.0   1.695  .001   1870   1851
15:21:19    3.8  2.57 .7645  2.0 572  .001   12.5   1.701  .001   2170   2152
15:22:19    2.3  2.68 .7401  2.0 382  .001    6.9   1.573  .001   1334   1323
```

T/V ratio confirmed to be high. SIE Intercept/s to SIE Instruct/s ratio is high also, which is consistent with high I/O instead of high CPU.

Resources That Affect Performance

Processor

Memory

I/O

Central Storage

- **z/VM 5.3 supports up to 256GB of real memory as central storage for:**
  - Virtual machine page frames
  - z/VM nucleus
  - Control blocks and other CP storage requirements
  - z/VM virtual disk blocks
  - Minidisk Cache
- **Prior to z/VM 5.2:**
  - Page frames used for I/O buffers and other similar uses were required to reside below 2GB
    - Movement of page frames from above 2GB to below 2GB is recorded in monitor records

Expanded Storage

- **Used for:**
  - **High speed paging**
  - **Minidisk cache**
  - **Virtual disks**
- **Page frames must be moved from expanded to central storage before addressing the contents**
- **May be dedicated to a virtual machine**
  - **Linux can use expanded storage as a swap device**
  - **Dedicated XSTORE not available for use by CP**
- **Even with large amounts of central storage, it's best to define some expanded storage**
  - **Creates paging hierarchy – page to xstore first, then migrate old pages to DASD**
  - **Good starting point is to define expanded storage equal to about 25% of central storage size (1GB-4GB)**

Virtual Machine Storage

- **Defined on USER statement in USER DIRECT file (default & maximum)**
  - z/Architecture virtual machines may have greater than 2 GB storage
- **Virtual machine storage can be "locked" or "reserved"**
  - CP LOCK USERID username firstpagenum lastpagenum
    - Least flexible method to prevent paging
    - Requires knowledge of what specific pages should stay in central storage
  - CP SET RESERVED userid nnn
    - Best choice to reduce paging for a virtual machine
    - Specify number of pages to maintain in central storage for virtual machine, but not the exact pages
    - CP will allow nnn pages to remain resident for the specified virtual machine at all times
- **For Linux guests, keep virtual machine size as small as possible to help reduce allocation of I/O buffer and file system cache**

Storage Planning

- **Amount of central storage allocated does not need to be equal to sum of all logged on virtual machine sizes**
  - Central storage size is a function of virtual machine working set sizes, and page turn over rate
  - Hard to estimate for a completely new system
- **Amount of DASD paging space allocated needs to be greater than the sum of all logged on virtual machine sizes**
  - DASD paging area utilization should not exceed 50%
  - Higher utilizations reduce z/VM paging efficiency

Storage Planning…

- **Monitor effectiveness of minidisk cache:**
  - If little benefit and storage is constrained, set maximum cap for MDC or turn off minidisk cache so that pages can be used for virtual machine paging
  - CP SET MDCACHE STORAGE 0M 128M – set cap for central storage
  - CP SET MDCACHE XSTORE 0M 0M – eliminates use of expanded storage
  - CP SET MDCACHE SYSTEM OFF – turns off MDC
- **Reduce size of CP TRACE table if storage is constrained**
  - CP SET TRACEFRAMES MASTER 100

Storage Control

- **Scheduler can control storage and paging device utilization**
- **Virtual machines in dispatch list will have pages resident in central storage**
  - Access to dispatch list limited by scheduler when central storage is constrained
  - Scheduler estimates working set size for each virtual machine placed in dispatch list
  - Virtual machines exceeding storage thresholds placed in eligible list
    - No access to physical processors from eligible list
    - Scheduler moves virtual machines from eligible list to dispatch list as central storage becomes available
- **CP SET SRM STORBUF and CP SET SRM LDUBUF commands influence scheduler behavior**

Storage Control…

- **STORBUF**
  - Changes scheduler view of the amount of storage available for virtual machine pages
  - Typically need to modify default for Linux guest environment
  - Modification over commits central storage, so a robust paging subsystem is necessary
  - CP SET SRM STORBUF 300 250 200
- **LDUBUF**
  - Changes scheduler view of what the paging subsystem can handle
  - Applies to virtual machines classified as “loading users” (i.e. more than 5 page faults in a minor time slice)
  - Leave at default or set all 3 values to 100, depending on number and size of paging devices available
    - CP SET SRM LDUBUF 100 100 100

Storage Control…

- **QUICKDSP**
  - Virtual machine attribute
    - CP SET QUICKDSP userid ON
    - OPTION QUICKDSP in USER DIRECT
  - Instructs scheduler to NEVER place the virtual machine in an eligible list, even if central storage is constrained
  - Only use this attribute for critical service machines that other virtual machines depend on
    - TCP/IP virtual machine
    - Linux virtual machines acting as routers, database servers, etc.
  - Specifying this attribute for all virtual machines inhibits scheduler's ability to manage central storage

Paging Device Guidelines

- **Keep DASD page space utilization less than 50%**
- **Monitor blocks read per paging request**
  - Greater than 10 is good
- **Allocate page space over multiple volumes and multiple paths for best performance**
- **Do not mix paging areas with minidisk or other DASD usage**
  - CP uses never ending channel programs for paging devices
- **For storage subsystems, make sure cache is enabled**

Monitoring Memory - Examples

1. CPU load and trans.

