



IBM Systems Technology Group

z/VM Performance Update MVMUA – April 24, 2007

Bill Bitner
VM Performance Evaluation
bitnerb@us.ibm.com

HA
VP

4/25/2007

© 2007 IBM Corporation

Trademarks

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: AS/400, DBE, e-business logo, ESCO, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/30, VM/ESA, VSE/ESA, Websphere, xSeries, z/OS, zSeries, z/VM

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 Linus 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.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

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

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

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

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

Introduction

- **Focus on z/VM 5.2.0**
 - Line Items that have an impact
 - See Performance Report on web for details
 - <http://www.vm.ibm.com/perf/reports/>
- **Additional Post z/VM 5.2.0 News**
- **A few comments on z/VM 5.3.0 Expectations**

z/VM 5.2.0

- **GA December 16, 2005**
- **Key Performance Line Items:**
 - Large Real Memory (64-bit)
 - Emulated FBA on SCSI Improvements
 - V=V QDIO Pass-through Stage II
 - Diagnose x'44' Processing
 - Directory Performance Improvements
- **Performance Toolkit for VM Enhancements**

Large Real Memory (64-Bit Enhancements)

- **Brief history**
- **How it works in z/VM 5.1.0**
 - What to look at to determine problems?
 - What was done to relieve problems?
- **How it works in z/VM 5.2.0**
 - What changed?
 - What did not change?

Large Real Memory – z/VM 5.1.0

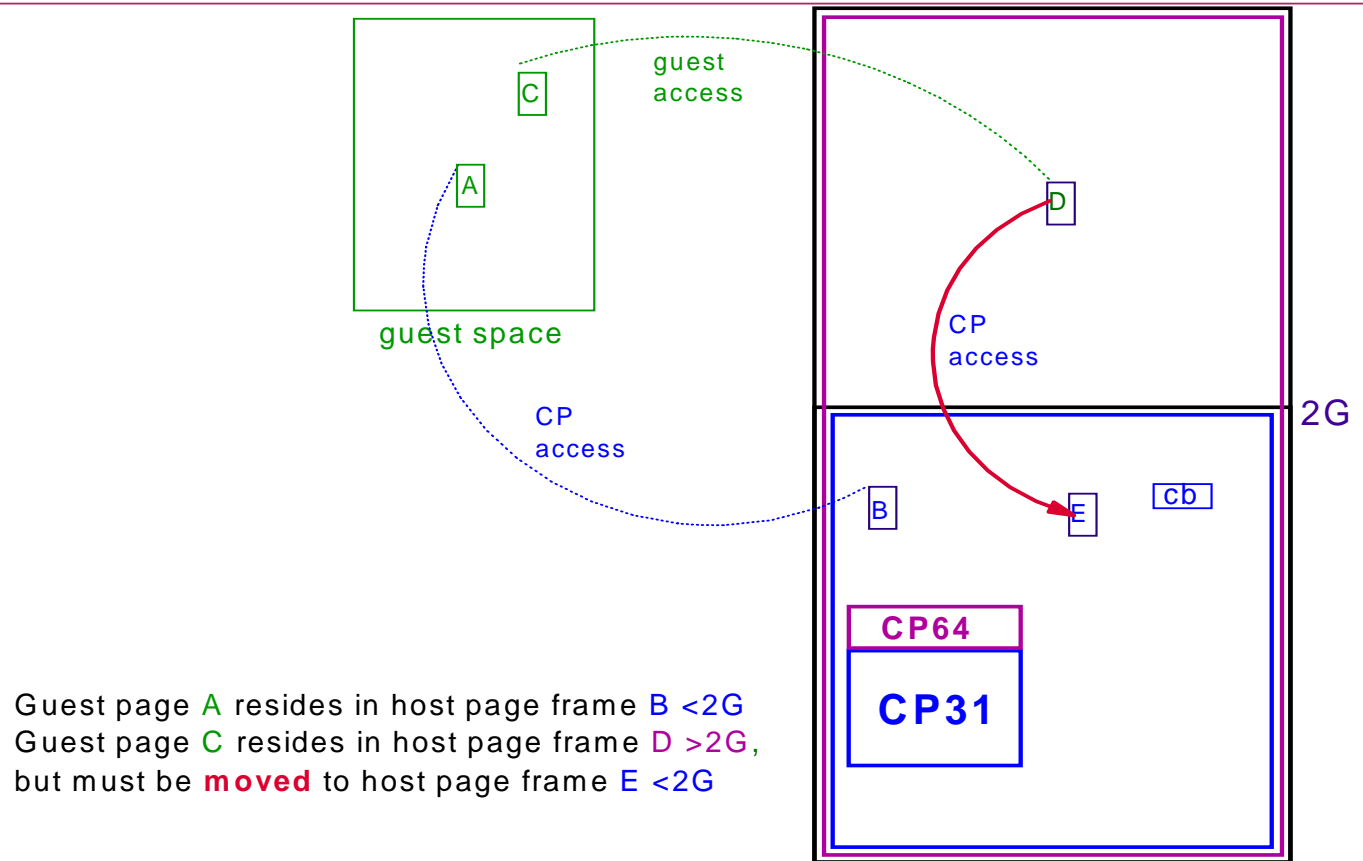
- While z/VM 5.1.0 supports 64-bit addressing and virtual 64-bit addressing, most of the Control Program (CP) still uses 31-bit addressing
 - Guest pages manipulated or used by CP need to be moved to below 2GB real (e.g. guest pages containing I/O channel programs and data buffers)
 - Can occur whether guest is 31-bit or 64-bit
- Permanent structures below 2GB:
 - CP Nucleus
 - Frame Table- maps all of real memory
 - CP control blocks (RDEVs, VDEVs, etc.)
 - Segment Tables
 - Data structures for each guest QDIO or FCP dedicated device
- Temporarily reside below 2GB:
 - Guest pages for channel programs and data buffers
 - Guest pages associated with Guest LAN and Virtual Switch
 - Page Management Blocks
 - Some exceptions are permanently locked (e.g. vdisk)
 - Others

Overview

z/VM 3.1.0 - 5.1.0 CP Storage Usage

IBM

- The Problem:
- Guest pages can reside >2G if referenced by the guest
- Guest pages must reside <2G if referenced by CP
- If a guest page resides >2G, and CP needs to reference it, the page must be moved from >2G (a relatively unconstrained resource) to <2G (a constrained resource)
- CP too large to convert all references to guest pages to 64-bit addressing
- CP cannot effectively exploit large amounts of real storage above 2G, and is constrained below 2G



Signs of Contention for <2GB Memory in z/VM

5.1.0

▪ **Basic CP Commands**

- INDICATE LOAD shows high paging rates **and** Q FRAMES shows large number of pages available above 2GB

▪ **Performance Toolkit**

- FCX100 CPU
 - “%SYS” – may be high if lots of CPU time spent in available list replenishment
- FCX103 Storage
 - “Page moves <2GB for trans.” in high 100s/second
- FCX114 Ustat
 - Significant users in “%IOW” and “%CPU” (when CPU available)
- FCX113 Upage
 - User with 100s pages/sec moved below 2GB or multiple users with large numbers “>2GB>”
 - Interesting to see who has pages below 2GB

FCX100 Data for 2005/09/01 Interval 01:02:39 - 01:03:40 Monitor Scan

CPU Load										Vector Facility			Status or
PROC	%CPU	%CP	%EMU	%WT	%SYS	%SP	%SIC	%LOGLD	%VTOT	%VEMU	REST	ded. User	
P00	48	24	24	52	13	6	100	57	0	0	.0	
P01	49	25	24	51	14	5	100	58	0	0	.0	
P02	49	23	25	51	13	4	100	57	0	0	.0	
P03	50	27	22	50	18	5	100	61	0	0	.0	
P04	48	21	27	52	11	4	100	56	0	0	.0	
P05	51	29	22	49	19	3	100	60	0	0	.0	
P06	51	27	24	49	17	3	100	62	0	0	.0	
P07	48	25	24	52	15	5	100	58	0	0	.0	

Total SSCH/RSCH	75/s	Page rate	78.6/s	Priv. instruct.	571k/s
Virtual I/O rate	38/s	XSTORE paging	7870/s	Diagnose instr.	567k/s
Total rel. SHARE	6200	Tot. abs SHARE	83%		

Queue Statistics:	Q0	Q1	Q2	Q3	User Status:	
VMDBKs in queue	27	0	0	1	# of logged on users	27
VMDBKs loading	0	0	0	0	# of dialled users	0
Eligible VMDBKs		0	0	0	# of active users	22
El. VMDBKs loading		0	0	0	# of in-queue users	28
Tot. WS (pages)	3330k	0	0	114714	% in-Q users in PGWAIT	1
Expansion factor		% in-Q users in IOWAIT	50
85% elapsed time	6.944	.868	6.944	41.66	% elig. (resource wait)	0

FCX103 Data for 2005/09/01 Interval 01:02:39 - 01:03:40 Monitor Scan

Main storage utilization:

Total real storage	19'456MB
Offline storage frames	0kB
SYSGEN storage size	19'456MB
CP resident nucleus	2'748kB
Shared storage	10'928kB
FREE storage pages	17'228kB
FREE stor. subpools	680kB
Total DPA size	1'969MB
Locked pages	179'924kB
Trace table	2'500kB
Pageable	1'791MB
Storage utilization	70%
Tasks waiting for a frame	9
Tasks waiting for a page	3/s

Paging / spooling activity:

Page moves <2GB for trans.	1521/s
Fast path page-in rate	505/s
Long path page-in rate	2912/s
Long path page-out rate	4453/s
Page read rate	79/s
Page write rate	0/s
Page read blocking factor	11
Page write blocking factor	...

