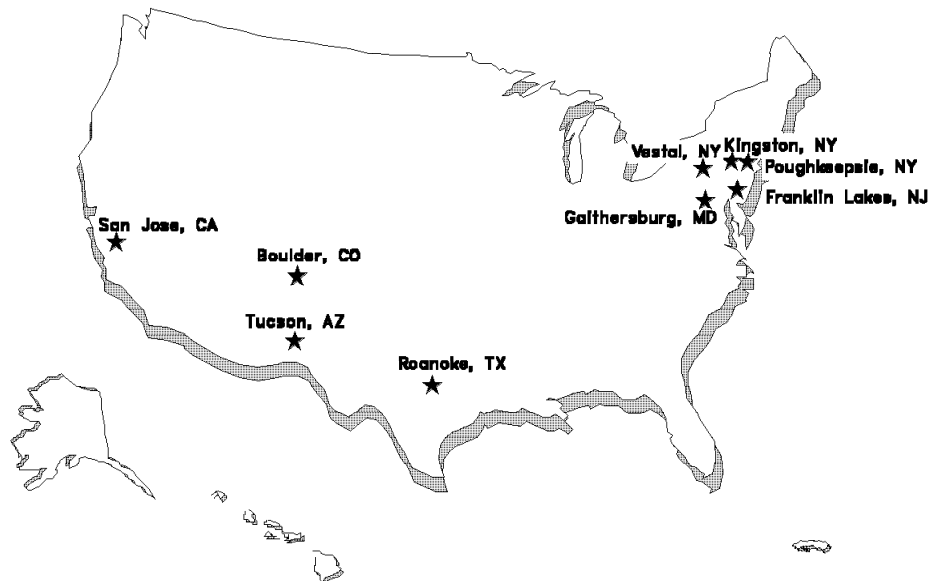


Enterprise Systems Services and Support

VM/ESA Release 2 Performance Report

C.L. Morse
J.L. Thrall



Washington Systems Center Technical Bulletin

GG66-3245-00
IBM Internal Use Only - This copy not for customer use

VM/ESA Release 2 Performance Report

C.L. Morse
J.L. Thrall

First Edition

This book applies to VM/ESA Release 2. (Program Number 5684-112).

Publications are not stocked at the address given below. Requests for IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

A form for readers' comment is provided at the back of this publication. If the form has been removed, comments may be addressed to: Chuck Morse, IBM Corporation, Washington Systems Center, 800 N. Frederick Avenue, Gaithersburg, MD 20879-3395.

© **Copyright International Business Machines Corporation 1993. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	v
Tables	v
Notices	ix
Programming Interfaces	ix
Trademarks	ix
Acknowledgements	x
Abstract	xi
Referenced Publications	xiii
Summary of Key Findings	1
Changes That Affect Performance	5
Performance Improvements	6
Performance Considerations	18
Performance Management	19
Measurement Information	25
Hardware Used	25
Software Used	25
Format Description	25
Tools Description	27
Migration to VM/ESA Release 2	29
CMS-Intensive Migration from VM/ESA Release 1.1	30
9021-900 / Minidisk	33
9021-900 / Minidisk / Storage-Constrained	38
9121-480 / Minidisk	42
9121-480 / Minidisk / Storage-Constrained	45
9121-480 / SFS	48
9221-170 / Minidisk	52
CMS-Intensive Migration from VM/SP HPO Release 5	56
9121-320 / Minidisk	57
9121-480 / Minidisk	62
CMS-Intensive Migration from VM/SP Release 5	67
CMS-Intensive Migration Summary	73
OfficeVision Migration from VM/ESA Release 1.1	79
9121-480	81
9021-720	85
MVS/ESA Guest Migration from VM/ESA Release 1.1	89
VSE/ESA Guest Migration from VM/ESA Release 1.1	92
9121-320 / V=V / MODE=ESA	94
9121-320 / V=V / MODE=VMESA	96
9121-320 / V=R / MODE=ESA	98
VSE/ESA Guest Migration from VM/SP Release 5	100
9221-170	102
9221-120	105
VMSES/E Migration from VM/ESA Release 1.1	108

9121-480 / Receive, Apply and Build Service	109
9021-720 / Receive, Apply and Build Service	112
New Functional Enhancements	115
FBA DASD Support	116
9221-170 / Minidisk	121
9221-170 / VSE Guest	131
9121-320 / VSE Guest	135
CP Configurability	138
DIAGNOSE Code X ϕ 250 ϕ	140
Extended CMS File System Interfaces	143
Tuning Considerations	145
CMS File Cache for SFS	146
I/O Assist for Guests	150
Additional Evaluations	161
Processor Capacity	162
Minidisk to SFS	168
Measurement Variability	173
LPAR Performance	179
RACF	185
9021-900 / RACF Release Comparison	186
9021-900 / RACF Security Levels	190
9021-900 / RACF Multiple Servers	194
VSE Guest Using Shared DASD	199

Appendixes	217
Appendix A. SFS Counter Data	219
Appendix B. Workloads	223
CMS-Intensive (FS7B)	224
IBM Office Benchmark (IOB V2.1)	234
VSE Guest (PACE)	237
VSE Guest (VSECICS)	239
MVS Guest	241
Appendix C. Configuration Details	243
Glossary of Performance Terms	245

Figures

1.	Internal Throughput for the CMS-Intensive Workload	31
2.	External Response Time for the CMS-Intensive Workload	32
3.	Internal Throughput Rate for OfficeVision/VM	79
4.	External Response Time for OfficeVision/VM	80
5.	Internal Throughput for VSE Guest Batch Workload	92
6.	External Throughput for VSE Guest Batch Workload, Migration from VM/SP 5	100
7.	Internal Throughput for VSE Guest Batch Workload, Migration from VM/SP 5	101
8.	External Response Time for FS7B0R FBA and CKD Measurements	116
9.	Internal Throughput for FS7B0R FBA and CKD Measurements	117
10.	External Throughput for VSE V=V Guest Batch FBA and CKD Measurements	118
11.	Internal Throughput for VSE V=V Guest Batch FBA and CKD Measurements	119
12.	MVS/ESA RMF Direct Access Device Activity report.	154
13.	RTM VM/ESA Device Screen.	154
14.	VMPRF DASD_BY_ACTIVITY.	154
15.	VMPRF DASD_IO_ASSIST.	155
16.	VMPRF DASD_BY_ACTIVITY_EF.	155
17.	RTM VM/ESA Device Screen.	157
18.	VMPRF DASD_BY_ACTIVITY.	157
19.	VMPRF DASD_IO_ASSIST.	157
20.	EXPLORE/VSE XA DEVICE STATISTICS	158
21.	RTM VM/ESA Device Screen.	158
22.	VMPRF DASD_BY_ACTIVITY.	159
23.	VMPRF DASD_IO_ASSIST.	159
24.	Internal Throughput Rate for the 9021-900 Processor.	162
25.	Internal Throughput for Single VSE Guest on 9221-200	200
26.	External Throughput for Single VSE Guest on 9221-200	200
27.	Internal Throughput for Single VSE Guest on 9221-170	204
28.	External Throughput for Single VSE Guest on 9221-170	204
29.	Internal Throughput for VSE CICS Workload on 9221-170	207
30.	Internal Throughput for Two VSE Guests on 9221-200	210
31.	External Throughput for Two VSE Guests on 9221-200	210
32.	Internal Throughput for Two VSE Guests on 9221-200	214
33.	External Throughput for Two VSE Guests on 9221-200	214

Tables

1.	Monitor Enhancement and Free Storage Calls	11
2.	Minidisk-Only CMS-Intensive Migration from VM/ESA Release 1.1 to VM/ESA Release 2 on the 9021-900.	36
3.	CMS-Intensive Minidisk Only Storage Constrained on the 9021-900.	40
4.	Minidisk Migration from VM/ESA Release 1.1 to VM/ESA Release 2 on the 9121-480.	43
5.	Minidisk Storage Constrained from VM/ESA Release 1.1 to VM/ESA Release 2 on the 9121-480	46

6.	SFS migration from VM/ESA Release 1.1 to VM/ESA Release 2 on the 9121-480	50
7.	Minidisk migration from VM/ESA Release 1.1 to VM/ESA Release 2 on the 9221-170	54
8.	CMS-Intensive Migration from VM/SP HPO Release 5 on the 9121-320	60
9.	CMS-Intensive Migration from VM/SP HPO Release 5 on the 9121-480	65
10.	CMS-Intensive Migration from VM/SP Release 5 on the 9221-120	70
11.	Approximate VM Relative Capacity: CMS-Intensive Environment	73
12.	Derivation and Supporting Data: VM Measurement Pairs	77
13.	OfficeVision Migration of VM/ESA and OV/VM on a 9121-480	83
14.	OfficeVision Migration of VM/ESA on a 9021-720	87
15.	MVS Batch Measurement Data: MVS 4.2.2 Guests on VM/ESA Release 1.1 and VM/ESA Release 2	90
16.	VSE Guest Comparison: V=V MODE=ESA on a 9121-320	94
17.	VSE Guest Comparison: V=V MODE=VMESA on a 9121-320	96
18.	VSE Guest Comparison: V=R MODE=ESA on a 9121-320	98
19.	VSE Guest Migration from VM/SP 5 on 9221-170	104
20.	VSE Guest Migration from VM/SP 5 on 9221-120	107
21.	VMFREC Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9121-480	110
22.	VMFAPPLY Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9121-480	111
23.	VMFBLD (STATUS Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9121-480	111
24.	VMFBLD (SERVICED Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9121-480	111
25.	VMFREC Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9021-720	112
26.	VMFAPPLY Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9021-720	113
27.	VMFBLD (STATUS Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9021-720	113
28.	VMFBLD (SERVICED Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9021-720	113
29.	9221-170 VM/ESA Release 2 CMS Migration on FBA and CKD Systems	124
30.	VM/ESA Release 2 CMS Migration on FBA and CKD Systems with MDC on 9221-170	129
31.	VM/ESA Rel. 2 VSE Guest Migration on FBA and CKD DASD on 9221-170	134
32.	VM/ESA Release 2 VSE Guest Comparison between FBA and CKD DASD on 9121-320	137
33.	System IPL Comparison	139
34.	DIAGNOSE Code X'250' Performance - Comparison to *BLOCKIO	141
35.	DIAGNOSE Code X'250' Performance - Synchronous Interface Benefits	142
36.	Path length comparison between FS macros and CSL functions.	143
37.	SFS - CMS File Cache Size Comparison: 9121-480	148
38.	Measurement-Block Updates for I/O	152
39.	VM/ESA Release 2 on the 9021-900 Using Different Numbers of Processors	166
40.	Minidisk to SFS Comparison / Equal Processor Utilization on 9121-480	171
41.	Measurement Variability of FS7B0R on the 9121-480 - Results	175
42.	Measurement Variability of FS7B0R on the 9121-480 - Statistics	177
43.	Shared LPAR CP monitor data summary (percent, rounded to tenths)	180
44.	Dedicated LPAR CP monitor data summary (percent, rounded to tenths)	181

45.	VM/ESA Release 2 in Shared and Dedicated LPARs on 9121-480	183
46.	VM/ESA Release 2 RACF Release Comparison on a 9021-900.	187
47.	VM/ESA Release 2 RACF SECLABELS Comparison on a 9021-900. . . .	192
48.	VM/ESA Release 2 RACF 1.9.2 for VM Multiple Server Comparison on a 9021-900.	196
49.	Single VSE Guest, PACEX8 Workload on 9221-200	201
50.	Single VSE Guest, PACEX8 Workload on 9221-170	205
51.	Single VSE Guest, VSECICS Workload on 9221-170	207
52.	Two VSE Guests, PACEX8 Workload on 9221-200	211
53.	Two VSE Guests, PACEX6 Workload on 9221-200	215
54.	FS7B Workload Characteristics	224
55.	FS7B Workload Script Summary	226
56.	IOB Workload Script Summary	235
57.	CICS/VSE Transaction Characteristics	240

Notices

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 expressed 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 customer’s 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 environments do so at their own risk.