```
FCX100      Data for 2003/12/02  Interval 16:45:37 - 16:46:37    Monitor Scan

CPU Load                                                   Vector Facility    Status or
PROC   %CPU  %CP %EMU  %WT %SYS  %SP %SIC %LOGLD           %VTOT %VEMU  REST  ded. User
P00     21    2   20   79    1    0   88     21               0     0    .0  ........
P01     28    1   27   72    0    0   82     28               0     0    .0  ........
P02     28    1   26   72    1    0   81     28               0     0    .0  ........

Total SSCH/RSCH      89/s       Page rate      140.7/s      Priv. instruct.   465/s
Virtual I/O rate      8/s       XSTORE paging  548.9/s      Diagnose instr.   192/s
Total rel. SHARE     8867       Tot. abs SHARE      0%

Queue Statistics:       Q0       Q1      Q2      Q3       User Status:
VMDBKs in queue          4        0       0      24       # of logged on users       42
VMDBKs loading           0        0       0       0       # of dialed users           0
Eligible VMDBKs                   0       0       0       # of active users          34
El. VMDBKs loading                0       0       0       # of in-queue users        28
Tot. WS (pages)     300011        0       0   1279k       % in-Q users in PGWAIT      0
Expansion factor                ...     ...     ...       % in-Q users in IOWAIT     97
85% elapsed time     1.472     .184   1.472   8.832       % elig. (resource wait)     0

Transactions        Q-Disp    trivial    non-trv          User Extremes:
Average users           .3         .0         .0          Max. CPU %    ZSITL001    36.1
Trans. per sec.         .3         .0         .0          Max. VECT %   ........   .....
Av. time (sec)       1.260       .053       .000          Max. IO/sec   MONWRITE     2.1
UP trans. time                   .053       .000          Max. PGS/s    LPWWPE02    60.4
MP trans. time                   .000       .000          Max. RESPG    LPWWPE02  120517
System ITR (trans. per sec. tot. CPU)         .5          Max. MDCIO    ........   .....
Emul. ITR (trans. per sec. emul. CPU)         .0          Max. XSTORE   ZDEVL014  101589
```

System has low CPU utilization, but is experiencing paging to both XSTORE and DASD (migration of pages out of XSTORE).

Monitoring Memory - Examples

2. Storage Utilization

```
FCX103       Data for 2003/12/02  Interval 16:43:37 - 16:44:37    Monitor Scan

Main storage utilization:
Total real storage              6'144MB
Total available                 6'144MB
Offline storage frames              0kB
SYSGEN storage size             6'144MB
CP resident nucleus             2'592kB
Shared storage                  2'184kB
FREE storage pages              7'532kB
FREE stor. subpools             1'944kB
Subpool stor. utilization          93%
Total DPA size                  2'021MB
Locked pages                   92'092kB
Trace table                     1'000kB
Pageable                        1'930MB
Storage utilization               109%
Tasks waiting for a frame           0
Tasks waiting for a page          0/s

V=R area:
Size defined                        0kB
FREE storage                        0kB
V=R recovery area in use          ...%
V=R user                      ........

Paging / spooling activity:
Page moves <2GB for trans.         2/s
Fast path page-in rate     17[illegible]/s
Long path page-in rate             5/s
Long path page-out rate          195/s
Page read rate                    17/s
Page write rate                    0/s
Page read blocking factor          6
Page write blocking factor       ...
Migrate-out blocking factor      ...
Paging SSCH rate                   5/s
SPOOL read rate                    0/s
SPOOL write rate                   0/s

XSTORE utilization:
Total available                 2'048MB
Att. to virt. machines              0kB
Size of CP partition            2'048MB
CP XSTORE utilization              99%
Low threshold for migr.         2'720kB
XSTORE allocation rate           200/s
Average age of XSTORE blks       4405s
Average age at migration          ...s

MDCACHE utilization:
Min. size in XSTORE                 0kB
Max. size in XSTORE             2'048MB
Ideal size in XSTORE          301'604kB
Act. size in XSTORE           166'664kB
Bias for XSTORE                   1.00
Min. size in main stor.             0kB
Max. size in main stor.         6'144MB
Ideal size in main stor.      559'916kB
Act. size in main stor.       409'832kB
Bias for main stor.               1.00
MDCACHE limit / user           32'144kB
Users with MDCACHE inserts          2
MDISK cache read rate              12/s
MDISK cache write rate          ...../s
MDISK cache read hit rate           7/s
MDISK cache read hit ratio         59%