XSTORE utilization:

Total available	5'120MB
Size of CP partition	5'120MB
CP XSTORE utilization	98%
Low threshold for migr.	2'000kB
XSTORE allocation rate	4453/s
Average age of XSTORE blks	359s
Average age at migration	...s
MDCACHE utilization:	
Min. size in XSTORE	0kB
Max. size in XSTORE	5'120MB
Ideal size in XSTORE	5'118MB
Act. size in XSTORE	0kB
Bias for XSTORE	1.00
Min. size in main stor.	0kB

MDISK cache write rate/s
MDISK cache read hit rate	0/s
MDISK cache read hit ratio	...%

VDISKS:

System limit (blocks)	7260k
User limit (blocks)	0
Main store page frames	0
Expanded stor. pages	0

```

FCX114      Data for 2005/09/01  Interval 01:02:39 - 01:03:40      Monitor
.           _____ . . . . . . . . . . . . . . . .
                                           <-SVM and->
Userid      %ACT  %RUN  %CPU  %LDG  %PGW  %IOW  %SIM  %TIW  %CFW  %TI  %EL  %DM  %IOA
>System<    68   11   9    0    1   30   23   1    0    4   0   3   19
LINUX02    100  10   0    0    0    0   90   0    0    0   0   0   0
LINUX03    100  20   7    0    0    0   73   0    0    0   0   0   0
LINUX08    100  20   7    0    0   70   0    0    0    0   0   0   3
LINUX09    100  13   3    0    0    0   83   0    0    0   0   0   0
LINUX10    100  20   7    0    0   70   0    0    0    0   0   0   3
LINUX11    100   3   3    0    3   73   0    0    0    0   0   0   17
LINUX12    100  27  20   0    0   37   0    0    0    0   0   0   17
LINUX13    100  13  13   0    0   70   0    0    0    0   0   0   3
LINUX14    100  17  10   0    0   73   0    0    0    0   0   0   0
LINUX15    100  17  17   0    0   37   0    0    0    0   0   0   30
LINUX17    100   0  10   0    0    0   83   0    0    0   0   0   7
LINUX18    100   7  13   0    3    0   70   0    0    0   0   0   7
LINUX19    100   0  23   0    0    0    0   0    0    0   0   0   77
LINUX21    100  17  10   0    7   53   0    0    0    0   0   0   13
LINUX22    100  17  10   0    0   70   0    0    0    0   0   0   3
    
```

```

FCX113          Data for 2005/09/01  Interval 01:02:39 - 01:0
.
.
.
.
.
Data <----- Paging Activity/s ----->
Spaces <Page Rate>  Page <--Page Migration-->
Userid  Owned  Reads Write Steals >2GB> X>MS MS>X X>DS
>System<   .0   2.9   .0  167  56.7  127  166   .0
LINUX08   0   2.2  .0   .0   363 231 488  .0
LINUX03   0    .7   .0   .0   288 375 294  .0
LINUX09   0    .2   .0   .0   156 189 242  .0
LINUX14   0   6.1  .0   .0   147 241 368  .0
LINUX21   0   5.0  .0   .0   108 309 348  .0
LINUX11   0  38.8  .0   .0   103 45.9 212  .0
LINUX02     0   4.4   .0   .0   71.7  336  448   .0
LINUX12     0    .9   .0   .0   69.5  430  404   .0
LINUX22     0   3.7   .0   .0   52.4  539  531   .0
LINUX13     0    .6   .0   .0   47.3  201  344   .0
LINUX18     0    .7   .0   .0   35.0  16.9  45.8   .0
LINUX17     0  15.4   .0   .0   29.0  21.0  55.5   .0
LINUX10     0    .0   .0   .0   28.8  399  567   .0
LINUX15     0    .0   .0   .0   25.7  84.0  120   .0
    
```

```

FCX113      Data for 2005/09/01  Interval 01:02:39 - 01:03:40
<----- Number of Pages ----->
              <-Resident->
Userid      WSS  Lockd  Resrvd  <2GB  >2GB  XSTOR  DASD  Stor
>System<   128025  393    0   15896 112212  47686  21093  763M
LINUX08    359641  711    0   47526 311895 132269  7659  2048M
LINUX03    218281  638    0   41653 176741  39585  7390  1024M
LINUX09    114714  640    0     641 113379  15952  4680   512M
LINUX14    341990  643    0   10866 329833 151751  8892  2048M
LINUX21    151962  646    0   13895 138277 107081  6375  1024M
LINUX11    155305  681    0     702 155018  94098 24566  1024M
LINUX02    300168  671    0   55301 246056 189731  8137  2048M
LINUX12    162194  687    0   14706 146827  98844  2484  1024M
LINUX22    355789  646    0   77663 278901 134762  6048  2048M
LINUX13    184656  691    0    1752 181825  77116  6468  1024M
LINUX18     99907  639    0     641  99750  12166 45911   512M
LINUX17     96013  639    0     641  95846  20931 50758   512M
LINUX10    415979  689    0 161184 255310  74808  4158  2048M
LINUX15    114048  646    0     672 114060 136837 17740  1024M

```

Methods to Lower Contention <2GB in z/VM 5.1.0

- **Minimize number of dedicated OSA, FCP, and HiperSockets devices**
 - Perhaps use a virtual router or the Virtual Switch (z/VM 4.4)
- **Minimize amount of virtual disks in storage space created**
- **Minimize size of Linux Guests**
 - The more pages available to Linux the more pages that can potentially be used for I/O data buffers
- **Explore use of Minidisk Cache**
- **Explore creating multiple VM systems (logical partitions) to increase amount of effective under 2GB memory**
- **Review application and product options to control impact**
- **See www.vm.ibm.com/perf/tips/2gstorag.html for more details and discussion.**

Mitigations for Performance in z/VM 5.1.0

- **Explore Linux Fixed I/O Buffers feature in SLES 9 SP1 or RHEL 4**
 - Also available in SLES 8
- **VM63729**
 - Improves efficiency of finding free frames below 2GB
- **VM63730**
 - Improves performance of contiguous frame management
- **VM63752**
 - More aggressive paging to above 2GB instead of expanded storage

- **Above is not a replacement for z/VM 5.2.0**

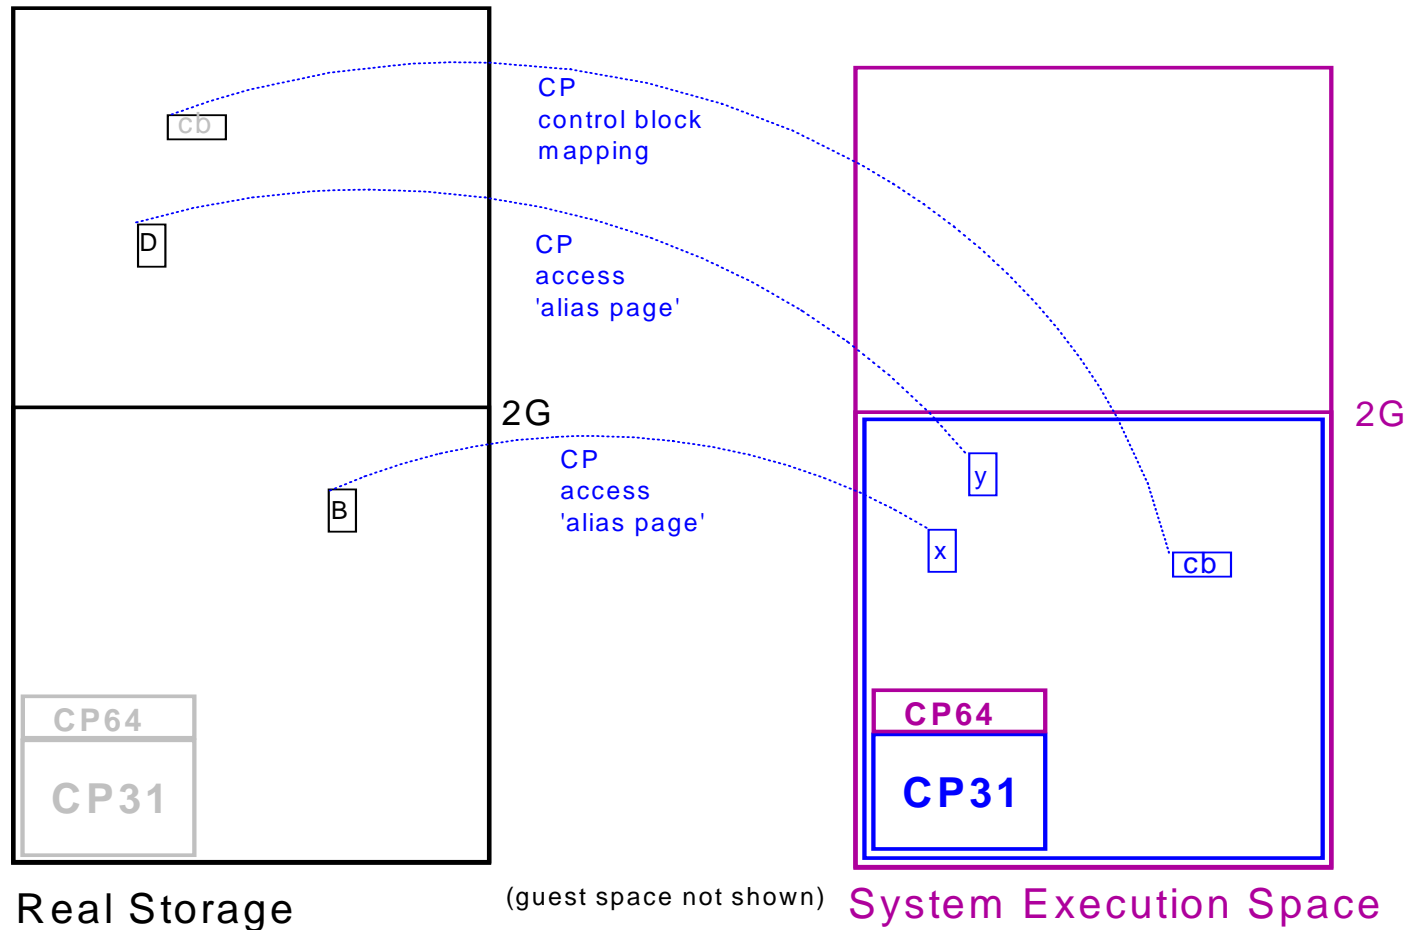
Large Real Memory – z/VM 5.2.0

- **Frame table moved above 2GB**
- **Key areas in CP converted to 64-bit**
 - Additional storage management
 - I/O Processing
- **Additional exploitation of access register mode**
- **For remainder, CP now runs in its own address space**
 - System Execution Space (SXS)
 - CP now runs DAT ON all the time
 - Guest pages no longer moved for CP access
 - Guest pages can be “aliased” to a page in the SXS