Performance data contained in this document were determined in various controlled laboratory environments and are for reference purposes only. **Customers should not adapt these performance numbers to their own environments as system performance standards.** The results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

This publication refers to some specific APAR numbers that have an effect on performance. The APAR numbers included in this report may have prerequisites, corequisites, or fixes in error (PEs). The information included in this report is not a replacement for normal service research.

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM licensed program in this publication is not intended to state or imply that only IBM’s program may be used. Any functionally equivalent program may be used instead.

Programming Interfaces

This publication is intended to help the customer understand the performance of VM/ESA Release 2 on various IBM processors. The information in this publication is not intended as the specification of any programming interfaces that are provided by VM/ESA Release 2. See the PUBLICATIONS section of the IBM Programming Announcement for VM/ESA Release 2 for more information about what publications are considered to be product documentation.

Trademarks

The following terms, denoted by an asterisk (*) in this publication, are trademarks of the IBM Corporation in the United States or other countries or both:

- ACF/VTAM
- CICS
- CICS/VSE
- DFSMS
- DisplayWrite
- ESA/370
- ESA/390
- IBM
- MVS/ESA
- OfficeVision
- OPC

PR/SM
RACF
S/390
VM/ESA
VM/XA
VSE/ESA
VTAM
3090

The following terms, denoted by a double asterisk (**) in this publication, are trademarks of other companies:

EXPLORE Legent Corporation

Acknowledgements

The following people contributed to this report:

Endicott Programming Laboratory

M. S. Bidwell	S. P. Lyons
W. J. Bitner	S. T. Marcotte
G. O. Blandy	V. E. Meredith
C. T. Bradley	D. L. Morrison
E. G. Branish	L. L. Quinn
P. A. Bronson	S. E. Shearer
R. A. Buck	A. M. Shepherd
P. A. Duncan	D. P. Spencer
W. G. Ernsberger	L. D. Stapleton
G. S. Gasper	C. R. Terry
L. H. Guyer	J. L. Thrall
W. J. Guzior	J. S. Tingley
G. A. Hine	A. M. Ward
R. R. Kaskan	R. C. Zinna
G. J. Kudamik	

Washington Systems Center

C. L. Morse

Mid-Hudson Valley Programming Lab

J. C. Dayka

B. C. Holmquist

Abstract

The *VM/ESA Release 2 Performance Report* summarizes the performance evaluation of VM/ESA Release 2. Measurements were obtained for the CMS-intensive, OfficeVision*, MVS guest, VSE guest, and VMSES/E environments on various ES/9000 processors.

This report provides performance and tuning information based on the results of the VM/ESA Release 2 performance tests conducted by the Endicott Programming Laboratory.

Discussion concentrates on the performance effects of migrating from a previous VM release (usually VM/ESA Release 1.1) to VM/ESA Release 2, exploiting the new functions provided in VM/ESA Release 2, and using certain tuning options. Some additional evaluations are also included.

Referenced Publications

VM/ESA Release 2 publications referred to in this report:

- *VM/ESA performance*, SC24-5642
- *SFS and CRR Planning, Administration, and Operation*, SC24-5649
- *Installation Guide*, SC24-5526-02

RACF publications referred to in this report:

- *Resource Access Control Facility (RACF) Program Directory for VM Installations*, GC28-1034
- *Resource Access Control Facility (RACF) Security Administrator's Guide*, SC28-1343.
- *System Programming Library: RACF*, SC28-1343

The following publication is unclassified and may be ordered by customers.

- *VM/ESA Release 1.1 Performance Report*, GG66-3236

The following publications are classified as IBM Internal Use Only. The information contained in these documents may be discussed with customers, but the documents may not be given to customers. They may ask their IBM representative for access to the information contained in these publications.

- *VM/ESA Release 1.0 Performance Report*, ZZ05-0469
- *OfficeVision/VM Performance comparison R1 SL101 to R2 SL200 on ES9021-500 and 4381¹*
- *OfficeVision/VM Performance comparison OV/VM R1 SL101 to OfficeVision/VM R2 SL200 and Current-OV/VM Direct Connect performance study on ES/9021-580²*

¹ IBM representatives can obtain this document by issuing the following command on their local office system:

```
TOOLS SENDTO DALVM41B TOOLS OFCTOOLS GET OVV12PC2 PACKAGE
```

² IBM representatives can obtain this document by issuing the following command on their local office system:

```
TOOLS SENDTO DALVM41B TOOLS OFCTOOLS GET OVPMULT PACKAGE
```

Summary of Key Findings

This report summarizes the performance evaluation of VM/ESA Release 2. Measurements were obtained for the CMS-intensive, OfficeVision*, VSE guest, MVS guest, and VMSES/E environments on various ES/9000 processors.

Performance Changes

VM/ESA Release 2 includes numerous performance enhancements. The twenty most significant ones are described in this report. Some of the most influential changes are in CP, but significant improvements have also been incorporated into CMS and SFS. The VM/ESA Release 2 performance enhancements particularly benefit the OfficeVision and CMS-intensive environments. They have little effect on the VSE and MVS guest environments.

The new functions in VM/ESA Release 2 and service updates since VM/ESA Release 1.1 have resulted in some unavoidable increases in storage requirements and processor usage. Although most VM computing environments are affected to some extent, these increases primarily affect the CMS-intensive environment. These increases reduce the net benefit of the performance enhancements.

Comparison to VM/ESA Release 1.1

Benchmark measurements showed the following performance results for VM/ESA Release 2 relative to VM/ESA Release 1.1:

CMS-intensive Internal throughput rate (ITR) improved 0.5% to 1.8%, while external response time ranged from equivalence to a 12% improvement. This environment was strongly influenced by both the performance improvements and the functional effects.

OfficeVision ITR improved 4.3% to 4.8% on the 9121-480 and 6.5% on the 9021-720. Response time improved 13% to 16% on the 9121-480 and 44% on the 9121-720. The improvements were greater on the 9121-720 due to performance enhancements that reduce master processor usage by certain functions that are heavily used by OfficeVision.

The performance of OV/VM 1.2 has also been significantly improved relative to OV/VM 1.1. Therefore, when a migration to VM/ESA Release 2 is combined with a migration to OV/VM 1.2, even larger performance benefits can result.

VSE Guest no significant change

MVS Guest no significant change

VMSES/E Overall performance improved. In addition, VMSES/E now automates certain checking that formerly had to be done manually. This can potentially result in significant reductions in the total time required to complete a given maintenance task.

Migration from older VM releases

Migrations from VM/XA and VM/ESA Release 1.0 (ESA Feature) should normally result in equivalent or improved performance.

Migrations from 370-based VM releases will see performance changes that vary depending upon the particular hardware, VM release, and workload involved. Migration to VM/ESA Release 2 can provide important constraint relief in areas such as real and virtual storage addressability, I/O, and multiprocessor processor support.

For migrations from nonconstrained systems, the measurement results, in general, show a response time improvement, similar or improved ITR when migrating on a multiprocessor, and a decrease in ITR when migrating on a uniprocessor. Special tuning, described in this report, is recommended for 9221 processors. A real storage increase is often advisable on the smaller processors in order to accommodate VM/ESA's larger real storage requirements.

Performance of New Functions

The results indicate that the FBA DASD support is equivalent to the existing CKD DASD support in terms of performance efficiency. Minidisk caching can be especially beneficial for FBA DASD but care must be taken (such as reblocking minidisks to 4KB) to ensure that MDC is actually used. Finally, there are special performance management considerations that apply to FBA DASD that are attached by integrated adapters.

The CP configurability function improves system availability and ease of administration without increasing system IPL time.

DIAGNOSE code X ϕ 250 ϕ extends the DASD I/O capabilities of the *BLOCKIO facility to VM Data Spaces. Measurements show that DIAGNOSE code X ϕ 250 ϕ uses 20% to 25% less processing time than *BLOCKIO.

Performance Management Improvements

A number of new records and fields were added to the CP monitor data. A new QUERY FILEPOOL REPORT command was provided to display SFS administration and performance data. All performance guidance in the VM manuals was consolidated into a new book, *VM/ESA Performance*. Finally, VMPRF and RTM/ESA were updated for VM/ESA Release 2.

Additional Information

This report contains useful information about the I/O assist feature. Topics include 1) when a guest is eligible for I/O assist, 2) what conditions cause an individual device or the entire guest to be taken out of I/O assist, and 3) how I/O assist affects the performance data provided by VM/ESA and its guests.

CMS-intensive measurements taken on the 9021-900 with 1, 2, 3, and 6 processors online show ITR scaling properly as the number of processors is increased.

LPAR measurements are presented to illustrate how running VM/ESA in an LPAR affects the CP monitor data.

Measurements were obtained for RACF/VM on VM/ESA Release 2. RACF/VM 1.9.2 showed improved performance relative to RACF/VM 1.9.1. With default access control and auditing, and with security label checking, RACF/VM 1.9.2 added 0.8% to overall processing requirements of a CMS-intensive workload relative to a VM/ESA Release 2 system without RACF. Measurements indicate that the multiple server support introduced by RACF/VM 1.9.2 can provide constraint relief and other benefits without increasing overall processing requirements.

VSE guest measurements were obtained on VM/ESA Release 2 to compare the performance of two guests, each with dedicated DASD, to two guests that share DASD through use of the VSE lock file. The results showed a 29% ITR decrease for preferred guests and equivalent ITR for V=V guests.

Changes That Affect Performance

This chapter contains descriptions of various changes to VM/ESA Release 2 that affect performance. Most of the changes are performance improvements and are listed under "Performance Improvements" on page 6. However, some have the potential to adversely affect performance. These are listed under "Performance Considerations" on page 18. The objectives of these two sections are as follows:

- Provide a comprehensive list of the significant performance changes.
- Allow installations to assess how their workloads may be affected by these changes.
- Describe new functions that applications could exploit to improve performance.

Throughout the rest of the report, various references are made to these changes. This further illustrates where these changes apply and how they may affect performance.

"Performance Management" on page 19 is the third section of this chapter. It discusses changes that affect VM/ESA* performance management.