VDISKs:
System limit (blocks)          Unlim.
User limit (blocks)            Unlim.
Main store page frames              0
Expanded stor. pages                0
Pages on DASD                      36
```

Check to see if there is a below 2GB constraint.

Monitoring Memory - Examples

3A. Paging activity

FCX143 Data for 2003/12/02 Interval 16:30:37 - 17:02:37 Monitor Scan

	<------------	Expanded	Storage	------------>			<---------	Real	Storage	---
Interval End Time	Paging Blocks	PGIN /s	Fast-Path %	PGOUT /s	Total /s	Est. Life sec	Page Migr /s	DPA Pgable Frames	Non-pgable Pages	Resid Shared Pages	Mean Avail List
>>Mean>>	484804	93.2	97.4	114.5	207.8	4233	20.5	494092	23353	497	485
16:20:37	483790	85.4	99.7	112.6	198.0	4298	.0	495354	22091	528	2029
16:21:37	483790	39.8	99.9	47.0	86.8	10301	.0	495398	22047	530	1725
16:22:37	483790	63.5	99.9	71.5	135.0	6768	.0	495224	22221	529	872
16:23:37	483790	34.5	99.7	42.4	76.9	11419	.0	495865	21580	529	806
16:24:37	483790	31.9	99.9	42.8	74.7	11304	.0	495354	22091	529	760
16:25:37	483790	40.7	99.9	42.4	83.0	11419	2.5	493610	23835	529	587
16:26:37	483790	36.3	99.8	43.8	80.2	11037	6.1	494505	22940	528	530
16:27:37	483790	487.8	99.3	483.9	971.7	1000	1.0	493371	24074	527	360
16:28:37	483790	123.2	85.3	148.6	271.8	3256	28.2	492728	24717	529	328
16:29:37	482112	29.8	97.3	10.2	40.0	47189	23.0	492809	24636	528	855
16:30:37	479523	98.3	98.4	71.8	170.1	6680	12.4	493690	23755	525	1647
16:31:37	487253	63.7	96.0	86.4	150.1	5637	.0	492884	24561	528	1249
16:32:37	487249	413.8	99.2	414.6	828.4	1175	.0	494534	22911	547	274
16:33:37	487249	50.6	94.7	88.3	138.9	5518	.0	494619	22826	527	306
16:34:37	487249	78.3	85.3	135.0	213.3	3608	.0	493812	23633	550	251
16:35:37	487249	48.9	98.1	61.6	110.4	7916	.0	494649	22796	530	285
16:36:37	484434	32.6	98.0	7.0	39.6	69536	32.1	493696	23749	551	999
16:37:37	484434	61.5	96.7	138.6	200.1	3496	51.0	495274	22171	530	660
16:38:37	484434	108.3	100.0	171.9	280.2	2819	68.1	495490	21955	548	604
16:39:37	484434	43.2	99.8	51.4	94.6	9425	9.3	493478	23967	529	575
16:40:37	484434	55.8	99.7	65.6	121.4	7383	6.2	493478	23967	531	258
16:41:37	484434	137.2	99.1	223.7	360.9	2166	91.4	493489	23956	534	270

Total paging into and out of XSTORE by time, shows spikes up to 971/s.

Monitoring Memory - Examples

3A. Paging activity

FCX143 Data for 2003/12/02 Interval 16:30:37 - 17:02:37 Monitor Scan

		Real Storage				Paging to DASD						
	>	Non-	Resid	Mean	Est.					<-Single Reads-->		
Interval	>	pgable	Shared	Avail	Page	Reads	Write	Total	Shrd	Guest	Systm	Total
End Time	>	Pages	Pages	List	Life	/s	/s	/s	/s	/s	/s	/s
>>Mean>>	>	23353	497	485	3681	26.6	19.7	46.3	4.4	1.0	.0	1.0
16:20:37	4	22091	528	2029	4401	35.9	.0	35.9	3.6	1.7	.0	1.7
16:21:37	8	22047	530	1725	10547	9.6	.0	9.6	3.4	.3	.0	.3
16:22:37	4	22221	529	872	6927	21.3	.0	21.3	3.5	.5	.0	.5
16:23:37	5	21580	529	806	11704	7.4	.0	7.4	3.5	.3	.0	.3
16:24:37	4	22091	529	760	11573	7.8	.0	7.8	3.4	.3	.0	.3
16:25:37	0	23835	529	587	11071	4.2	2.2	6.4	3.3	.2	.0	.2
16:26:37	5	22940	528	530	9903	4.6	6.1	10.7	3.4	.2	.0	.2
16:27:37	1	24074	527	360	1016	5.4	1.3	6.7	4.8	.2	.0	.2
16:28:37	8	24717	529	328	2785	26.6	28.3	54.9	6.7	2.4	.4	2.8
16:29:37	9	24636	528	855	14836	8.3	23.0	31.3	3.6	.2	.0	.2
16:30:37	0	23755	525	1647	5863	16.6	12.4	29.1	4.8	.7	.0	.7
16:31:37	4	24561	528	1249	5702	25.0	.0	25.0	5.2	1.4	.0	1.4
16:32:37	4	22911	547	274	1192	10.6	.1	10.6	5.8	.4	.0	.4
16:33:37	9	22826	527	306	5601	39.0	.0	39.0	5.1	1.8	.0	1.8
16:34:37	2	23633	550	251	3656	52.4	.0	52.4	6.0	.8	.0	.8
16:35:37	9	22796	530	285	8036	9.1	.0	9.1	3.5	.4	.0	.4
16:36:37	6	23749	551	999	13544	5.2	29.5	34.7	4.1	.3	.0	.3
16:37:37	4	22171	530	660	2614	89.9	50.9	140.8	4.9	6.0	.0	6.0
16:38:37	0	21955	548	604	2067	67.4	67.8	135.1	3.6	9.5	.0	9.5
16:39:37	8	23967	529	575	8129	10.2	9.3	19.5	3.0	.2	.0	.2
16:40:37	8	23967	531	258	6874	8.6	6.2	14.8	3.4	.4	.0	.4
16:41:37	9	23956	534	270	1567	86.9	91.1	178.1	3.8	.4	.0	.4
16:42:37	7	23598	547	369	2766	40.2	49.2	89.4	3.8	.3	.0	.3
16:43:37	7	23008	532	1797	8559	19.0	35.2	54.3	5.0	.6	.0	.6
16:44:37	2	23023	546	254	2537	17.3	.0	17.3	4.1	.7	.0	.7
16:45:37	6	25189	551	269	1165	29.1	.0	29.1	6.3	1.5	.0	1.5
16:46:37	4	22051	531	321	1305	73.8	67.0	140.7	6.3	1.3	.0	1.3

Total paging to DASD per second by time – shows spikes also. Mean Avail List is quite small at times (274*4096~1MB). Storage is constrained.

Monitoring Memory - Examples

14. CP owned disks