Overview z/VM 5.2 Storage Usage

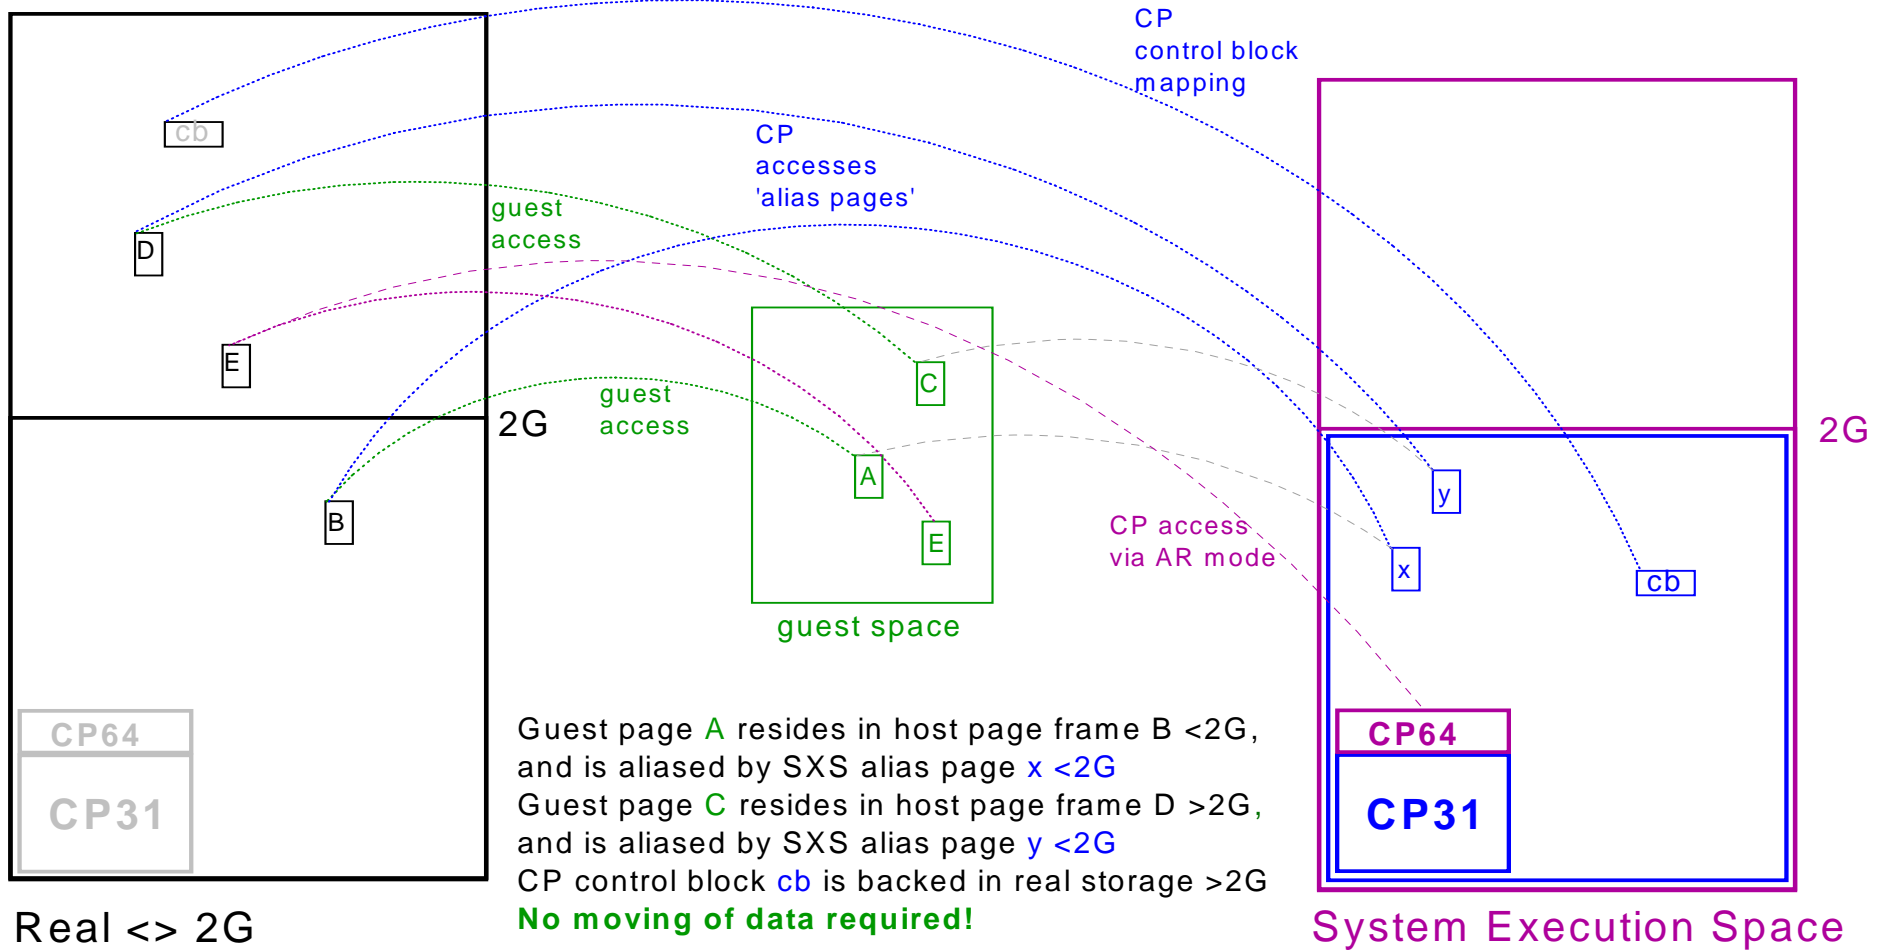
IBM

- CP executes:
 - DAT on
 - In a non-identity mapped System Execution Space
- Most of CP still:
 - Executes in 31-bit addressing mode
 - Can only touch storage below 2G in the System Execution Space (aka "Host Logical")
 - But **CAN** touch storage above 2G in real storage (implicitly)
- Parts of CP which touch real storage:
 - Mostly execute in 64-bit addressing mode and can touch real storage above 2G



Overview z/VM 5.2 Storage Usage

IBM



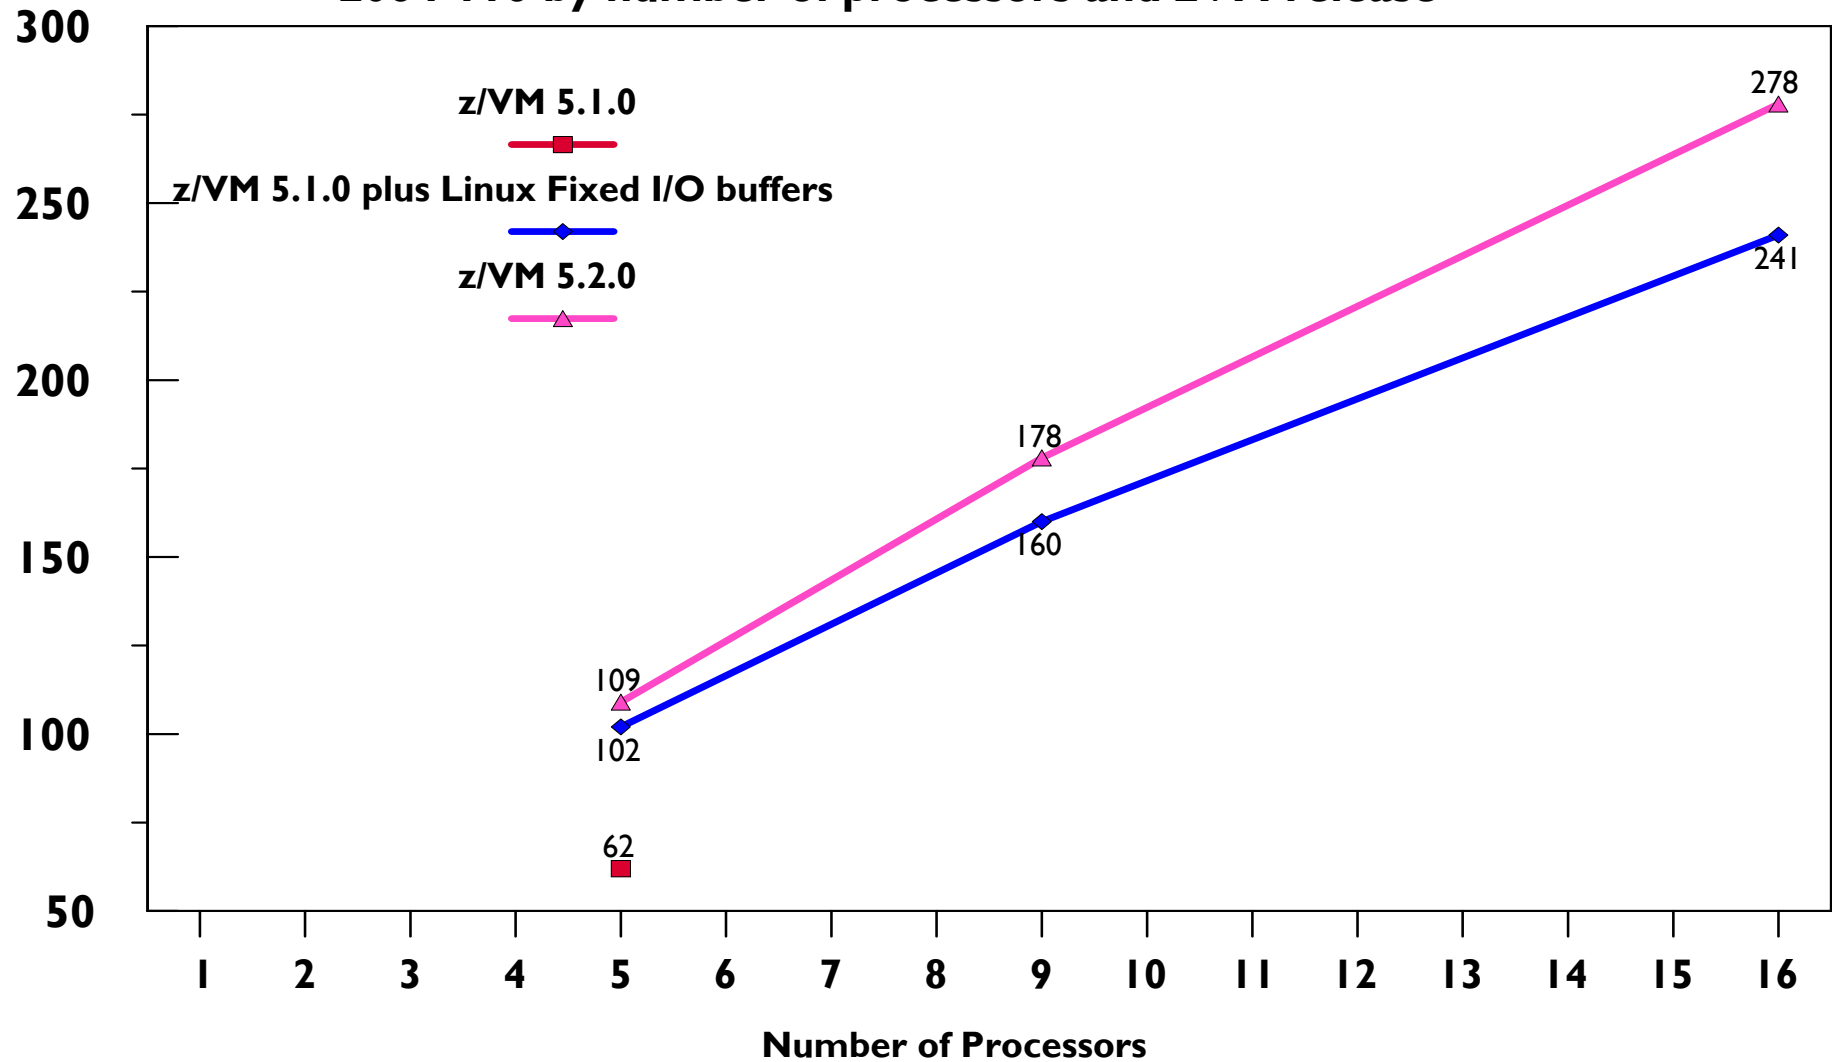
Large Real Memory – z/VM 5.2.0 - Results