Performance Improvements

This section describes performance improvements made to CP, CMS and SFS.

CP Improvements

This section describes performance improvements made in CP.

CP Free Storage

The code that handles short-term free storage requests was redesigned to reduce the processor requirements of this CP function. The changes to free storage for VM/ESA Release 2 significantly reduce the processor usage of HCPFRE and reduce the amount of storage HCPFRE displaces from the processor cache. HCPFRE's path lengths are cut by more than 1/2. The processing time in HCPFRE was reduced by over 1/3. The reason for the discrepancy in path length and processing time reduction is that the two operations that require the most processing time (TRACE and clearing the allocated free storage to zero) could not be removed.

For the minidisk CMS-intensive workload, this resulted in a 2% increase in internal throughput rate (ITR). The degree of benefit is proportional to the frequency of free storage requests. The number of free storage calls can be found in 2 sources: RTM's FRE_SUB field and in monitor domain 3, record 2 STORSP_PLSFSPCT field.

In addition, two new macros were created that normally obtain or release a block of free storage without calling HCPFRE. These macros are intended for use by performance-critical CP components. They are currently exploited by three other performance enhancements: HCPCFR (described in "Reduced Master Processor Usage" on page 9), IUCV (described in "IUCV Storage Management" on page 8), and fast path CCW translation (described in "Fast Path CCW Translation Extensions" on page 12).

CP Free Storage Implementation Details

In VM/ESA, free storage requests fall into one of two categories based on the length of time the storage is likely to be held. Blocks held for several minutes or more (long-term) take up the majority of free storage space but account for a small fraction of the free storage calls. Therefore, for this type of request, storage efficiency is far more important than the amount of time required to locate the storage.

The opposite is true for short-term storage. The request rates are tremendous, while the actual amount of storage held at any given time is far less than long-term storage. Therefore, path lengths to acquire such blocks must be kept to a minimum. For this reason, VM maintains several pools of available blocks where all the blocks in a given pool are of identical size. All short-term requests are satisfied by selecting a block from the pool that most closely matches (equal to or larger than) the requested size. When the targeted pool is empty, a frame (4096 bytes on a 4KB boundary) is carved into blocks equal to the pool size. One block is used to satisfy the request, and the remainder are placed in the pool to satisfy subsequent requests. Any fragment at the end of the frame (if the size did not divide evenly into 4096) is never used. When storage is returned, it is placed back in the pool from which it came.

As the requirements for a given size ebb, it is important to return the unused storage to the system so that it can be used for other purposes (such as free storage pools of different sizes, or to hold the contents of virtual storage). The

means by which this storage is returned is known as garbage collection. For VM/ESA, the garbage collector must be able to detect when an entire frame is freed so that it can be returned to the global available list. The garbage collector also takes actions that make it more likely that all the blocks within a frame are freed simultaneously.

The free storage enhancements in VM/ESA Release 2 improve the performance of the free storage pool manager, including garbage collection.

Prior to VM/ESA Release 2, steps were taken in the mainline of the free storage pool manager to assist garbage collection. All available blocks from the same frame were isolated by chaining them from the frame table entry (FRMTE) that represented that frame. In turn, the FRMTEs were chained together, and the anchor of that chain was in the subpool data area block (SUBBK). When all the blocks for a single frame were returned, the frame was time-stamped and removed from the queue. When the last block in a frame was given out (no blocks available), that FRMTE also was removed from the queue to avoid examining it each time. When the first block was subsequently returned, the FRMTE was placed back on the queue. This involved a good deal of bookkeeping and extra storage references.

With VM/ESA Release 2, the blocks are directly queued from the anchor, so that they can be located and removed with very few instructions. All garbage collection work is performed by a separate function that completes once every 30 seconds (as opposed to many thousand times a second).

Prior to VM/ESA Release 2, a rather inflexible algorithm was employed to convert the size of the request to a SUBBK address. The nature of the algorithm dictated what size pools were used, and many of the pools were actually synonyms for each other. For example, there were four sizes (58, 60, 62, and 64 double words) that yielded 8 blocks per frame. The difference was the size of the unusable fragment at the end of the frame.

With VM/ESA Release 2, the new HCPPOOLS macro, is used to generate the SUBBKs and to generate a simple table to convert size to SUBBK displacement. Now any size pool can be created and only the optimal sizes are selected.

HCPPOOLS also provides an assembly time conversion of size to SUBBK address. This is used by another new macro, HCPFGST, that acquires free storage very quickly for tasks that are extremely performance sensitive. All the code is generated inline, no TRACE entry is cut, and each block has only a header (no trailer).

These garbage collection techniques require no setup time during the mainline acquisition and release of storage. Additionally, much of the remaining statistics collection is now performed at garbage collection time rather than at each free storage call.

HCPSTP calls HCPFRFDX every 30 seconds to perform garbage collection and to accumulate statistics. HCPFRFDX processes each SUBBK by performing the following:

1. Any FRMTEs on the empty frame table (SUBMTFTE) queue are returned to the available list. These are frames that have remained empty since they were queued 30 seconds ago.

2. For each frame returned, the count of frames allocated to the pool is decremented.
3. All blocks anchored at SUBANCH are queued on their respective FRMTE, and a counter in the FRMTE is incremented for each block queued. When the first block is queued on a FRMTE, that FRMTE is in turn queued on a FRMTE anchor.
4. Empty FRMTEs (representing frames that currently have no allocated blocks) are queued on the SUBMTFTE queue.
5. The blocks from partially allocated FRMTEs are placed back on the SUBBLOCK anchor, but now they are isolated by frame.
6. If more than one frame is returned to the SUBBLOCK anchor, the last block is flagged to be returned to the collection queue.
7. The amount of storage currently in use is accumulated for the subpool and is calculated as

$$\text{SUBFRMCT} \times \text{SUBINUSE} - \text{unallocated_blocks} \times \text{SUBSIZE} \times 8$$

After processing each SUBBK, global statistics are stored in the appropriate counters.

When HCPFRE attempts to allocate from a subpool, and there are no blocks on the SUBBLOCK anchor, a check is made to see if a preformatted frame already exists on its *empty* queue.

IUCV Storage Management

In VM/ESA Release 1.1, one of the several changes to IUCV that led to improved performance was the more efficient storage management of control blocks MSGBK and IUSBK. Storage management was improved by making the control blocks semi-permanent and by exploiting stack management (see the *VM/ESA Release 1.1 Performance Report* for more details). This reduced the processor usage of HCPFRE.

However, this improvement caused short term garbage collection to be less efficient, because the control blocks were being held for several minutes (long-term). This caused storage fragmentation. In an attempt to improve short term garbage collection in VM/ESA Release 2, the control blocks were obtained from long-term storage. But the 9021-900 and 9021-720 experienced large-system effects in the long-term storage algorithm when obtaining long-term storage. This caused unacceptable response times and internal throughput rate.

APAR VM54161 to VM/ESA Release 2 corrects three problems by using inline macros to get the control blocks out of short term storage.

- Large-system effect

This is avoided because the control blocks come out of short-term storage. There is no net improvement over VM/ESA Release 1.1. However, avoiding the large-system effect is the primary reason for this APAR.

- HCPFRE calls

By using the inline macros, calls are avoided to HCPFRE. This is an improvement over VM/ESA Release 1.1.

- Short-term garbage collection

Storage fragmentation does not occur, because storage is no longer held long-term. This is an improvement over VM/ESA Release 1.1.

For more information on the inline macros, see “CP Free Storage” on page 6.

For the minidisk CMS-intensive workload, this improvement resulted in a 0.3% increase in internal throughput rate on the 9021-900, comparing VM/ESA Release 1.1 to VM/ESA Release 2 with this APAR.

Reduced Master Processor Usage

In VM/ESA, one method of serialization is to limit certain functions to run on only the master processor. Only one processor is designated as the master in VM/ESA. The IPLed processor is usually the master. However, if the IPLed processor has a special feature (vector or crypto) installed, then another processor without a special feature is selected as the master. The master can also be changed by disabling processors (using the VARY OFFLINE PROCESSOR command). There are three ways to determine which processor is the master. One is by issuing the CP QUERY PROCESSORS command, another is finding the first processor listed on RTM's CPU display screen, and yet another is locating the processor with the nonzero value in the *Moved to Master* field in VMPRF's PRF003 PROCESSORS_BY_TIME report.

Because there is only one master processor, it has the potential to be a bottleneck if the demand for master-only work is great. (OfficeVision environments were shown to be heavier in master-only work than most CMS-intensive or guest environments). As the number of processors in the complex increases, the potential for a master processor bottleneck increases. For example, a master processor bottleneck is more likely on a 6-way system than on a 2-way system.

Two changes were made in VM/ESA Release 2 that reduce the master processor requirements. Neither change involved moving work off the master, because moving work off the master would involve replacing the serialization methodology, a nontrivial task. Instead, these changes improve the efficiency of processing on the master. As a result, they also have a positive effect for environments that are not master processor constrained. The two items are as follows:

- HCPCNF improvement - DIAGNOSE code X'08' handling
- HCPCFR improvement - storage management for the console formatter

The HCPCNF module formats console output from DIAGNOSE code X'08' processing. This DIAGNOSE code issues CP commands. The modification applies to console output that goes to a buffer, not the user's screen. This would apply to programs that use EXECIO or CMS Pipelines to capture output from CP commands.

CP takes special precautions to prevent a single user from flooding the system if the user issues DIAGNOSE code X'08's that generate a lot of output. Prior to VM/ESA Release 2, each time this DIAGNOSE code was issued, CP opened a dispatch window from HCPCNF to allow another user to run. This kept the virtual machine that repeatedly issued the DIAGNOSE code from flooding the system with work.

In VM/ESA Release 2, the frequency of HCPCNF opening the dispatch window is changed. Instead of opening the dispatch window every time, it is opened

approximately every sixteenth time. This eliminates most of the overhead while maintaining the protection feature.

A 3.1% ITR improvement and a 16% improvement in external response time was observed on a VM/ESA Release 1.1 prototype for the OfficeVision workload on a 9021-720. A 0.2% improvement in ITR was projected for VM/ESA Release 2 for the CMS-intensive workload.

The HCPCFR module is not a master-only module. However, it is usually called by master-only modules, and therefore ends up completing on the master processor. This module handles console output formatting related to CP commands. In the OfficeVision environment, it was called frequently. The improvement in HCPCFR involved replacing costly CP free storage calls with the new inline macros. These inline macros are new in VM/ESA Release 2. See the "CP Free Storage" on page 6 section for additional information on them.

A 1.3% ITR improvement was observed on a prototype for the OfficeVision workload that was implemented on VM/ESA Release 1.1 running on a 9021-720. Because this enhancement works by replacing calls to the free storage manager with inline macros, and because free storage management itself was made more efficient in VM/ESA Release 2, a smaller improvement was seen on a VM/ESA Release 2 base. Minimal benefit was observed for the CMS-intensive workload on VM/ESA Release 2.