```
FCX109      Data for 2003/12/02  Interval 16:45:37 - 16:46:37    Monitor Scan

Page / SPOOL Allocation Summary
PAGE slots available       4831200         SPOOL slots available       632880
PAGE slot utilization          59%         SPOOL slot utilization         15%
T-Disk cylinders avail.    .......         DUMP slots available             0
T-Disk space utilization      ...%         DUMP slot utilization          ..%

____ .          .                    >    .     .     .     .     .     .
< Device Descr. ->                   > User             Serv MLOAD Block %Used
             Volume Area   Area        Used >Inter Queue  Time  Resp Page  for
Addr Devtyp Serial Type   Extent         % >feres Lngth /Page  Time Size Alloc
0638 3390-3 PAG003 PAGE     1-3338      78      1     0   1.7   1.7    4    52
063A 3390-3 PAG005 PAGE     1-3338      79      1     0    .9    .9    3    62
084A 3390-3 SPL001 SPOOL    1-3338      12      0     0   1.2   1.2  ...   ...
085D 3390-3 PAG001 PAGE     1-3338      83      1     0    .6    .6    4    90
085E 3390-3 430RES SPOOL   79- 256      73    ...   ...   ...   ...  ...   ...
                   PAGE   257- 390     100      1     0   5.1   5.1  ...   ...
0A00 3390-9 PAG006 PAGE     802880      24      1     0   1.0   1.0    9    90
0C10 3390-3 PAG004 PAGE     1-3338      79      1     0    .9    .9    4    71
0C29 3390-3 PAG002 PAGE     1-3338      77      1     0    .7    .7    5    57
```

Since the system is paging, check paging efficiency. Block page size is small – the result of over allocated paging extents. A block page size greater than 10 is good. Add paging extents to correct this problem.

Monitoring Memory - Examples

22. User paging load

FCX113 Data for 2003/12/02 Interval 16:45:37 - 16:46:37 Monitor Scan

	Data	<---------- Paging	Activity/s	----------	---------->				<--------------			Numbe
	Spaces	<Page	Rate>	Page	<--Page	Migration	-->					
Userid	Owned	Reads	Write	Steals	>2GB>	X>MS	MS>X	X>DS	WSS	Lockd	Resrvd	
>System<	.0	1.8	1.6	7.4	.1	5.6	7.4	1.8	37745	423	0	
LPWWPE02	0	60.4	.0	.0	.2	2.0	4.1	.0	129591	2546	0	
LPWWPE04	0	.2	.0	.0	.2	.1	3.8	.0	126167	2883	0	
ZDEVL007	0	.6	.0	.0	.0	1.0	66.6	.0	113718	9	0	
LUWWPE01	0	.0	.0	.0	.0	2.8	19.4	.0	111966	9	0	
LPWWPE03	0	.0	.0	.0	.0	.0	.6	.0	111367	2274	0	
ZSITL001	0	3.1	14.7	.0	.0	.3	34.3	14.8	102406	9	0	
ZDEVL004	0	.0	.0	.0	.0	.0	7.5	.0	99803	1181	0	1
LDWWAI01	0	.0	50.1	.0	.4	3.3	27.0	50.1	98736	9	0	
ZDEVL017	0	.0	.0	.0	.0	5.8	.1	.0	95041	9	0	
LUWWPE02	0	8.0	.0	.0	.0	138	31.0	.0	94253	9	0	
ZDEVL014	0	.0	.0	.0	.0	2.7	18.4	.0	89763	9	0	
LPWWPE01	0	1.0	.0	.0	.2	.2	.2	.0	88841	3058	0	
ZDEVL006	0	.0	.0	.0	.0	.1	.0	.0	83700	12	0	
ZSITL002	0	.0	.0	.0	.0	.0	.0	.0	56792	9	0	
LNXSYS01	0	.0	.0	.0	.0	.0	.0	.0	29374	9	0	
ZSYSL002	0	.2	.0	.0	.0	.1	.0	.0	24078	12	0	
ZDEVL015	0	.0	.0	.0	.0	.0	.0	.0	21886	9	0	
ZDEVL008	0	.0	.0	.0	.0	.0	.0	.0	20749	9	0	
LIWWAI01	0	.1	2.2	.0	.0	.6	.0	8.6	19961	9	0	
LPWWAE01	0	.0	.0	.0	.0	.5	.0	.0	19572	659	0	
ZPRDL006	0	.1	.0	.0	.4	2.6	.9	.0	12500	9	0	
ZDEVL016	0	.0	.0	.0	.0	.5	7.0	.0	10213	9	0	
LPSLDI01	0	.0	.0	.0	.0	1.8	8.1	.0	4703	1610	0	
LDSLDI01	0	.0	.0	.0	.0	.0	.0	.0	4207	9	0	
LUSLDI01	0	.0	.0	.0	.0	8.2	9.3	.0	4129	9	0	

Since system is paging, check to see which guests are paging.

Monitoring Memory - Examples

23. User wait states

FCX114 Data for 2003/12/02 Interval 16:45:37 - 16:46:37 Monitor Scan

											<-SVM	and->			
Userid	%ACT	%RUN	%CPU	%LDG	%PGW	%IOW	%SIM	%TIW	%CFW	%TI	%EL	%DM	%IOA	%PGA
>System<	67	3	0	0	0	0	0	0	0	0	0	0	96	0
LPWWPE02	100	0	0	0	12	0	0	0	0	0	0	0	88	0
DATAMOVE	0	...	...	...	...	...	...	...	...	...	...	...	...	...
DIRMAINT	0	...	...	...	...	...	...	...	...	...	...	...	...	...