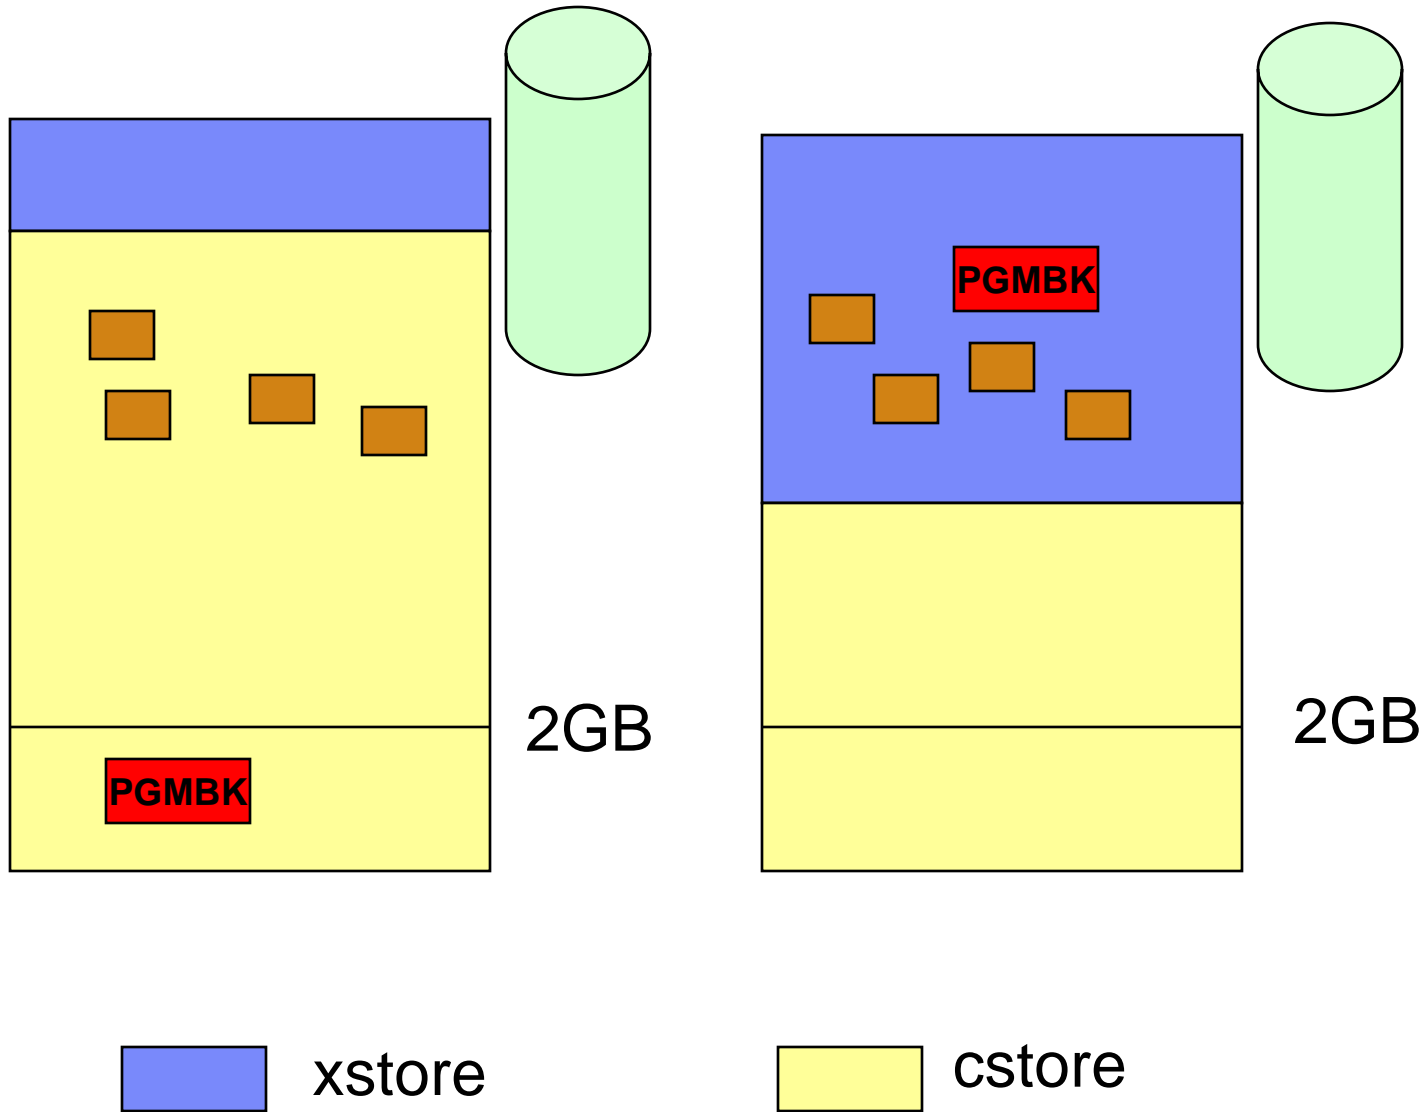
- **Large improvement for systems constrained for memory under 2GB.**
 - Will **not** see claims such as “Can run nn% more virtual servers”.
 - Think of it as “now scales as one would expect”
 - Can consider reversing mitigations
 - Turn off Linux Fixed I/O Buffer
 - Consolidating LPARs
 - Example: Informix workload didn’t scale past 6GB on z/VM 5.1.0; scales as expected on z/VM 5.2.0 (measured up to 16GB guest)
- **Encouraging results from various test environments**

Large Real Memory – z/VM 5.2.0 - Results

- **Unconstrained environments may see a processor time increase.**
 - Goal is an increase <5% on average. Current range is up to ~10% increase. Biggest hit for very storage constrained environments (but not 2GB constrained)
 - Hit from instruction mix, linkage, and mode changes
- **Best results when 4GB or more of memory exists**
 - 2GB configurations have all the disadvantages and few of the advantages
- **Still recommend configuring some expanded storage in paging environments**
 - The greater the mix of idle and active virtual machine, the greater the benefit to expanded storage for smoothing page response time.
 - No longer need xstore to mitigate below 2GB constraints

Throughput (Transaction per second) Apache Workload (5000 files,5 DASD,primed MDC XSTORE) 2064-116 by number of processors and zVM release





Emulated FBA on SCSI Disk Improvements


- **Paging to Emulated FBA on SCSI**
 - Processor time down from z/VM 5.1.0
 - Still more expensive in processor time than paging to ECKD (~ 5 times more expensive)
 - Throughput (pages per second) up from z/VM 5.1.0
 - Higher paging rates achievable than ECKD (~70% more pages/second)
- **Improvements to pathlength for of Emulated FBA on SCSI**
 - Benchmark results show additional improvements in processor time
 - Still significant overhead compared to traditional ECKD
- **Performance report has additional comparison data for different disk I/O configurations.**

V=V QDIO Pass Through Stage II

- **QDIO Enhanced Buffer State Management (QEBSM)**
- **HW requirement:**
 - z990 or z890 with QEBSM Enablement or z9
 - <http://www.vm.ibm.com/perf/qebsm.html> for more details
- **z/VM 5.2.0**
- **Linux Changes Required**
- **Minimizes CP involvement for FCP and QDIO I/O**
- **Reduces Total CPU 5% to 30% for network and disk building block workloads.**
 - Higher virtual CPU time (Linux) because more is done under SIE
 - CP CPU time (VM) decreases greatly


QEBSM Improvements

OSA Streaming Workload 50 Clients (1492 MTU)