Improving the efficiency of master processor work reduces processor consumption. In addition, master processor constrained environments may also see significant improvement in response times. However, the great majority of systems are not master processor constrained. A system probably has a meaningful master processor constraint if the processor utilization of the master processor is much higher than the utilization of the other processors, and if the percent emulation time and idle time on the master processor is less than 5%. These values can be found in RTM on the CPU screen, in VMPRF's PRF003 PROCESSORS_BY_TIME report and VMPRF's PRF015 PROCESSOR_COMPLEX_BY_TIME report, and in VMPAF from VMPRF summary record type 02.

Monitor Enhancement

This enhancement applies to VM systems that run with user state sampling (MONITOR SAMPLE ENABLE USER). When user state sampling is enabled, HCPMOU is called once for every logged on or disconnected user at the end of each high frequency sample interval. In VM/ESA Release 2, HCPMOU no longer calls HCPFRE to obtain or release free storage for a module work area. Instead, HCPMOU uses leftover storage in the save block (used for linkage). Reducing the number of free storage calls reduces processor usage and improves response time.

The number of calls avoided is proportional to the number of users and the monitor sample rate. Each entry in Table 1 on page 11 column 3 (Reduction in FREE TOTL/SEC (Rate)) represents the reduction in free storage calls per second between VM/ESA Release 1.1 and VM/ESA Release 2. Notice that it decreases (6000, 1145, 198) as the number of users decreases (11420, 1870, 290). The sample rate for all measurements was 2 seconds.

Table 1 on page 11 also shows the percent reduction in free storage calls per command increasing as the system becomes smaller. This is due to the larger

base costs per command on smaller systems. There is a certain amount of overhead for monitor regardless of system throughput.

System	Users (Number)	Reduction in FREE TOTL/SEC (Rate)*	Reduction in FREE TOTL/CMD (Percent)	Source Data (Page)
9021-900	11420	6000	17.95	**
9121-480	1870	1145	20.91	42
9221-170	290	198	21.45	52
Note: * These numbers are based on the values for the FRE_TOTL field in RTM on the DISPLAY SYSTEM screen and STORSP_PLSFRECT and STORSP_PLSFSPCT of monitor domain 3 record 2.				
Note: ** This was a measured value on VM/ESA Release 2 prototype code.				

Table 1. Monitor Enhancement and Free Storage Calls

For one minidisk CMS-intensive measurement on a 9221-170 with 250 users and a 2-second high frequency sampling rate, internal throughput rate (ITR (H)) increased 0.5%.

DASD Slot Allocation

The DASD slot allocation algorithm, which gives out page and spool slots, was redesigned to take advantage of block paging and to reduce I/O seek times by increasing the number of contiguous DASD slots.

As in the past, a moving cursor approach is used. However, the new scheme allocates contiguous slots as much as possible by scanning ahead of the cursor for groups of contiguous available slots of appropriate sizes that reside on the same cylinder. This is accomplished in two ways:

1. These groups of contiguous available slots have a minimum size of two. Anything less than this minimum size is ignored. The one with the best fit is selected.
2. Allocation no longer switches volumes in the middle of allocating a block. Instead it may switch a little prematurely or a little late.

Environments that do significant paging to DASD should see an improvement in response time because of faster page fault resolution and a reduction in I/Os on the paging volumes. The benefits increase as the system ages. In the CMS-intensive workload, the average pages per page I/O increased 13% as a result of this improvement.

DASD Slot Allocation Implementation Details

In previous releases, when the slot allocation scheme received a request for a number of slots, it searched the page allocation blocks (PALBK) for an available slot. When it found one, that slot was allocated. It then continued searching from that point to fill the next slot in the request. As time went on, and more slots were allocated and deallocated, the chance of allocating contiguous slots decreased. Furthermore, allocation might have been stuck in a cylinder or area of DASD where there were very few slots available and may have been forced to allocate slots that were not contiguous. Meanwhile, there may have been large groups of contiguous available slots in other areas of the DASD that remained unused.

The new slot allocation scheme locates, ahead of time, groups of contiguous available slots for both paging and spooling space. These groups have a

minimum size and a maximum size. For paging, an appropriate size is any group larger than two. Anything less is considered a *crumb* and is not included in the buffer. When a request for slots arrives, the groups of slots are searched, and the one with the best fit is selected. Only those slots from the group that are needed to fill the request are used.

Slot allocation continues in this fashion. Each time a request comes in, the groups of slots are used. When all slots in a group are used, CP goes out to the bit map and searches for another group of slots to replace the one that was used. If only part of a group is used, the buffer entry is adjusted to indicate the new starting address and length, unless that remaining length is *crumb* size. In that case, the *crumb* is discarded and the entry is refilled.

Allocation no longer switches volumes in the middle of allocating a block (contiguous slots to hold related pages for block paging). Instead, it switches either a little prematurely or a little late. If the current request puts it over the limit (30), and at least half the limit was allocated, then it switches volumes. Otherwise, allocation switches after allocating the entire block on this volume.

This scheme is better because there is a greater chance of allocating contiguous slots. As time goes by and space becomes fragmented, the groups of available slots become smaller, but some contiguous slots are still allocated.

Fast Path CCW Translation Extensions

Fast path CCW translation was introduced as APAR VM51012 to VM/ESA Release 1.1. The APAR was not in the VM/ESA Release 1.1 software used in the CMS-intensive or OfficeVision measurements, but it was in the VM/ESA Release 1.1 software used in the VSE guest and MVS guest measurements.

Fast path CCW translation was extended in VM/ESA Release 2 to include support for FBA DASD. In addition, it now applies to more types of DASD channel programs. Fast path CCW translation also takes advantage of two new macros that normally obtain or release free storage without having to call HCPFRE.

Typical DASD I/O such as SSCH, DIAGNOSE code X'0A8', SIO, and SIOF will benefit from this. However, not all I/O is DASD I/O. For example, if RTM is reporting a high rate of DIAGNOSE code X'0A8's, many of these could be for unit record devices; therefore, the system does not benefit from this improvement.

A VSE guest environment running the PACEX8 workload (heavy DASD I/O) showed a 28% ITR improvement in the dedicated DASD environment and a 37% ITR improvement in the minidisk DASD environment. This is probably a best case and should be used as an upper bound. In a CMS-intensive environment, where most minidisk I/Os are DIAGNOSE X'0A4's and SFS I/Os are block I/O, ITR improved less than 1%. Fast path CCW translation does not apply to these cases because they take a different path through CP's I/O subsystem. The input to CP is not an actual channel program, but rather input from which CP builds the channel program.

For more information on this enhancement, see the *VM/ESA Release 1.1 Performance Report* and Washington System Center Flash 9220.

For VM/ESA Release 2, there are four new counters on the RTM/ESA SYSTEM screen:

FTR_DONE The number of virtual CCW programs successfully translated using fast path translation in CP.

FTR_ABOR The number of virtual channel program fast path translations that were terminated because they required full-function translation.

FTR_NE The number of virtual CCW programs presented to CP for fast path translation that were ineligible and, therefore, not processed through the fast path.

FTR_TOTL The number of virtual CCW programs presented to CP for fast path translation ($FTR_TOTL = FTR_DONE + FTR_ABOR + FTR_NE$).

DSPSLICE Calculation Change

CP initialization establishes a minor time slice for the dispatcher to use when running virtual machines. This code calculates the amount of time to complete about 100 000 selected instructions. The result is used as the dispatching minor time slice. Installations can change the size of the time slice with the CP SET SRM DSPSLICE command, but a reasonable default value is important. As the minor time slice becomes too small, the cost to run and to stop running users goes up. As the time slice is increased, the responsiveness of the system may decrease.

Before VM/ESA Release 1.1, the minimum value for the minor time slice was 1 millisecond. If the CP initialization code calculated a value less than 1 millisecond, CP used 1 millisecond instead. On some large S/390* processors, this led to minor time slices of 2 milliseconds or less. Such a small value caused a noticeable increase in the CP processor time needed to run and to stop running users. In VM/ESA Release 1.1, the minimum default time slice was changed to 5 milliseconds because of this problem.

Even after this change, the time slice value for some of the smaller S/390 processors remained too small. Many of the instructions used by the initialization code to determine the minor time slice were too short relative to a typical instruction mix and this discrepancy was particularly pronounced on smaller S/390 processors. Although the default minor time slice for these processors fell above the 5 milliseconds minimum introduced in VM/ESA Release 1.1, the value was still too low for efficient dispatching of work.

VM/ESA Release 2 changes the loop of instructions used to determine the minor time slice. The number of instructions completed is still about the same, but on average the instructions are longer. This loop produces a reasonable minor time slice value for the whole range of S/390 processors. The minimum default time slice is still 5 milliseconds, and the maximum remains at 100 milliseconds.

There are two principal benefits of this change:

1. Installations with small S/390 processors do not have to issue the CP SET SRM DSPSLICE command to achieve reasonable dispatching costs.
2. Installations with small S/390 processors that do not issue CP SET SRM DSPSLICE to increase the minor time slice may see lower CP processor costs. When the DSPSLICE is increased, workloads consisting of long-running transactions experience the largest dispatching efficiency benefits.

The DSPSLICE value is reported by monitor data in VMPRF's PRF072 SYSTEM_CONFIGURATION report. It is also reported by the QUERY SRM DSPSLICE command, with less precision.

Free Storage Limit Detection

In prior releases, a virtual machine could request excessive amounts of CP free storage, which could disable the system and be seen as performance degradation or an empty available list. This situation could be caused by:

- Repeatedly issuing CP commands that consume free storage
- Performing tasks in a disconnected machine that caused large amounts of console output to be routed to a secondary user
- Accounting or EREP records not being retrieved by the accounting or EREP virtual machine.

In VM/ESA Release 2, the free storage limit detection function tracks free storage so that a virtual machine cannot request excessive amounts. Three thresholds are calculated based on the size of the dynamic paging area. If the virtual machine exceeds the first threshold, it receives a warning message. If the virtual machine exceeds the second threshold, it is put into stopped state. If the virtual machine exceeds the third threshold, it is forced off the system. Messages are sent to the virtual machine and the operator user ID for each threshold crossed.

Four new CP commands were added to the system and documented in *CP Command and Utility Reference*. The QUERY STGLIMIT command (privilege classes: A, B, C, and E) checks the state of free storage limit detection for the system. The SET STGLIMIT command (privilege classes: A, B, and C) controls the state of free storage limit detection for the system. The QUERY STGEXEMP command (privilege classes: A, B, C, and E) checks the state of free storage limit detection for virtual machines. The SET STGEXEMP command (privilege classes: A, B, and C) controls the state of free storage limit detection for virtual machines.

Exempting a user ID ensures that the user ID is not subject to being stopped or forced off due to the amount of free storage it causes CP to consume. This is recommended for special purpose user IDs that are vital to the installation, user IDs running trusted code, and user IDs that should never be forced off the system. If no action is taken, free storage limits are enforced for all users.