EREP	0	...	...	...	...	...	...	...	...	...	...	...	...	...
FTPSERVE	0	...	...	...	...	...	...	...	...	...	...	...	...	...
LDSLDI01	100	0	0	0	0	0	0	0	0	0	0	0	100	0
LDWWAI01	100	0	0	0	0	0	0	0	0	0	0	0	100	0
LIWWAI01	100	0	0	0	0	0	0	0	0	0	0	0	100	0
LNXSYS01	100	0	0	0	0	0	0	0	0	0	0	0	100	0
LPSLDI01	100	0	0	0	0	0	0	0	0	0	0	0	100	0
LPSLDI02	100	0	0	0	0	0	0	0	0	0	0	0	100	0
LPWWAE01	100	0	0	0	0	0	0	0	0	0	0	0	100	0
LPWWPE01	93	4	0	0	0	0	0	0	0	0	0	0	96	0
LPWWPE03	100	0	0	0	0	0	0	0	0	0	0	0	100	0
LPWWPE04	100	2	0	0	0	0	0	0	0	0	0	0	98	0
LUSLDI01	100	0	2	0	0	0	0	0	0	0	0	0	98	0
LUWWPE01	100	0	0	0	0	0	0	0	0	0	0	0	100	0
LUWWPE02	100	7	0	0	0	0	0	0	0	0	0	0	93	0
MAINT	0	...	...	...	...	...	...	...	...	...	...	...	...	...
MONWRITE	7	0	0	0	0	50	0	50	0	0	0	0	0	0
OPERATOR	0	...	...	...	...	...	...	...	...	...	...	...	...	...
OPERSYMP	0	...	...	...	...	...	...	...	...	...	...	...	...	...
TCPIP	50	3	0	0	0	0	0	0	0	0	0	0	97	0

Since guests are paging and the block page size is small, check to see if guests are frequently in page wait.

IBM Advanced Technical Support

Resources That Affect Performance

Processor

Memory

I/O

Physical Devices

- **Physical devices can be exploited by CP or supported for guest use, for example:**
  - CP supports allocation of system areas on 3390 DASD
  - CP supports OSA devices for guest use only
- **Devices supported for guest use only must be dedicated to a virtual machine**
  - Devices exploited by CP may also be dedicated or attached to a service virtual machine
  - Dedicated devices may only be used by the virtual machine controlling the device

Virtual Devices

- **Virtualized devices**
  - CP manages the underlying physical device
  - CP provides appearance that each virtual machine has the device
  - The underlying physical device must be present on the system
  - E.g. minidisks, crypto cards
- **Simulated devices**
  - Complete representation of a physical device
  - No physical device present
  - E.g. 2540 card punch, Guest LAN, virtual disk
- **Virtual devices are defined in the virtual machine directory entry, or dynamically created using CP commands**

Controlling I/O

- **CP does not manage general I/O as a tunable resource**
- **Limited tuning for real devices provided by:**
  - CP SET THROTTLE command
  - Slow I/O's per second from guest operating system to a particular real device
- **Indirect control through SHARE setting, or other scheduler controls such as STORBUF, DSPBUF, LDUBUF**
  - Virtual machines not in the dispatch list, cannot execute I/O instructions
- **I/O Priority Queuing**
  - CP can effect queue placement for DASD devices
  - HW can effect priority in channel usage
- **Queued I/O Assist**
  - Interpretive execution assist for QDIO devices

Additional I/O Considerations

- **Dedicated I/O is not eligible for Minidisk Cache (MDC)**
- **MDC read performance is as good as virtual disk performance**
- **Both virtual disks and MDC require sufficient real memory**

Monitoring I/O

- **Privileged CP commands**
  - INDICATE I/O – shows users in I/O wait state, and real device number to which most recent virtual I/O operation was mapped
  - INDICATE USER userid EXP – displays total number of virtual I/O operations started since logon, repeated displays provide a rough idea of I/O activity for particular virtual machine
- **CP Monitor**
  - Detailed I/O information available within monitor for each real device, including seeks information
  - Total I/O information for individual virtual machines
- **Performance Toolkit for VM**
  - Detailed device I/O information and virtual machine I/O information

Monitoring I/O - Examples

1. CPU load and trans.