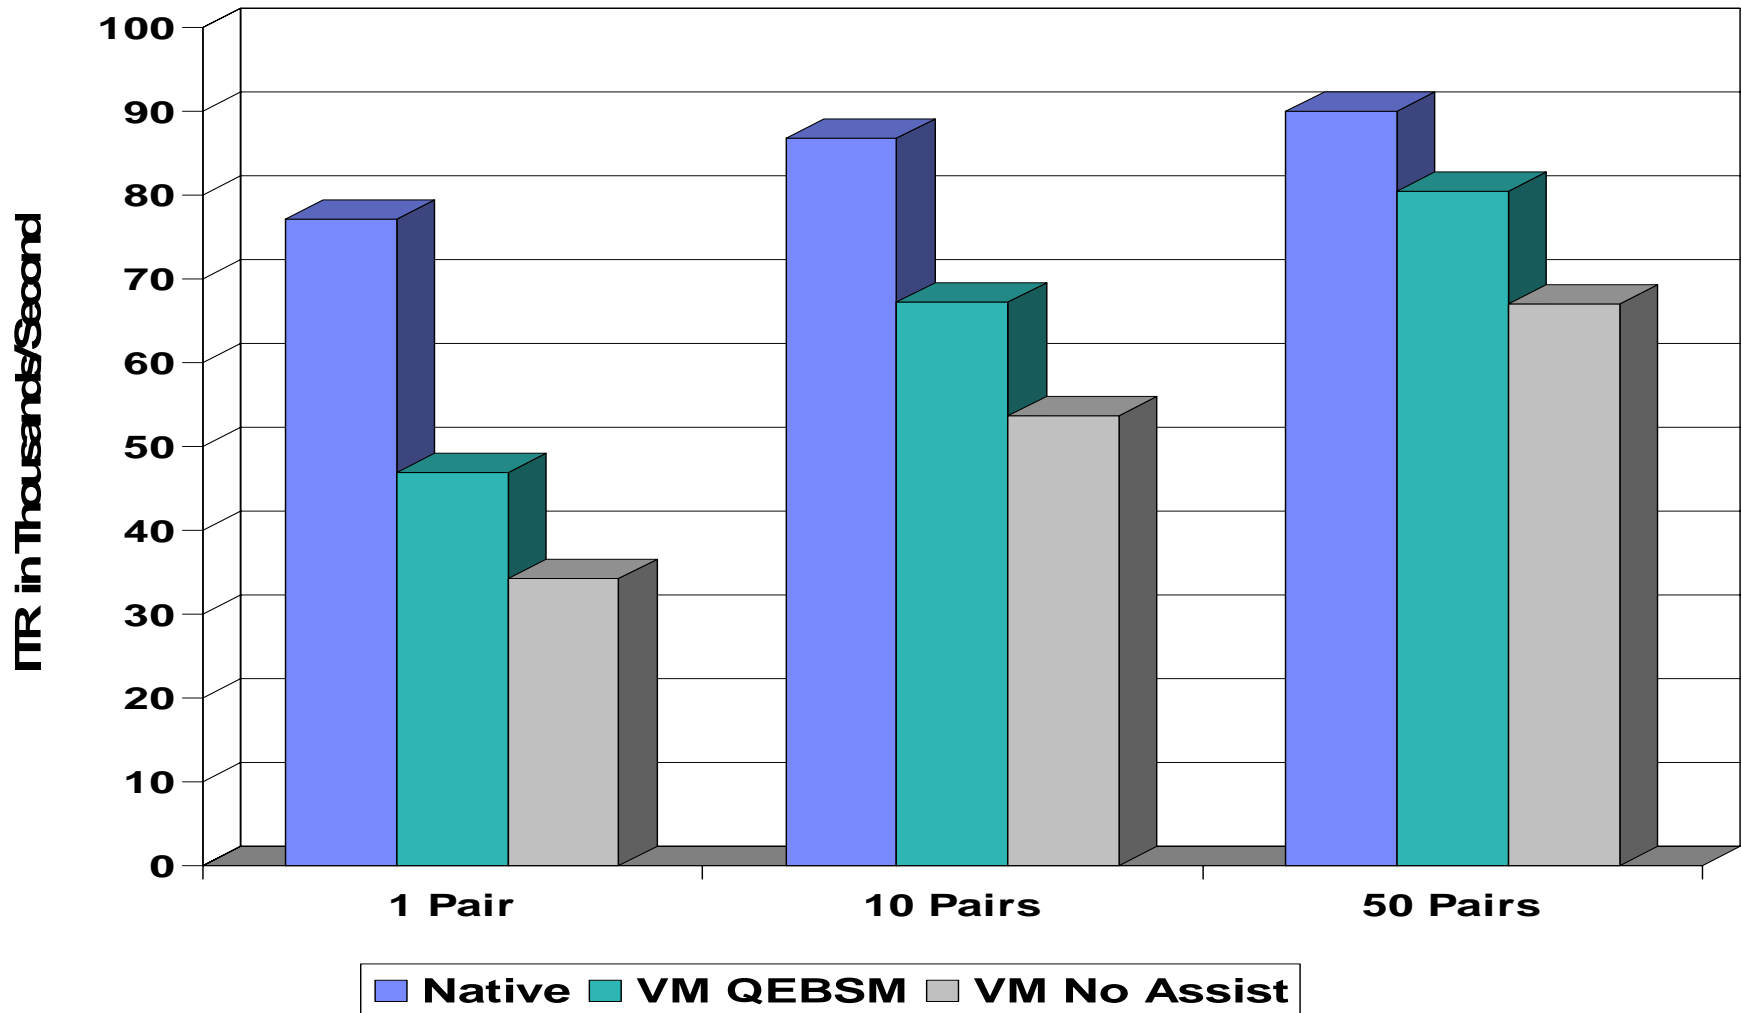


QEBSM Improvements

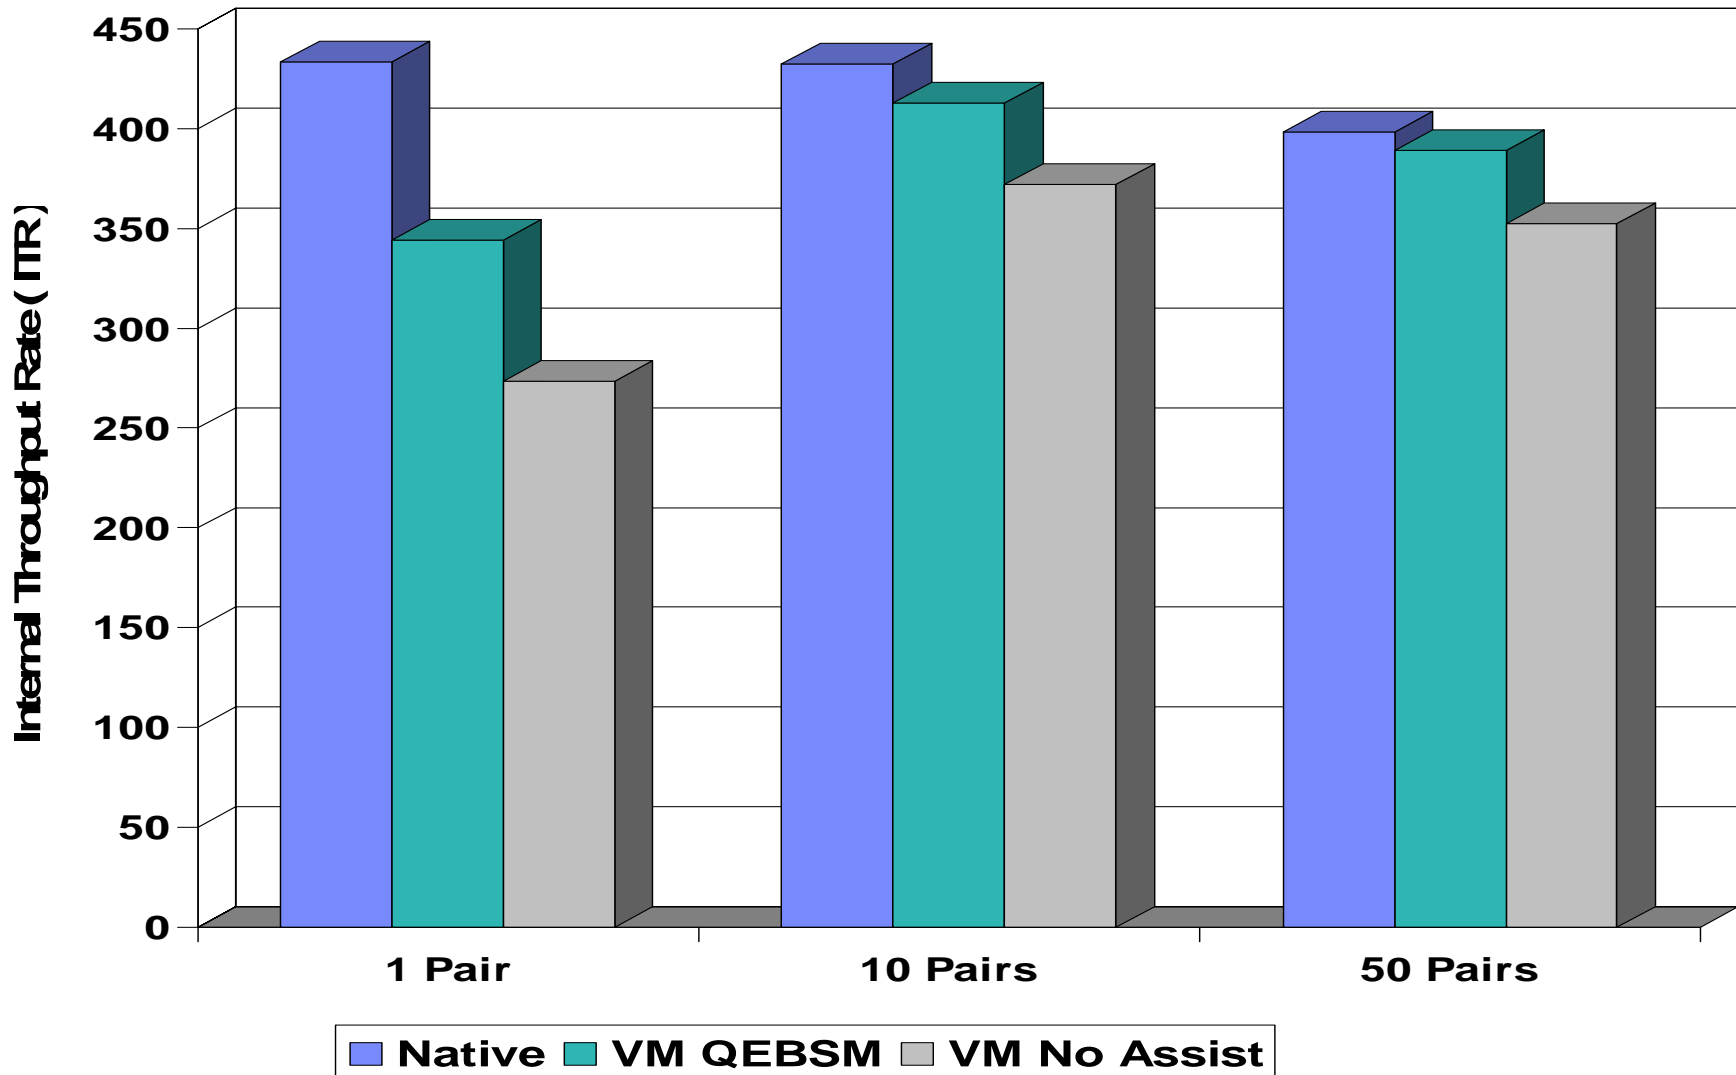
OSA Streaming Workload 50 Clients (8992 MTU)



Results with QEBSM – Network RR Workload



Results with QEBSM – Network Streams Workload



Diagnose x'44' Processing

- **Diagnose x'44': Voluntary Time Slice End**
- **Linux uses in spin lock processing to yield to other virtual processors in the complex**
- **z/VM 5.1.0 fast path to check case where on a virtual 2-way where the other processor is already running.**
- **z/VM 5.2.0 extends the fast path to cover any size virtual n-way**
- **Lowers stress on VM scheduler lock**
- **Potential to lower System CPU time and Scheduler Lock time in z/VM 5.2.0**
 - Can be seen in Performance Toolkit