CP Configurability and the Frame Table

For VM/ESA Release 2, the value that CP uses to determine the size of the frame table to build has changed. Prior to VM/ESA Release 2, CP used the RMSIZE value generated in HCPSYS. In VM/ESA Release 2, CP builds the frame table for the smaller value of either the actual real storage or RMSIZE of the SYSTORE macro in HCPSYS. With this improvement, the system may use less real storage.

The reported SYSGEN value for the QUERY FRAMES command was changed to reflect the smaller value of either actual real storage or RMSIZE.

The frame table is no longer built as part of the CP nucleus. It is built in the dynamic paging area. This reduces the size of the CP nucleus, which reduces the time to read the nucleus at system IPL. For more details see "CP Configurability" on page 138.

CMS Improvements

This section describes the performance improvements made in CMS.

Read-Only Minidisks

The format for accessed read-only CMS minidisk directories (FSTs) was improved. The old layout, called hyperblock format, alternated a 30-byte header with a block of FSTs, where the size of the block equaled the minidisk block size. The new layout has one header followed by all the FSTs. These FSTs are organized so that the block of FSTs in the hyperblock map are aligned by page boundary. This structure was in use, for some time, for the S and Y disks.

During FST lookup, the hyperblock map points to the page of desired FSTs. If the desired page doesn't exist, only one page of FSTs is referenced. Previously, the hyperblock of FSTs spanned pages, causing two pages to be referenced. For 4KB-blocked minidisks with an average-to-large number of files, this new layout typically eliminates one referenced nonshared page per read-only minidisk (other than S and Y) in the search order. The minidisk version of the workload (FS7B0) experienced this improvement because it has four such minidisks. The SFS version of the workload (FS7BMAX) experienced no benefits because these minidisks are replaced with SFS directories. For more details, see section "Minidisk to SFS" on page 168.

CMS Record Manager

Improvements were made to the CMS record manager (DMSRCM) to increase from 8 to 500 the maximum number of 4KB blocks that DMSRCM reads or writes with one DIAGNOSE code X'44'. This can result in reduced DIAGNOSE code X'44's, processor usage, I/O time, and device usage. This improvement applies to minidisk files used by applications (like XEDIT) that specify a large buffer. Most applications are unaffected by this improvement, and it does not apply to SFS. It has little effect on the CMS-intensive minidisk workload.

CMS Storage Management

Since CMS 5.5, page boundary alignment was forced for storage requests greater than or equal to 4KB. In VM/ESA Release 2, this is no longer done. As a result, applications that do a large number of requests greater than or equal to 4KB may see a reduction in virtual or real storage requirements and paging. However, if no such requests are made, it costs the CMS user 1 more referenced nonshared page.

For those applications that require page boundary alignment, CMSSTOR has an option to do this.

CMS Application Multitasking

CMS provides a multitasking environment for applications and servers. The multitasking facilities are available only at the application programming level and are provided as routines in a callable service library. The CMS user still runs one application at a time, but the application can split itself into multiple threads, and the multiple threads can be dispatched on multiple virtual processors. These multitasking facilities allow applications to harness the power of the underlying multiprocessor complex and to overlap operations, achieving high performance. Multiprocessor exploitation is supported in ESA, XA, and XC mode virtual machines.

NAMEFIND Improvement

In VM/ESA Release 2, improvements were made to the processing of the NAMES file. These improvements include caching the NAMES file in storage and opti-

mizing the search for tags. The degree of improvement depends on the size of the NAMES file and the number of names requested; the larger the NAMES file and number of requested names, the larger the improvement.

REXX Storage Management

REXX storage management APARs VM47302 and VM50916 for VM/ESA Release 1.1 can improve performance. These REXX storage management improvements are part of VM/ESA Release 2, so the VM/ESA Release 2 measurements in this report reflect the improvements. The APARs were not available for the VM/ESA Release 1.1 base level used in this report. Most of the performance change in REXX from VM/ESA Release 1.1 to VM/ESA Release 2 is due to these changes.

The improvements are to obtain or release storage in larger, page-aligned areas. By obtaining *larger* areas, there are fewer calls to CMS storage management. This reduces path length, but depending on the REXX program, some additional virtual storage may be required. One REXX instruction trace showed a 3% path length reduction when comparing VM/ESA Release 2 to VM/ESA Release 1.1. This reduction was primarily due to fewer CMS storage management calls.

By obtaining *page-aligned* areas, a header address can be calculated by clearing the low order bits rather than scanning a chain. Removing this scan can reduce path length and thrashing when REXX releases variable storage. One REXX trace of thousands of variables showed a 25% path length reduction when comparing VM/ESA Release 2 to VM/ESA Release 1.1. This reduction was primarily due to removing the scan.

The increased REXX usage in VM/ESA Release 2 (see "Productivity Aids" on page 18) resulted in a 50% increase in requests to REXX storage management. The combination of the REXX storage management improvements and the increased REXX usage yielded a 0.1% ITR improvement in the CMS-intensive minidisk workload.

REXX and Global Variables

The REXX VALUE function was extended in VM/ESA Release 2 to permit manipulation of global variables (as is done by the CMS GLOBALV command). Prior to this, manipulation of global variables from REXX required use of the GLOBALV CMS command. One measurement showed the VALUE function to be 27% faster when setting a global variable (GLOBALV SELECT *group* PUT *name1*), and 40% faster when getting a global variable. When working with a list of variables, the GLOBALV command has an advantage because it supports a list of variable names.

SFS Improvements

This section describes performance improvements made in SFS.

SFS Checkpoint

The checkpoint process consists of writing back to DASD all changed catalog blocks and unallocating all shadow blocks allocated since the last checkpoint. The checkpoint routine serializes the operation of the server before beginning this process. The checkpoint process was changed in two ways to reduce the amount of time that the server is serialized. First, the writeback of changed catalog buffers now exploits multiblock block I/O. Second, writeback of changed catalog buffers is now done before checkpoint serialization starts. This is called preflush. Any catalog buffer changes that may have occurred after the preflush

and before the checkpoint is serialized are still written back during the serialization period.

These improvements benefit users of read/write SFS filepools. The benefit in reduced checkpoint serialization time results in a lower average response time, a more consistent response time, and a reduced SFS server working set. In addition, the reduced checkpoint time allows a larger CATBUFFERS setting for better exploitation of large real memories.

From QUERY FILEPOOL STATUS, checkpoint duration can be calculated by dividing checkpoint time by checkpoints taken.

Catalog Insert Algorithm

The SFS catalog insert algorithm was changed to remember the location of the last inserted row and begin the search for space at that location. The benefits from this change are reduced processor usage and, in some cases, such as very large catalog spaces, reduced I/Os. The degree of benefit is proportional to the rate of inserts into the catalog and the catalog size. Typically, most catalog inserts come from GRANT AUTH and files that are larger than 32KB.

Log Manager

The log manager was enhanced to exploit multiblock block I/O. This enhancement eliminates log writes due to a full log buffer. Instead, the full log buffer is written at the next commit along with the buffer containing the commit log record. In addition, the log manager routines were combined and streamlined for reduced path length.

These changes apply to users of read/write file pools and result in fewer log I/Os and reduced processor usage. In the CMS-intensive SFS version of the workload, the estimated increase in ITR was 0.3%.

Default File Cache Size Increased

The default file cache size for SFS files (a CMS nucleus generation option) was increased from 12KB to 20KB. This has the advantage of reduced server communication and I/Os, but increases virtual storage requirements and may increase paging. The 20KB default should be a better trade-off for most installations.

Performance Considerations

This section describes changes in VM/ESA Release 2 that may adversely affect performance the most.

CMS Application Multitasking

The CMS application multitasking code is in a callable service library called VMMLIB. It is important to save VMMLIB in a shared segment whether or not CMS application multitasking is in use in order to keep virtual and real storage requirements to a minimum.

Virtual Storage Increase

CMS virtual storage requirements for VM/ESA Release 2 increased by about 25 pages. The majority of this increase is due to CMS application multitasking. These pages are referenced at CMS initialization or when CMS multitasking is used. For workloads that do not use multitasking, these pages are migrated out to DASD, causing an increase in DASD slot usage (reported by VMPRF) and an increase in paging.

In the CMS-intensive workload, one application could not run in a 2MB virtual machine because of this virtual storage increase. Consequently, the virtual machine sizes were increased to 3MB.

Productivity Aids

The performance of the CMS productivity aids changed. Most of the commands were rewritten in REXX, causing an increase in path length. The combined effect of the NAMEFIND improvement and the productivity aids consideration is negligible for the CMS-intensive workload.

Using Minidisk Cache with FBA

Because FBA devices are, in general, slower than CKD devices, minidisk caching is particularly beneficial for reducing the response times of FBA DASD. To use minidisk caching with FBA minidisks, it is important to format the CMS minidisks with a block size of 4KB. However, the CMS default block size for FBA devices is 1KB. Therefore, 4KB blocking must be specified explicitly. In addition to this requirement, when CMS minidisks are allocated on an FBA DASD volume, they must begin on a 4KB page boundary (that is, a block address that is divisible by 8).

Performance Management

This section discusses changes that affect VM/ESA performance management.

Accounting Data Changes

Values in accounting records may change in relationship to other changes in VM/ESA Release 2 that involve resource consumption. The degree of change in accounting data is workload dependent.

Virtual Machine Resource Usage Record

See "Performance Improvements" on page 6 and "Performance Considerations" on page 18 for details on changes that could affect this accounting record. The following list describes fields in the virtual machine resource usage accounting record (type 01) that may be affected by performance changes in VM/ESA Release 2. The columns where the field is located are shown in parentheses.

Milliseconds of processor time used (33-36) This is the total processor time charged to a user and includes both CP and emulation time. Some CMS changes resulted in increased processor usage while several CP improvements resulted in decreased processor usage. Some system overhead improvements do not show up in normal user type 01 records, but do affect the type 01 record for the system.

Milliseconds of Virtual processor time (37-40) This is the virtual time charged to a user and maps to emulation time in the various tables in this report. As mentioned earlier, changes in CMS resulted in increased processor usage. Therefore, the value for this field will probably increase for CMS users.

Total Page Reads (41-44) CMS storage requirements increased this release, so this field will increase for CMS users.

Total Page Writes (45-48) CMS storage requirements increased this release, so this field will increase for CMS users.

Requested Virtual nonspooled I/O Starts (49-52) This is a total count of requests. All requests may not complete. The value of this field will decrease in proportion to the benefit of the CMS record manager (DMSRCM) improvement. See "CMS Record Manager" on page 15 for details.

Completed Virtual nonspooled I/O Starts (73-76) This is a total count of completed requests. All requests may not complete. The value of this field will decrease in proportion to the benefit of the CMS record manager (DMSRCM) improvement.

Temporary Disk Space Record

The accounting record for temporary disk space (record type 03) was modified for FBA support. For FBA devices, the number of FBA blocks is given. For CKD and ECKD devices, cylinders remain as the given units. A new field was added that lists the size of the temporary disk space in pages for either type of device.