```
FCX100      CPU 7060  SER 10431  Interval 11:52:19 - 11:53:19    ZVMV4R40

CPU Load                                                 Vector Facility    Status or
PROC   %CPU   %CP %EMU   %WT %SYS   %SP %SIC %LOGLD      %VTOT %VEMU  REST   ded. User
P00     13     4    9    87    1     0   98     13        not installed    Master
P01     12     4    9    88    1     0   98     13        not installed    Alternate

Total SSCH/RSCH     809/s     Page rate         .0/s       Priv. instruct.   2461/s
Virtual I/O rate    808/s     XSTORE paging     .0/s       Diagnose instr.    25/s
Total rel. SHARE     3200     Tot. abs SHARE     0%

Queue Statistics:      Q0      Q1      Q2      Q3      User Status:
VMDBKs in queue         5       0       0       0      # of logged on users       14
VMDBKs loading          0       0       0       0      # of dialled users          0
Eligible VMDBKs                 0       0       0      # of active users           7
El. VMDBKs loading              0       0       0      # of in-queue users         5
Tot. WS (pages)     12493       0       0       0      % in-Q users in PGWAIT      0
Expansion factor                1       1       1      % in-Q users in IOWAIT     35
85% elapsed time    6.864    .858   6.864   41.18      % elig. (resource wait)     0

Transactions        Q-Disp    trivial    non-trv       User Extremes:
Average users           .4         .0         .0       Max. CPU %    WOJLINUX    22.6
Trans. per sec.         .4         .2         .0       Max. VECT %   ........   .....
Av. time (sec)       1.261       .002       .002       Max. IO/sec   WOJLINUX     806
```

System shows I/O activity (virtual I/O almost equal to SSCH/RSCH rate). Also note that the privileged instruction rate is high. This is consistent with virtual I/O.

Monitoring I/O - Examples

13. I/O device load