Directory Performance Enhancements

- **Changes in VM virtual machine directory processing and DIRMAINT**
 - New online via Delta option
 - Allows virtual machine directory entries to be added, deleted, or updated without having to rewrite the entire object directory.
- **Significant improvements for large directories**
- **No noticeable improvement for smaller directories**

Service Related to z/VM 5.2.0

- **Always check the install bucket for latest info**
- **A few key performance APARs:**
 - PK11392 – Improves TCP/IP performance when using VIPA with multiple interfaces.
 - VM64009 - Corrects Monitor memory leak when running on processors without crypto devices.
 - VM63961 - SPXTAPE LOAD filling up spool with unaccounted pages.
 - VM63958 - Slow IPL of z/OS guests
 - VM63955 - SYSTEM HANG ON SRMSLOCK SCHEDULER LOCK AND RSAAVLLK
 - VM63935 - Avoids unnecessary clearing of memory in CMS
 - VM63916 - Corrects ASCCTPGS field updating and bad values in monitor regarding paging counts.
 - VM63893 - Improves minidisk cache performance.
 - VM63892 - Corrects locked page counts used by Q FRAMES and CP monitor.
 - VM63877 - Corrects an error in the value reported for the LockedRS field in the Q FRAMES output.
 - VM63845 - Corrects high demand scan CPU time that occurs for certain paging workloads
 - VM63841 - Improves performance of emulated FBA on SCSI. This also applies to z/VM 5.1.0

PAV Exploitation for VM Minidisks

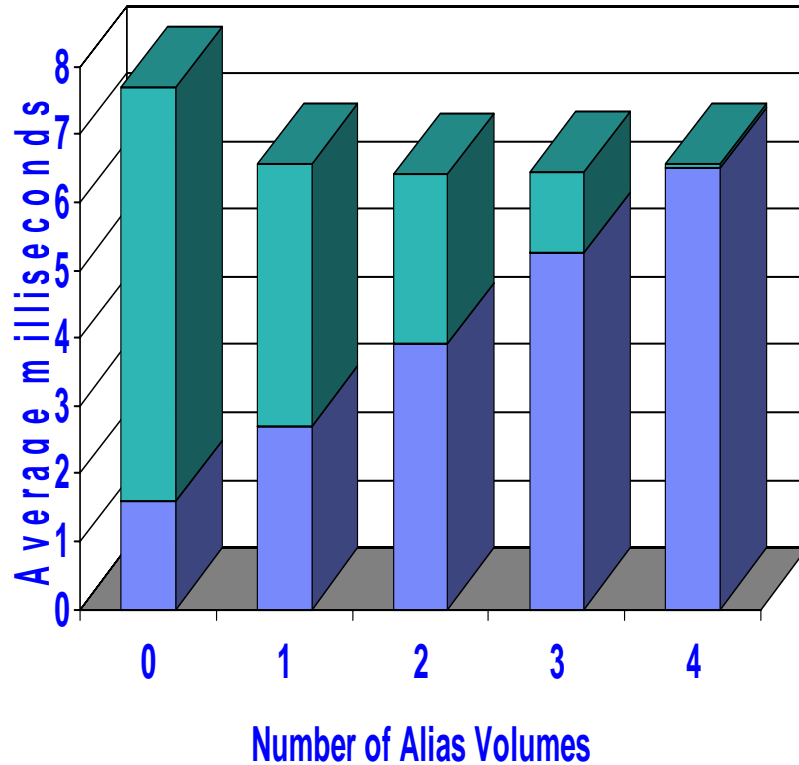
- **Previously only supported as dedicated disks.**
- **VM63855 (for z/VM 5.2.0, available May 2006):**
 - CP uses PAV to potentially decrease response time on minidisk I/O
 - We tightened the rules about ATTACHing or DEDICATEing PAV devices
- **VM63855 – virtualizes PAV for minidisks.**
- **Useful for environments where queuing on I/O occurs for minidisk I/O.**
- **Sometimes referred to as SYSTEM-owned PAV volumes**
- **PAV Base and Alias volumes defined on the Storage CU**
- **Summary of Results**
 - Varies depending on DASD CU Model
 - Varies depending on read-write mix
 - Helpful when I/O queuing occurs
 - Law of diminishing return; that is, defining more Alias than needed can lower performance

PAV – Rules of Thumb


- **Symptom:**
 - I/O wait queue forming at real volume where minidisks are
 - See Performance Toolkit FCX168 reports (or equivalent)
- **Remedy:**
 - Configure a PAV alias device in the storage controller
 - Make sure the alias device is varied online
 - Make sure the alias device is ATTACHed to SYSTEM
- **Measure:**
 - Re-run your workload
 - Look again at those disk performance reports
- **Success criterion:**
 - Response time equals service time (no wait queue)