FBA DASD Considerations

For integrated adapter and control unit DASD, including all models of the 9332, 9335, and 9336 DASD, the hardware updates only the subchannel measurement block I/O request count. It does not update device connect time, device disconnect time, control unit queueing time (pending time), or device active-only time. As a result, performance products such as RTM/ESA and VMPRF show zero or inaccurate values for device service times, utilizations, or response times.

However, MLOAD statistics contain accurate queueing time for page and spool volumes. CP computes MLOAD statistics by recording the number of paging or spooling requests and the time required to process the requests. These statistics are kept for each page or spool volume. They are maintained by the paging subsystem in CP internal control blocks for the purpose of load balancing. This information is not useful if volumes contain page or spool and are mixed with nonpage or nonspool.

For volumes that do not contain page or spool space, the average queue length for the device can provide some information. This is reported in the VMPRF PRF012 DASD_BY_ACTIVITY report. The queue length value is from a field in the Real Device (RDEV) control block. However, the RDEV queue length field is not updated for page or spool I/Os. A high queue length could indicate a performance problem. A performance problem due to seek time for I/Os from a single server would not result in a high queue length. Therefore, a low queue length value does not always mean there are no performance problems with the given volume.

The SYSDASD screen was added to RTM. It reports the MLOAD statistics and roughly corresponds to the VMPRF PRF088 DASD_SYSTEM_AREAS report. VMPRF was updated to report that a device is FBA. This information is contained in several monitor records.

Monitor Enhancements

The MONITOR LIST1403 file (on MAINT's 194 minidisk) describes the content of each monitor record. Revision bars show which records and fields are new in VM/ESA Release 2. These enhancements are summarized below.

LPAR Monitor Enhancements

Prior to VM/ESA Release 2, monitor reported partition processor consumption, but did not report LPAR management time (the processor busy time that is not charged to any given partition). Monitor was enhanced to report LPAR management time. A flag in monitor domain 0 record 15 (D0/R15) tells whether the information is reported by the hardware or not, and a new record (D0/R17) reports the management time per physical processor.

There is a new flag in D0/R15 that reports whether capping is in effect. Capping limits the processor resources that a partition may consume.

ESCON Multiple Image Facility (EMIF) Monitor Enhancements

EMIF provides the capability to share a physical ESCON channel among multiple logical partitions running on the same processor. Prior to EMIF, processor channels were dedicated to individual logical partitions.

The new EMIF monitor data (found in D0/R18) reports busy time due to the partition in which VM/ESA is running.

Additional Monitor Enhancements

The following fields were added:

- User configuration fields were added: virtual machine size; account number; RACF* group name; count of reserved pages; logon time; QUICKDSP status; flag for dialed or SNA user and a V=V, V=F, or V=R flag. The records that were affected are D1/R15, D4/R1, D4/R3, and D4/R9.

- Relative and absolute SHARE settings by user ID were added. These affect records D2/R5, D4/R3, and D4/R9.
- User logoff data and resource consumption statistics from the user activity record (D4/R3) were added to the logoff event record (D4/R2). Now user resource consumption statistics are not lost for the time between the last sample interval and when the user logs off.
- The processor version (determines the type of processor) was added to D1/R5 (for example: 9021-580, not 9021).
- A count of dialed users and a count of SNA users was added to D0/R8.
- SET SRM DSPBUF settings were added to D1/R16 and D2/R7.

VMPRF Enhancements

VMPRF 1.2.1 with APAR VM53430 contains the following enhancements.

VMPRF Service

VMPRF APAR VM52375 replaced several values in VMPRF's PRF098 DASD_IO_ASSIST report with "*****" because of anomalies with the frequency of monitor sampling versus the frequency of CP updates to the counters. CP updates timers and counters when the state of a device changes, while monitor takes a snapshot of counters on a fixed interval.

FBA DASD

VMPRF was updated to report that a device is FBA.

Counters Added to the SFS Reports

PRECOORD	This field holds the number of precoordination requests (counter ID 140).
RENAMEUS	This field holds the number of rename user ID requests (counter ID 141).
FILEPOOLCTL	This field holds the number of filepool control backup requests (counter ID 142).
SFSLOGIO	This field holds the number of block I/O requests to write log blocks (counter ID 143).
SFSCPLOGIO	This field holds the number of I/O requests to write log blocks (counter ID 144).
ADDMDISK	This field holds the number of add minidisk requests (counter ID 145).
QDISABLE	This field holds the number of query disable requests (counter ID 146).
LOCKTIMEOUT	This field holds the number of locks denied due to timeout (counter ID 147).

SFS Field Changes

FPR	This field, which appears in all 3 SFS reports, includes counter IDs: 140, 141, 142, 145, and 146.
------------	--

The following fields are in VMPRF's PRF091 SFS_IO_BY_TIME report.

TOTAL	This field holds the number of total I/O requests. Counter ID 83 is replaced with counter ID 144.
LOG WRITE	This field holds the number of log write requests. Counter ID 83 is replaced with counter ID 144.
BLOCKS per BLOCK I/O	This field holds the number of I/O blocks. Counter ID 83 is replaced with counter ID 143.
BLOCKS per I/O	This field holds the number of I/O requests. Counter ID 83 is replaced with counter ID 144.
MEAN TIME, BLOCK I/O	This field holds the number of total block I/O requests. Counter ID 83 is replaced with counter ID 143.

RTM/ESA 1.5.2 Enhancements

RTM/ESA 1.5.2 contains the following service and enhancements.

RTM/ESA Service

- GC04859** *BLOCKIO was included in the ISEC counter of the SYSTEM screen. New counters IUCVTBIO, IUCV_FBIO, and IUCV_UBIO were added to the PRIVOPS screen.
- GC04769** ISEC on user screen was changed to use the same counters as used by monitor. DIAGNOSE code X'98' and DIAGNOSE code X'14' in ISEC were included on the SYSTEM screen.
- GC04908** The USER option was added to the SET PRINT command.

Counters Added to the RTM/ESA SYSTEM Screen

- RSCH** This field holds the number of real RESUME SUBCHANNEL instructions issued by CP.
- R_I/O** This field holds the number of real SSCH and RSCH instructions issued by CP. It is the sum of the SSCH and RSCH fields found on the SYSTEM screen.
- FTR_DONE** This field holds the number of virtual CCW programs successfully translated using fast path translation in CP.
- FTR_ABOR** This field holds the number of virtual CCW program fast path translations that were abandoned (requiring the full function translation).
- FTR_NE** This field holds the number of virtual CCW programs presented to CP for fast path translation that were ineligible; therefore, fast path translation was not attempted.
- FTR_TOTL** This field holds the number of virtual CCW programs presented to CP for fast path translation (FTR_TOTL = FTR_DONE + FTR_ABOR + FTR_NE).
- FST_DG44** This field holds the number of fast path simulations of DIAGNOSE code X'44' instructions.
- FST_SIGP** This field holds the number of fast path simulations of SIGP external call instructions.
- FST_EXTC** This field holds the number of fast path reflections of guest external call interrupts.

- FST_PARX** This field holds the number of fast path processing of partial execution intercepts.
- SIE_EX** This field holds the number of SIE instructions issued by CP.
- SIE_INT** This field holds the number of SIE interceptions.
- GUESTWT** This field holds the number of entries into guest enabled wait state.

Additional Enhancements Made to RTM/ESA

The SYSDASD screen was added. It includes minidisk-mapped devices (DASD that have a minidisk mapped to a VM Data Space; therefore, paging I/O is used instead of the usual I/O).

- The GENERAL, USER, and IDLE screens were added to the PRINT ALL command.
- An overall average line was added to SCLOG screen.
- A new CPSESV screen shows CP programming services usage, such as *BLOCKIO and *LOGREC.
- The PRIVSZ field was added to the VMX screen. This information gives the user's private address space size and includes the size of the VM Data Spaces along with the base virtual machine size.

SFS Enhancements

Nine new QUERY FILEPOOL commands were added to display the output of the QUERY FILEPOOL STATUS command in a more readable format. QUERY FILEPOOL REPORT displays all of the information that is contained in the QUERY FILEPOOL STATUS output plus additional information. The 8 other new commands each display a specified subset of that output.

The additional information provided by QUERY FILEPOOL REPORT includes the following:

- The date and time when the file pool server was last started
- The date and time that this query report was requested
- The date and time of the last control data backup
- The maximum number of IUCV and APPC connections that are allowed to the file pool server machine
- The number of addressable 4KB blocks in the file pool that are currently defined
- Total number of agents
- The number of storage groups and minidisks that are in use
- Block usage information, aggregated by storage group
- Virtual storage size of the file pool server machine
- The control minidisk size in 512-byte blocks
- Virtual addresses of the control minidisk and the log minidisks

SFS administrator authority is no longer required to issue the QUERY FILEPOOL STATUS command except when the CATALOG option is specified. The same applies to the new QUERY FILEPOOL REPORT command. Of the eight subset commands, only QUERY FILEPOOL CATALOG requires SFS administrator authority.

For more information on the QUERY FILEPOOL REPORT and its associated subset commands, see *SFS and CRR Planning, Administration, and Operation*.

Several new SFS counters were added in VM/ESA Release 2. These show up in QUERY FILEPOOL STATUS, QUERY FILEPOOL REPORT, and the domain XçAç monitor data. They are described in *VM/ESA Performance*, Appendix D.

VM/ESA Performance Publication

VM/ESA Performance is a new manual in the VM/ESA Publication Library. It contains performance information previously found in the following VM/ESA Release 1.1 books: *CP Planning and Administration*; *CP Programming Services*; *Features Summary*; *CMS Planning and Administration Guide*; and *Connectivity Planning, Administration, and Operation*. It discusses:

- Issues involved in managing, planning, measuring, and tuning a VM/ESA system for improved performance
- The various CMS and CP commands used in performance management
- A discussion of monitor system service (*MONITOR)
- CP monitor records
- SFS and CRR file pool server monitor records
- the Monwrite program

The order number is SC24-5642.

Measurement Information

This chapter discusses the types of processors used for measurements in the report, the level of software used, an explanation of the configuration details that are associated with each measurement, and the licensed programs and tools that were used in running and evaluating the performance measurements.

Hardware Used

The following processors and Licensed Internal Code levels were used for this report.

- 9021-900, Licensed Internal Code level 234784
This processor was used with 1, 2, 3, and 6 processors online and whatever processors were not being used were varied offline.
- 9021-720, Licensed Internal Code level 229914
- 9021-580, Licensed Internal Code level 229914
This processor was used for the 9021-340 measurement. To run as a 9021-340, 2 processors were varied offline.
- 9121-480, Licensed Internal Code level C22840
This processor was used for the 9121-480 and 9121-320 measurements. To run as a 9121-320, one processor was varied offline.
- 9221-200, Licensed Internal Code level 10D+
- 9221-170, Licensed Internal Code level 95D + MCL35
- 9221-120, Licensed Internal Code level 95D + MCL35

Software Used

Unless otherwise noted, a pre-general-availability level of VM/ESA Release 2 was used for the measurements in this report. All performance enhancements discussed in this report are part of the pre-general-availability level of code unless otherwise noted.