```
FCX108      CPU 7060  SER 10431  Interval 11:52:19 - 11:53:19    ZVMV4R40
.    .                     .           .       .    .    .    .    .    .    .
<-- Device Descr. -->  Mdisk Pa- <-Rate/s-> <------- Time (msec) -------> Req.
Addr Type   Label/ID   Links ths  I/O Avoid Pend Disc Conn Serv Resp CUWt Qued
>> All DASD <<          ....       .3    .0   .0  8.4  2.6 11.0 11.0   .0  .00
A00B 3390-3 WOJ001         2   1  8.0    .0   .0  9.7  3.0 12.7 12.7   .0  .00
A00C 3390-3 WOJ002         2   1   .2    .0   .0  1.0   .1  1.1  1.1   .0  .00
A005 3390-3 440W02 CP      8   1   .2    .0   .0   .9   .1  1.0  1.0   .0  .00
0124 3380   OS39H7         0   1   .0    .0  ...  ...  ...  ...  ...  ...  .00
0125 3380   HFSUS1         0   1   .0    .0  ...  ...  ...  ...  ...  ...  .00
0A80 3390-3 OS39R7         0   1   .0    .0  ...  ...  ...  ...  ...  ...  .00
0A82 3390-2 OS3R7A         0   1   .0    .0  ...  ...  ...  ...  ...  ...  .00
0A85 3390-1 M3KPLX         0   1   .0    .0  ...  ...  ...  ...  ...  ...  .00
4340 3390-3 240RES         0   2   .0    .0   .3   .0   .3   .6   .6   .0  .00
4341 3390-3 LNX1  X        0   2   .0    .0   .1   .0   .3   .4   .4   .0  .00
4342 3390-3 SUS8           0   2   .0    .0   .3   .1   .3   .7   .7   .0  .00
4343 3390-3 VMLIN2         0   2   .0    .0   .1   .0   .3   .4   .4   .0  .00
4344 3390-3 DEBIAN         0   2   .0    .0   .1   .0   .3   .4   .4   .0  .00
A003 3390-3 440RES CP    132   1   .0    .0   .1   .1   .0   .2   .2   .0  .00
A004 3390-3 440W01 CP     47   1   .0    .0   .0   .0   .0   .0   .0   .0  .00
A009 3390-3 PAG001 CP      0   1   .0    .0   .1   .0   .0   .1   .1   .0  .00
A00A 3390-3 BASRES         0   1   .0    .0   .0   .1   .0   .1   .1   .0  .00
A00D 3390-3 440U01         3   1   .0    .0   .2   .1   .0   .3   .3   .0  .00
A00E 3390-3 BASW01         0   1   .0    .0   .0   .0   .0   .0   .0   .0  .00
A00F 3390-3 M2K353         0   1   .0    .0   .1   .1   .0   .2   .2   .0  .00
A010 3390-3 BASW02         0   1   .0    .0   .0   .0   .0   .0   .0   .0  .00
A011 3390-3 430RES         0   1   .0    .0   .1   .0   .0   .1   .1   .0  .00
A012 3390-3 430W01         0   1   .0    .0   .1   .0   .0   .1   .1   .0  .00
```

Since I/O is occurring, look to see which devices are affected and whether response time is poor. Note the device address and number of paths for next screen.

Monitoring I/O - Examples

19. I/O configuration

```
FCX131      CPU 7060  SER 10431       Status   11:53:19        ZVMV4R40

<----  Ranges  ---->   Device       <- Channel Path Ids   ->   Control
Device-No  Subch.-ID   Type          1  2  3  4  5  6  7  8    Unit      Status
000E-000F  01B8-01B9   1403         FC  .  .  .  .  .  .  .    2821-01   Online
0124-0125  01BC-01BD   3380   (E)   FC  .  .  .  .  .  .  .    3880-23   Online
0700-0702  01C6-01C8   3270-2       FC  .  .  .  .  .  .  .    3274-1D   Online
0900-0901  01CB-01CC   3270-2       FC  .  .  .  .  .  .  .    3274-1D   Online
0A80       01CD        3390-3 (E)   FC  .  .  .  .  .  .  .    3990      Online
0A82       01CF        3390-2 (E)   FC  .  .  .  .  .  .  .    3990      Online
0A85       01D2        3390-1 (E)   FC  .  .  .  .  .  .  .    3990      Online
0C20-0C23  01D3-01D6   3490         05 07  .  .  .  .  .  .    3490-20   Online
0C24-0C27  01D7-01DA   3490         05 07  .  .  .  .  .  .    3490-20   Offline
0C28-0C2B  01DB-01DE   3490         05 07  .  .  .  .  .  .    3490-20   Online
0C2C-0C2F  01DF-01E2   3490         05 07  .  .  .  .  .  .    3490-20   Offline
0E60       01E7        3270         0E  .  .  .  .  .  .  .    3174-1D   Online
0E61       01E8        3422         0E  .  .  .  .  .  .  .    3422-A1   Online
0F40-0F41  0207-0208   CTCA         FC  .  .  .  .  .  .  .    3088      Online
1200-1207  020B-0212   CTCA         0B  .  .  .  .  .  .  .    3088      Online
1300-1307  0213-021A   CTCA         0B  .  .  .  .  .  .  .    3088      Online
1400-1407  0223-022A   CTCA         0B  .  .  .  .  .  .  .    3088      Online
1500-1507  023B-0242   CTCA         0B  .  .  .  .  .  .  .    3088      Online
4340-4344  02B3-02B7   3390-3 (E)   04 06  .  .  .  .  .  .    9393 RVA  Online
A003-A005  02C3-02C5   3390-3 (E)   FD  .  .  .  .  .  .  .    3990      Online
A009-A017  02C9-02D7   3390-3 (E)   FD  .  .  .  .  .  .  .    3990      Online
A018       02D8        3390-1 (E)   FD  .  .  .  .  .  .  .    3990      Online
```

Look to see what chpids are associated with device A00B.

Monitoring I/O - Examples

11. Channel load

FCX107 CPU 7060 SER 10431 Interval 11:29:19 - 11:53:19 ZVMV4R40

CHPID (Hex)	Chan-Group Descr	Chan-Group Qual	<%Busy> Cur	<%Busy> Ave	<----- Channel %Busy Distribution 11:29:19-11:5 0-10	11-20	21-30	31-40	41-50	51-60	61-70	71-80
04	ESCON	00	0	0	100	0	0	0	0	0	0	0
05	ESCON	00	0	0	100	0	0	0	0	0	0	0
06	ESCON	00	0	0	100	0	0	0	0	0	0	0
07	ESCON	00	0	0	100	0	0	0	0	0	0	0
08	ESCON	00	0	0	100	0	0	0	0	0	0	0
09	ESCON	00	0	0	100	0	0	0	0	0	0	0
0A	ESCON	00	0	0	100	0	0	0	0	0	0	0
0B	ESCON	00	0	0	100	0	0	0	0	0	0	0
0C	ESCON	00	0	0	100	0	0	0	0	0	0	0
0D	ESCON	00	0	0	100	0	0	0	0	0	0	0
0E	ESCON	00	0	0	100	0	0	0	0	0	0	0
0F	ESCON	00	0	0	100	0	0	0	0	0	0	0
10	ESCON	00	0	0	100	0	0	0	0	0	0	0
40	ESCON	00	..	..	..	..	..	..	..	..	..	..
FC	FICON	00	..	..	..	..	..	..	..	..	..	..
FD	ESCON	00	..	..	..	..	..	..	..	..	..	..

Chpid FD shows all '…', which means utilization data is not available (internal channels on 7060).

Monitoring I/O - Examples

13. I/O device load (select device)