System Owned PAV Results – DS8100 – 100% Writes

Device Response Time



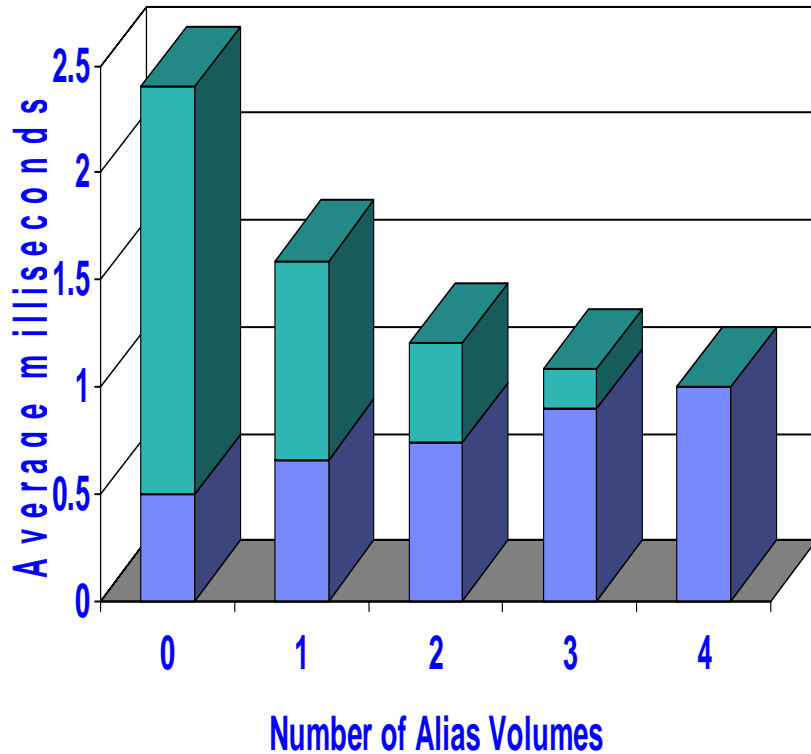
I/O Rate



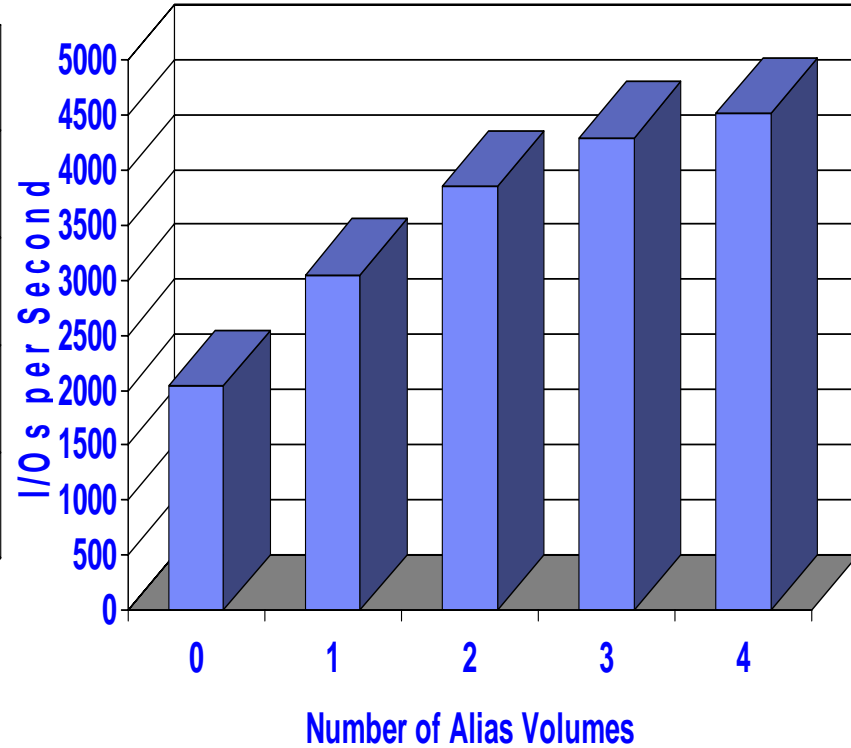
■ Serv Time ■ Queue Time

System Owned PAV Results – DS8100 – 100% Reads

Device Response Time




I/O Rate




■ Serv Time
 ■ Queue Time

System Owned PAV Results – DS6800 – 100% Writes

Device Response Time




I/O Rate




■ Serv Time ■ Queue Time

System Owned PAV Results – DS6800 – 100% Reads

Device Response Time



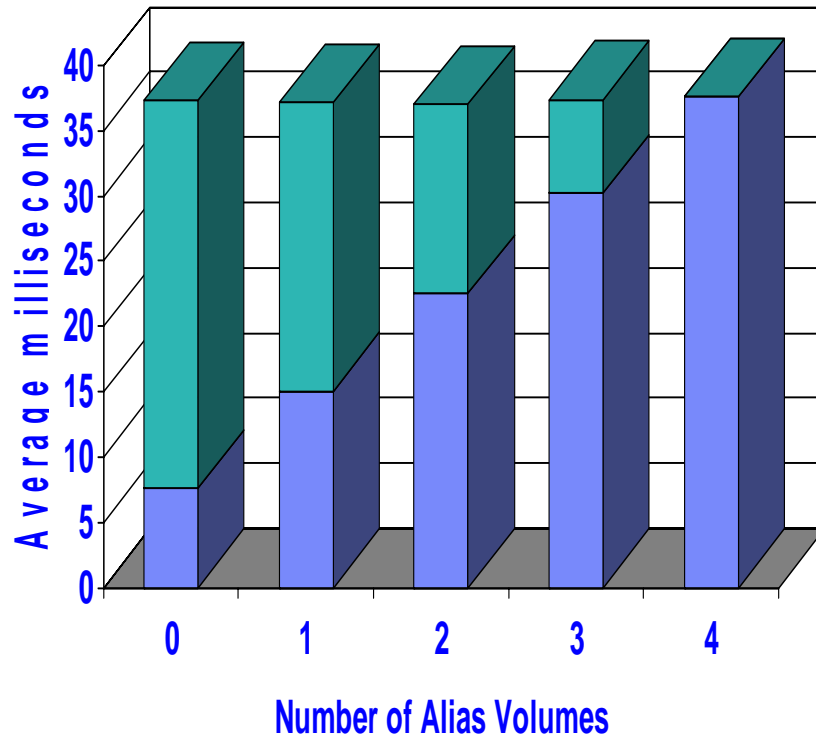
I/O Rate



■ Serv Time ■ Queue Time

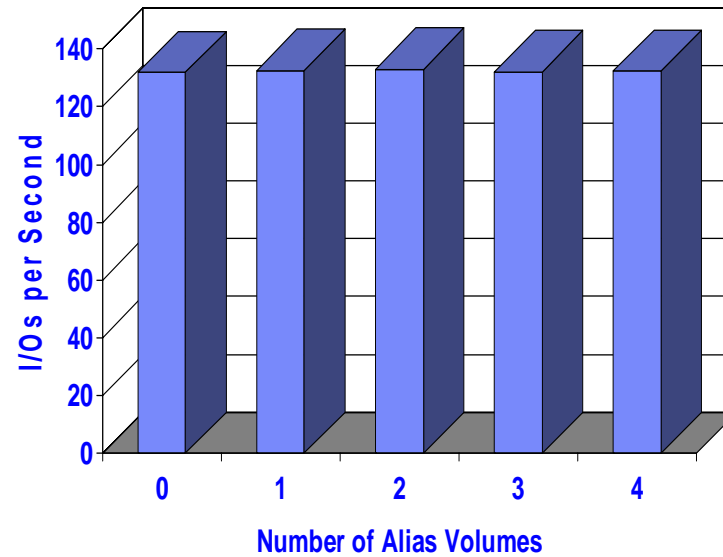
System Owned PAV Results – ESS F20 – 100% Writes

Device Response Time




■ Serv Time ■ Queue Time

I/O Rate




System Owned PAV Results – ESS F20 – 100% Reads

Device Response Time



I/O Rate




■ Serv Time ■ Queue Time

Cooperative Memory Management

- **VMRM Support of Linux CMM API**
- **Evaluates z/VM memory usage and notifies Linux guests to release memory**
- **z/VM 5.2.0 APAR VM64085**
 - <http://www.vm.ibm.com/sysman/vmr/vmrmmcmm.html>

Transaction Rate vs. Real Storage

for various Storage Management Products using 32 servers with a virtual memory size of 1.5GB