Other releases of VM were used in the report. VM/ESA Release 1.1 was at the GA+first-RSU level. The service that was part of VM/ESA Release 1.1 after the first RSU level and forwarded to the pre-general-availability code can account for some of the difference between VM/ESA Release 1.1 and VM/ESA Release 2. VM/ESA Release 1.0 (370 Feature) was at service level 104, VM/SP Release 5 CP was at service level 523, and VM/SP Release 5 CMS was at service level 533.

See the appropriate workload section in Appendix B, "Workloads" on page 223 for the other operating system and licensed programs' software levels.

Format Description

This part of the report contains a general explanation of the configuration details that are associated with each measurement.

For each group of measurements there are five sections:

1. **Workload:** It specifies the name of the workload associated with the measurement. For more detail on the workload, see Appendix B, "Workloads" on page 223.
2. **Hardware Configuration:** It summarizes the hardware configuration, and it contains the following descriptions:
 - **Processor model:** The model of the processor.
 - **Processors used:** The number of processors used.
 - **Storage:** The amount of real and expanded storage used on the processor.
 - **Real:** The amount of real storage used on the processor.
On the 9221, any real storage not defined for the specific measurement and not configured as expanded storage was attached to an idle user.
 - **Expanded:** The amount of expanded storage used on the processor.
The 9021 processors have dedicated expanded storage, while the other processor models configure part of real storage for use as expanded storage.
 - **Tape:** The type of tape drive and what the tape was used for.
 - **DASD:** The DASD configuration used during the measurement.
The table indicates the type of DASD used, type of control units that connect these volumes to the system, the number of paths between the DASD and the control unit, and the distribution of these volumes for PAGE, SPOOL, TDSK, USER, SERVER and SYSTEM. R or W next to the DASD counts means Read or Write caching enabled.
 - **Communications:** The type of control unit, number of communication control units, number of lines per control unit, and the line speed.
Because the 3745-410 control units used for this report are split in half and run in twin-dual mode, the number specified is the number of halves used for each of the processors (that is, the TPNS driver and the processor being measured). Each of these halves has a maximum of 50 lines available and can support a maximum of 3 000 users.
3. **Software Configuration:** This section contains pertinent software information.
 - **Driver:** The tool used to simulate users.
 - **Think time distribution:** The type of distribution used for the user think times.
 - Bactrian** This type of think time distribution represents a combination of both active and inactive user think times. The distribution includes long think times that occur when the user is not actively issuing commands. Actual user data were collected and used as input to the creation of the Bactrian distribution. This type of mechanism allows the transaction rate to vary depending on the command response times in the measurement.
 - IOB** This type of think time distribution represents the think time defined by the IBM Office Benchmark (IOB V2.1) workload. The think time includes an average 2-second delay between commands issued by TPNS, the built-in think times that are part of

the IOB scripts, and the IOB script scheduling algorithm. The average message rate per user stays constant across all of the measurements. See “IBM Office Benchmark (IOB V2.1)” on page 234 for more details.

- CMS block size: The block size of the CMS minidisks.
- Virtual Machines: The virtual machines used in the measurement.

For each virtual machine, the table indicates the following: name, number used, type, size and mode, relative share of the system resources scheduled, number of pages reserved, and any other options that were set.

4. Measurement Discussion: Contains an analysis of the performance data in the table and gives the overall performance findings.
5. Measurement Data: Contains the table of performance results. These data were obtained or derived from the tools listed in the section “Tools Description.”

There are several cases where the same information is reported from two sources because the sources calculate the value in a slightly different manner. As an example, consider the external throughput rate measures, ETR (T) and ETR, that are based on the command rate calculated by TPNS and RTM respectively. TPNS is external to the system and can directly count the command rate as it runs the commands in the scripts. Because RTM uses information that is internal to the system, it has to make assumptions about when transactions begin and end. This can make the counts reported by RTM vary in meaning from run to run and vary from the values reported by TPNS. As a result, the analysis of the data relies more on the TPNS command rate. Furthermore, some values in the table (like TOT INT ADJ) were normalized to the TPNS command rate in an effort to get the most accurate performance measures possible.

Performance terms listed in the tables and discussed in this part of the document are defined in the glossary.

Tools Description

A variety of licensed programs were used to evaluate the performance measurements. The licensed programs used in the measurements are listed below.

EXPLORE**	monitors and reports performance data for VSE systems.
RMF	Resource Measurement Facility monitors and reports MVS performance.
RTM	Real Time Monitor records and reports performance data for VM systems. The appropriate release level of VM/RTM was used for VM/SP Release 5, VM/SP HPO Release 5.0, and VM/ESA Release 1.0 (370 Feature) systems, and the appropriate release level of RTM VM/ESA was used for VM/ESA Release 1.1 and VM/ESA Release 2 systems.
TPNS	Teleprocessing Network Simulator is a terminal and network simulation tool.
VMMAP	VM Monitor Analysis Program is the VM/370 Monitor reduction program.

VMPAF	VM Performance Analysis Facility data are not in the measurement tables but the data are used for performance analysis of VM systems.
VMPRF	VM Performance Reporting Facility is the VM Monitor reduction program.

When comparing the VMMAP data for the VM/ESA Release 1.0 (370 Feature) measurements to the corresponding VMPRF and RTM data for the VM/ESA Release 1.1 and VM/ESA Release 2 measurements, be aware that many of these measures do not have precisely the same meaning.

The tools internal to IBM that were used in conjunction with the licensed programs in the measurements are listed below:

FSTTAPE	Reduces hardware monitor data for the 9121 and the 9021-720.
Hardware Monitor	Collects branch, event, and timing data.
INSREDH2	Reduces hardware monitor data for the 9021-900.
MONFAST	Collects branch, event, and timing data on a 9221 in addition to reducing the data it collects.
REDFP	Consolidates the QUERY FILEPOOL STATUS data.
TPNS Reduction Program	Reduces the TPNS log data to provide performance, load, and response time information.

Because each workload used a different subset of licensed programs and tools, the licensed programs and tools are listed at the bottom of each measurement table following the measurement discussion.

Migration to VM/ESA Release 2

This chapter explores the performance effects of migrating to VM/ESA Release 2. The following environments were measured: CMS-intensive, OfficeVision, MVS guest, VSE guest, and VMSES/E.

DASD fast write was used for the OfficeVision environment, because OfficeVision's workload had heavy write DASD I/O activity and a constraint on the amount of available DASD.

Each workload experiences some run variability. The degree of variability depends on the workload and configuration. A study of run variability was performed on the 9121-480 with the CMS-intensive minidisk version of the workload. The results are summarized in "Measurement Variability" on page 173.

CMS-Intensive Migration from VM/ESA Release 1.1

The new functions in VM/ESA Release 2 caused a growth in storage and path length. This had some adverse impact on processor usage. However, the increased processor usage was more than offset by improvements made in VM/ESA Release 2, most notably in CP. Consequently, internal throughput improved. The following improvements influenced the CMS-intensive environments the most:

- CP Free Storage

This improvement had the largest impact, decreasing CP processor usage. All measurements in this section show about the same degree of benefit from this improvement.

- IUCV Storage Management

This improvement is provided by an APAR (VM54161) to VM/ESA Release 2. It decreases CP processor usage. Some measurements were made on the larger systems such as the 9021-900 with this improvement.

This APAR corrects three problems by using inline macros to get the control blocks out of short-term storage.

- Large-systems effect

This is avoided because the control blocks come out of short-term storage. There is no net improvement over VM/ESA Release 1.1. However, avoiding the large-systems effect is the primary reason for this APAR.

- HCPFRE calls

By using the inline macros, calls to HCPFRE are avoided. This is an improvement over VM/ESA Release 1.1.

- Storage fragmentation with short-term garbage collection

Storage fragmentation does not occur because storage is no longer held long-term. This is an improvement over VM/ESA Release 1.1.

- Monitor Enhancement

This improvement decreases CP processor usage. The systems that benefit the most collect high-frequency user data. All measurements in this section benefit from this improvement.

- DASD Slot Allocation

This improvement decreases page I/Os. It mostly benefits storage-constrained environments with paging to DASD.

- DSPSLICE Calculation Change

This improvement decreases CP processor usage. Smaller systems such as the 9221-170 get the most benefit.

- Read-Only Minidisks

This improvement decreases a user's working set size and paging rate. The minidisk measurements indicated a savings of one page per read-only minidisk (other than S and Y) in the search order. Because the SFS measurements have no read-only minidisks (other than S and Y) in their search order, they received no benefit.

The following improvements apply only to the SFS measurements:

- SFS Checkpoint
This improvement decreases response times and makes them more consistent.
- Improved Catalog Insert Algorithm
This improvement reduces processor usage. If the catalog is very large, it may also reduce I/Os.
- Log Manager
This improvement reduces log I/Os and processor usage.
- Default File Cache Size Increased
This improvement reduces file I/Os and processor usage. The default was changed from 12KB to 20KB, which should be a better trade-off for most installations.

For more information on these and other performance-related improvements in VM/ESA Release 2, see “Changes That Affect Performance” on page 5.

The improved ITRs for VM/ESA Release 2 are due to decreased processor usage which was brought about by these changes. The internal and external throughputs for these measurements are shown in the following 2 graphs.




Figure 1. Internal Throughput for the CMS-Intensive Workload. CMS-intensive environments compared between VM/ESA Release 1.1 and VM/ESA Release 2




Figure 2. External Response Time for the CMS-Intensive Workload. CMS-intensive environments compared between VM/ESA Release 1.1 and VM/ESA Release 2

9021-900 / Minidisk

Workload: FS7B0R

Hardware Configuration

Processor model: 9021-900
 Processors used: 6
 Storage
 Real: 1024MB
 Expanded: 4096MB (400MB retained for MDC)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3380-D	3880-03	4	18	30	30			
3380-E	3880-03	4	2		4			
3380-A	3880-03	4		2	6			
3380-D	3880-23	4				16 R		4 R
3380-K	3990-02	4	10	10		60		
3390-2	3990-02	4				40 R		

Note: *R* or *W* next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3745-410	4	44	56Kb
3088-02	1	NA	4.5MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVED	Other Options
VTAMXA	1	VTAM	64MB/XA	10000	350	QUICKDSP ON
VSCSn	6	VSCS	64MB/XA	10000	OFF	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
Users	11420	User	3MB/XC	100	OFF	

Measurement Discussion

Table 2 on page 36 summarizes the results of the migration measurements on the 9021-900 comparing VM/ESA Release 1.1 and VM/ESA Release 2 for the minidisk-only CMS-intensive environment. This environment showed improvement in response time and processor capacity.

Overall processor requirements (PBT/CMD (H)) decreased by 1.8%, resulting in a 1.8% increase in internal throughput rate (ITR (H)). The decrease in processor time was from the combined effects of a 6.2% decrease in the CP processor time (CP/CMD (H)) and a 1.1% increase in emulation processor time (EMUL/CMD (H)). The drop in CP processor time was due mainly to the following performance improvements, which offset the increase in processor usage due to new function in VM/ESA Release 2:

- CP Free Storage
- IUCV Storage Management

The VM/ESA Release 2 measurement was run with this APAR (VM54161). It is very important for the performance of this environment because it removes the large-system effect.