```
FCX110      CPU 7060  SER 10431  Interval 11:52:03 - 11:54:30    ZVMV4R40

Detailed Analysis for Device A00B ( SYSTEM )
Device type :  3390-3      Function pend.:     .0ms      Device busy    :     9%
VOLSER      :  WOJ001      Disconnected  :    9.5ms      I/O contention:     0%
Nr. of LINKs:       2      Connected     :    3.0ms      Reserved       :     0%
Last SEEK   :    1224      Service time  :   12.5ms      SENSE SSCH     :     0
SSCH rate/s :     7.9      Response time :   12.5ms      Recovery SSCH  :     0
Avoided/s   :      .0      CU queue time :     .0ms      Throttle del/s:   ...
Status: MDCACHE USED

Path(s) to device A00B:    FD
Channel path status   :    ON

Device                 Overall CU-Cache Performance           Split
DIR ADDR VOLSER        IO/S %READ %RDHIT %WRHIT ICL/S BYP/S    IO/S %READ %RDHIT

  MDISK Extent         Userid    Addr IO/s VSEEK Status     LINK   VIO/s %MDC MDIO/s
+---------------------------------------------------------------------------------+
C      1 -   50        WOJLINUX 0191   .0      0 WR            1      .0  ...     .0 C
C     51 - 3338        WOJLINUX 0200  7.9   1173 WR            1     7.9    0    7.9 C
+---------------------------------------------------------------------------------+
```

Detail device screen shows virtual machine doing I/O to a minidisk (0200) on physical device A00B.

Monitoring I/O - Examples

23. User wait states

FCX114 CPU 7060 SER 10431 Interval 11:53:19 - 11:54:19 ZVMV4R40

										<-SVM	and->			
Userid	%ACT	%RUN	%CPU	%LDG	%PGW	%IOW	%SIM	%TIW	%CFW	%TI	%EL	%DM	%IOA	%PGA
>System<	24	5	2	0	0	0	0	45	0	0	0	18	28	0
PERFSVM	63	0	0	0	0	3	0	0	0	0	0	97	0	0
DISKACNT	0	...	...	...	...	...	...	...	...	...	...	...	...	...
EREP	0	...	...	...	...	...	...	...	...	...	...	...	...	...
FTPSERVE	0	...	...	...	...	...	...	...	...	...	...	...	...	...
MAINT	10	0	0	0	0	0	0	100	0	0	0	0	0	0
OPERATOR	0	...	...	...	...	...	...	...	...	...	...	...	...	...
OPERSYMP	0	...	...	...	...	...	...	...	...	...	...	...	...	...
TCPIP	67	0	0	0	0	0	0	0	0	0	0	0	100	0
VMRTM	0	...	...	...	...	...	...	...	...	...	...	...	...	...
VMSERVR	0	...	...	...	...	...	...	...	...	...	...	...	...	...
VMSERVS	0	...	...	...	...	...	...	...	...	...	...	...	...	...
VMSERVU	0	...	...	...	...	...	...	...	...	...	...	...	...	...
WOJLINUX	100	18	3	0	0	0	0	47	0	0	0	0	30	0
WOJLIN02	100	0	5	0	0	0	0	95	0	0	0	0	0	0

Check to see if the virtual machine performing the I/O, is consistently waiting on I/O.

Summary

Summary

- **Primary resources that can affect z/VM guest performance are:**
  - Processor
  - Memory
  - Paging I/O
- **z/VM controls for managed resources include:**
  - SET SRM STORBUF – memory
  - SET SRM LDUBUF – paging I/O
  - SET SHARE – processor
- **Monitoring z/VM performance**
  - CP commands – watch out for smoothed averages
  - Performance Toolkit for VM
  - Other performance software
  - Start monitoring now!