Bigger is Better

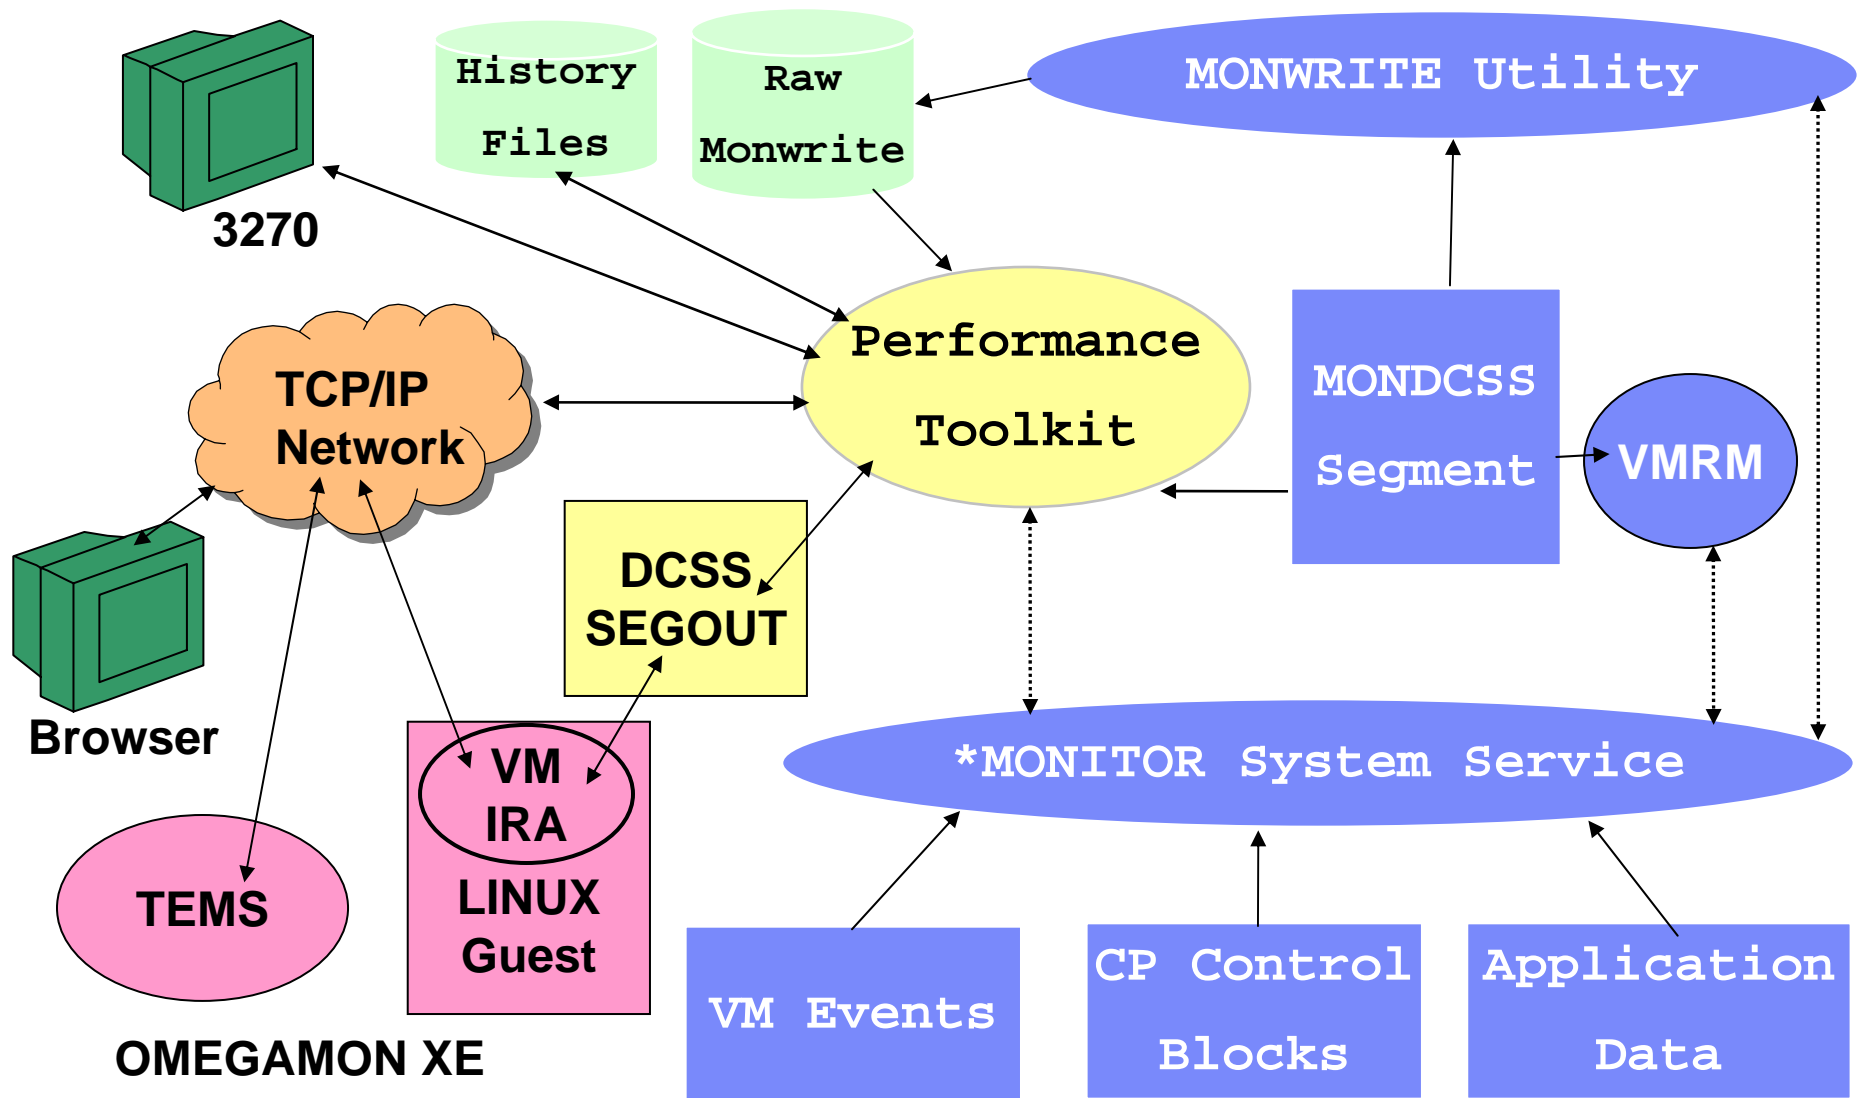
Smaller is more interesting

Performance Toolkit

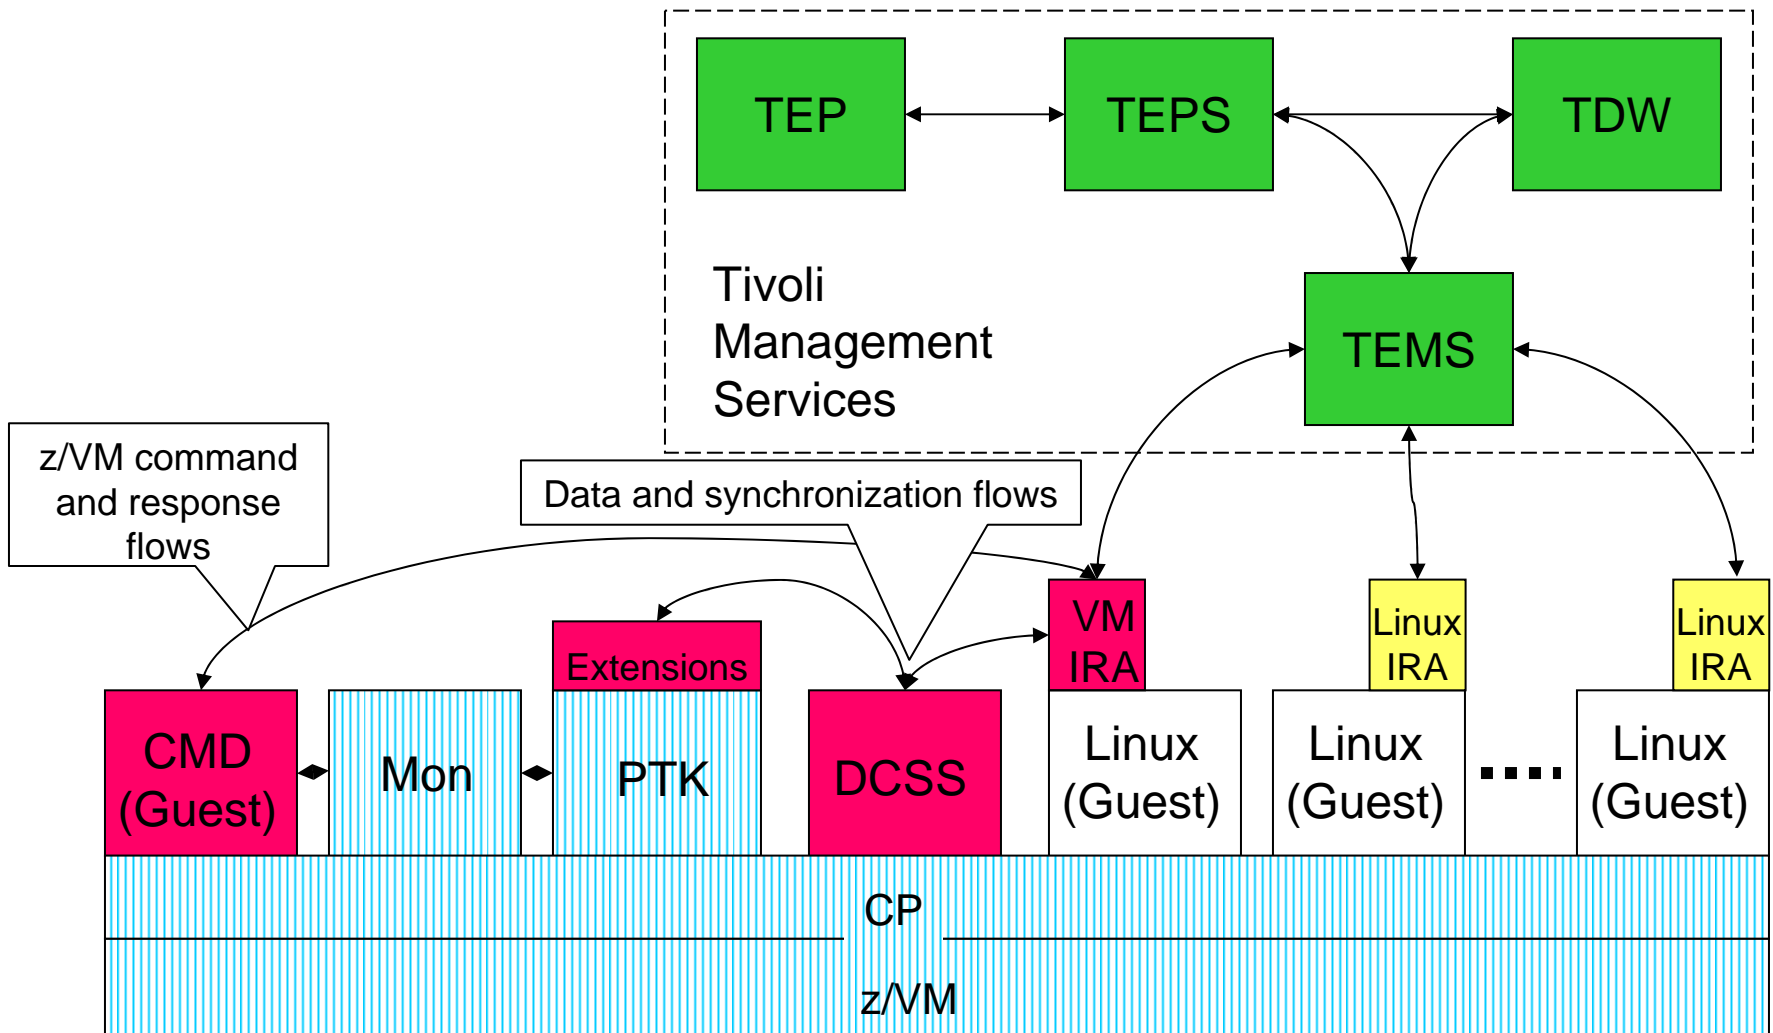
- **New screens/reports**

- FCX260 - STORMENU (option 3A from main menu)
- FCX253 - STORLOG
- FCX254 - AVAILLOG
- FCX259 - DMNDLOG
- FCX261 - SXS AVAIL
- FCX262 - SXSPAGE
- FCX263 - SXSDEFER
- FCX264 - SXSUTIL
- FCX257 - UQDIO
- FCX258 - UQDIOLOG
- FCX255 - QEBSM
- FCX256 - QEBSMLOG

5,000 Foot View



Basic Architecture



PAGING and SPOOLING Utilization

Paging_Spooling - PHKMSM - SYSADMIN *ADMIN MODE*

File Edit View Help

View: Physical

Enterprise

- Linux Systems
 - vmInx10
 - Linux OS
 - z/VM Linux Systems
 - CP-Owned_Devices(Page_Spool)
 - DASD
 - LPAR
 - Network
 - Real_Storage
 - System

Physical

Paging and Spooling Space

Top 5 Page Extent Utilization

Top 5 Dump Extent Utilization

Top 5 Spool Extent Utilization

CP-Device Table (Paging and Spooling)

Time	System ID	LPAR Name	Device VOLSER	Device Address	PAGING/SPOOLING	Allocation	Avilable Slots	Device Type	Device End Extent	Device Percent Full	Device Start Extent	Device Slots Used
06/21/06 09:49:40	WLAVMXA	LPAR001	VMSY03	2412	PAGING	10	1	3390	540	1	5	3
06/21/06 09:49:40	WLAVMXA	LPAR001	VMSY03	2114	SPOOLING	23	12	3390	540	52	5	3
06/21/06 09:49:40	WLAVMXA	LPAR001	VMSY03	2213	SPOOLING	33	33	3390	540	0	5	3
06/21/06 09:49:40	WLAVMXA	LPAR001	VMSY03	1423	UNKNOWN	12	1	3390	540	8	5	3

Hub Time: Wed, 06/21/2006 09:50 AM Server Available

Paging_Spooling - PHKMSM - SYSADMIN *ADMIN MODE*

Start Windows Task Manager CNPS Manage Tivoli Enterprise ... ~ 9:50 AM

Paging_Spooling - PH... ~ ~ Document1 - Microsoft W...

REAL STORAGE Utilization

Real_Storage - PHKMSM - SYSADMIN *ADMIN MODE*

File Edit View Help

View: Physical

Enterprise

- Linux Systems
 - vmInx10
 - Linux OS
 - z/VM Linux Systems
 - CP-Owned_Devices(Page_Spool)
 - DASD
 - LPAR
 - Network
 - Real_Storage**
 - System

Physical

Administer Users... (Ctrl+U)

Available Frames Mean

System Page Rate

System Resource Utilization

Page Wait Queue

z/VM Storage Utilization

Time	System ID	LPAR Name	Number of Frames	Number of Frames > 2GB	Available Frames High Thresh	Available Frames High Thresh > 2GB	Available Frames Mean	Available Frames Mean > 2GB	Available Pages Low Thresh	Available Pages Low Thresh > 2GB	System Paging Rate	Number of Dynamic Frames	Demand Scan Fails	Free Use
06/21/0...	WLAVMXA	LPAR001	3	3	28000	28000	28	333	71	55	332	0	22	

Hub Time: Wed, 06/21/2006 09:49 AM Server Available

Real_Storage - PHKMSM - SYSADMIN *ADMIN MODE*

Windows Task Manager, CNPS, Manage Tivoli Enterprise ...

Real_Storage - PHKM..., Document1 - Microsoft W...

9:49 AM

SYSTEM Utilization

Time	System ID	LPAR Name	Active Users	Average Number of Users	Tasks in Wait	CP Percent of CPU	Percent of CPU	Number Active CPU's	Number of Dialed Users	Average Number of Eligible Users	Expand User Queue 1	Expand User Queue 2	Expand User Queue 3	LPAR Busy Percent	Normal T P
06/21/06 09:49:40	WLAVMXA	LPAR001	123	5	4	24	91	2	44	22.00	112.00	2,113.00	443.00	74.00	

z/VM 5.3.0 Performance Sneak Peek

- **Greater than 24 CPU Support**
- **Greater than 128GB Memory Support**
- **Virtual Switch Link Aggregation (SDO)**
- **HyperPAV Support**
- **Monitor Enhancements**

Greater than 24 CPU Support

- **While z/VM 5.2.0 would run on up to 31 processors, it only supported 24 due to performance limitations**
- **Serialization Changes**
 - General support for exclusive and shared formal spin locks
 - First to exploit is the Scheduler Lock
 - New lock associated with Processor Local Dispatch Vector (PLDV) for dispatching
- **z/VM 5.3.0 will support 32 processors**

Greater than 128GB Memory Support

- **PGMBKs allowed to be allocated above 2GB**
- **Required Control Register 64-bit support**
- **Also enhanced contiguous frame management**
- **Limitations dependent on workload & configuration**
 - 256GB Real memory
 - 8TB of in use virtual machine memory

Virtual Switch Link Aggregation (SDO)

- **Ability to attach multiple OSAs to a single virtual Switch**
 - Aggregate bandwidth
 - Failover
- **Requires z9 OSA2 Express Support**
- **Dynamic Load Balancing**
 - Can be influenced by distribution of MAC addresses

Summary

- **z/VM 5.2.0 is an awesome release!**
 - Does not solve all performance problems!
 - 64-bit exploitation
 - Minimal impact to regression environments
 - Lifts configuration limitations (number of dedicated OSA or FCP)
 - Removes contention for guest pages below 2GB
 - Emulated FBA over SCSI
 - Reductions in path length and processor requirements
 - But ECKD still has performance advantage
 - QEBSM - Hardware Assists
 - FCP and QDIO I/O
 - z990, z890, and z9
 - Reduces VM involvement in I/O
- **Post z/VM 5.2.0 Additional Improvements**
- **z/VM 5.3.0 looks to be another great release**
- **See z/VM Performance report for details when available**
<http://www.vm.ibm.com/perf/reports/>