- Monitor Enhancement
- Reduced Master Processor Usage

See “Changes That Affect Performance” on page 5 for more detail on these and other performance improvements for VM/ESA Release 2.

Real storage requirements increased for VM/ESA Release 2. This can be seen in the overall increase in paging (PAGE/CMD plus XSTOR/CMD). Some of the growth in real storage requirements was in CP, as evidenced by the growth in the storage required for CP control blocks (FREEPGS) and the decrease in the number of pageable pages available (PGBLPGS). User working set size (WKSET (V)) grew by 5%, reflecting the combined effects of growth in CMS storage usage and the benefits of the read-only minidisks performance enhancement. See “Virtual Storage Increase” on page 18 and “Read-Only Minidisks” on page 15 for more information.

External response time (AVG LAST (T)) decreased by 12.4% and internal response time (TOT INT ADJ) decreased by 15.5%. Response times are affected by changes to processor usage, real storage requirements, and I/Os. In this nonconstrained environment, the benefits of reduced processor usage greatly outweighed the effects of the increased real storage requirements.

The CP nucleus size (NUCLEUS SIZE (V)) decreased for this measurement pair. This was primarily due to the frame table (FRMTES) being moved from the CP nucleus to the dynamic paging area (DPA). See “CP Configurability and the Frame Table” on page 14 for details.

There was an increase in the runnable portion of the nucleus, but for this measurement pair the increase was more than offset by the removal of the frame table. See “9221-170 / Minidisk” on page 52 for a measurement pair where the frame table savings was not enough to offset the runnable portion increase. With the smaller memory, the frame table was smaller.

The minidisk cache size was fixed at 400MB (RETAIN XSTORE MDC 400M 400M). This was done because preliminary VM/ESA Release 1.1 base measurements showed significant adverse performance effects when the cache size was allowed to vary dynamically. Fixing the cache size, then, allowed the remaining VM/ESA Release 1.1 to VM/ESA Release 2 performance effects to be observed without being obscured by this problem.

An additional VM/ESA Release 2 measurement (not shown) was done where the arbiter was given complete flexibility to adjust the minidisk cache size. The

results were equivalent to the VM/ESA Release 2 400MB minidisk cache run shown in this section. This indicates that the problem previously observed on VM/ESA Release 1.1 has been corrected in VM/ESA Release 2. However, the specific change that implements this has not been identified.

Because FS7B0R's minidisk cache requirements are quite uniform over time, this workload does not require the arbiter's ability to dynamically adjust the minidisk cache size. However, for production systems, the workload is usually much more variable. Therefore, for production workloads, it is usually best to give the arbiter flexibility (possibly within a specified size range) to adjust the size of the minidisk cache so as to balance I/O and paging performance.

The number of VSCS machines that are required for good response time is proportional to the number of users the system supports. For these 11420 user measurements, 6 VSCS machines were used (1903 users per VSCS) based on extrapolations from data published in *VM/ESA Release 1.0 Performance Report*. With careful tuning, it may be possible to support 11420 users with 5 external VSCSs (2284 users per VSCS) and obtain response times similar to the VM/ESA Release 2 measurement in Table 2 on page 36. However, 4 external VSCSs are apparently insufficient. Earlier 11420-user measurements of VM/ESA Release 2 (not shown) with 4 external VSCS machines (2855 users per VSCS) showed excessive IUCV queuing and poor response times.

In summary, the performance of VM/ESA Release 2 showed improvement over VM/ESA Release 1.1 for the minidisk-only CMS-intensive environment. This environment showed lower response times and increased processor capacity.

RELEASE RUN ID	VM/ESA 1.1 264RB424	VM/ESA 2 265RB426
Environment		
REAL STORAGE	1024MB	1024MB
EXP. STORAGE	4096MB	4096MB
USERS	11420	11420
VTAMs	1	1
VSCSs	6	6
PROCESSORS	6	6
Response Time		
TRIV INT	0.096	0.080
NONTRIV INT	0.304	0.266
TOT INT	0.223	0.198
TOT INT ADJ	0.227	0.192
AVG FIRST (T)	0.374	0.319
AVG LAST (T)	0.463	0.406
Throughput		
AVG THINK (T)	25.95	25.98
ETR	410.10	394.82
ETR (T)	402.00	406.79
ETR RATIO	1.020	0.971
ITR (H)	446.86	454.96
ITR	76.16	73.73
EMUL ITR	121.16	113.52
ITRR (H)	1.000	1.018
ITRR	1.000	0.968
Proc. Usage		
PBT/CMD (H)	13.427	13.188
PBT/CMD	13.408	13.176
CP/CMD (H)	5.359	5.028
CP/CMD	4.975	4.622
EMUL/CMD (H)	8.068	8.160
EMUL/CMD	8.433	8.555
Processor Util.		
TOTAL (H)	539.77	536.47
TOTAL	539.00	536.00
UTIL/PROC (H)	89.96	89.41
UTIL/PROC	89.83	89.33
TOTAL EMUL (H)	324.34	331.93
TOTAL EMUL	339.00	348.00
MASTER TOTAL (H)	95.76	94.21
MASTER TOTAL	96.00	94.00
MASTER EMUL (H)	29.96	32.02
MASTER EMUL	32.00	34.00
TVR(H)	1.66	1.62
TVR	1.59	1.54
Storage		
NUCLEUS SIZE (V)	6192KB	2376KB
TRACE TABLE (V)	1200KB	1200KB
WKSET (V)	58	61
PGBLPGS	212 k	209 k
PGBLPGS/USER	18.6	18.3
FREEPGS	27912	30652
FREE UTIL	0.95	0.90
SHRPGS	2759	2609
Paging		
READS/SEC	471	666
WRITES/SEC	251	446
PAGE/CMD	1.796	2.734
PAGE IO RATE (V)	90.500	142.600
PAGE IO/CMD (V)	0.225	0.351
XSTOR IN/SEC	2596	3021
XSTOR OUT/SEC	2977	3670
XSTOR/CMD	13.863	16.448
FAST CLR/CMD	5.868	6.409

RELEASE RUN ID	VM/ESA 1.1 264RB424	VM/ESA 2 265RB426
Environment		
REAL STORAGE	1024MB	1024MB
EXP. STORAGE	4096MB	4096MB
USERS	11420	11420
VTAMs	1	1
VSCSs	6	6
PROCESSORS	6	6
Queues		
DISPATCH LIST	202.30	174.63
ELIGIBLE LIST	0.00	0.00
I/O		
VIO RATE	3481	3506
VIO/CMD	8.659	8.619
RIO RATE (V)	976	1016
RIO/CMD (V)	2.428	2.498
MDC READS	2264	2313
MDC WRITES	1089	1097
MDC MODS	898	905
MDC HIT RATIO	0.92	0.92
PRIVOPs		
PRIVOP/CMD	15.820	15.886
DIAG/CMD	22.544	24.092
DIAG 08/CMD	0.739	0.767
DIAG 14/CMD	0.017	0.022
DIAG 58/CMD	1.234	1.219
DIAG 98/CMD	0.261	0.258
DIAG A4/CMD	3.973	3.938
DIAG A8/CMD	1.883	1.900
DIAG 214/CMD	11.716	13.260
SIE/CMD	47.263	46.707
SIE INTCPT/CMD	31.194	29.425
FREE TOTL/CMD	84.576	86.039
VTAM Machines		
WKSET (V)	12992	12025
TOT CPU/CMD (V)	2.2887	2.2249
CP CPU/CMD (V)	1.2182	1.1384
VIRT CPU/CMD (V)	1.0704	1.0865
DIAG 98/CMD (V)	0.262	0.260
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 2. Minidisk-Only CMS-Intensive Migration from VM/ESA Release 1.1 to VM/ESA Release 2 on the 9021-900.

9021-900 / Minidisk / Storage-Constrained**Workload: FS7B0R****Hardware Configuration**

Processor model: 9021-900
 Processors used: 6
 Storage
 Real: 768MB
 Expanded: 2048MB (200MB retained for MDC)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3380-D	3880-03	4	18	30	30			
3380-E	3880-03	4	2		4			
3380-A	3880-03	4	30	2	6			
3380-D	3880-23	4				16R		4R
3380-K	3990-02	4	10	10		60		
3390-2	3990-03	4				40R		

Note: *R* or *W* next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3745-410	4	44	56Kb
3088-02	1	NA	4.5MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVED	Other Options
VTAMXA	1	VTAM	64MB/XA	10000	350	QUICKDSP ON
VSCSn	6	VSCS	64MB/XA	10000	OFF	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
Users	10800	User	3MB/XC	100	OFF	

Measurement Discussion

This section summarizes the results of the measurements on the 9021-900 comparing VM/ESA Release 1.1 and VM/ESA Release 2 in a storage-constrained environment. Also, some comparisons are made between these storage constrained measurements and the corresponding nonconstrained measurements

described in "9021-900 / Minidisk" on page 33. This environment showed improvement in internal throughput and response time.

The 9021-900 storage-constrained environment experienced a 1.6% increase in internal throughput (ITR (H)). This is from the combined effects of a 5.8% decrease in CP time per command (CP/CMD (H)) and a 1.5% increase in emulation time per command (EMUL/CMD(H)).

External response time (AVG LAST (T)) decreased by 5.4% and internal response time (TOT INT ADJ) decreased by 8.7%. A contributing factor to the response time improvement was the decrease in page I/O requests per command (PAGE IO/CMD(V)). This decreased, even though the paging rate increased (PAGE/CMD), because of the slot allocation enhancement (see "DASD Slot Allocation" on page 11).

When comparing these constrained measurements to the corresponding nonconstrained measurements, the paging rate was 4-5 times higher in the constrained measurements. The number of available pages (PGBLPGS/USER) was about 26% lower in the constrained measurements. This demonstrates the trend that paging increases as storage decreases.

RELEASE RUN ID	VM/ESA 1.1 264RA802	VM/ESA 2 265RA800
Environment		
REAL STORAGE	768MB	768MB
EXP. STORAGE	2048MB	2048MB
USERS	10800	10800
VTAMs	1	1
VSCSs	6	6
PROCESSORS	6	6
Response Time		
TRIV INT	0.123	0.111
NONTRIV INT	0.376	0.352
TOT INT	0.277	0.261
TOT INT ADJ	0.288	0.263
AVG FIRST (T)	0.446	0.414
AVG LAST (T)	0.546	0.517
Throughput		
AVG THINK (T)	25.90	25.86
ETR	395.40	386.89
ETR (T)	380.53	384.33
ETR RATIO	1.039	1.007
ITR (H)	420.45	427.27
ITR	72.98	71.92
EMUL ITR	121.87	116.14
ITRR (H)	1.000	1.016
ITRR	1.000	0.986
Proc. Usage		
PBT/CMD (H)	14.270	14.043
PBT/CMD	14.243	14.024
CP/CMD (H)	6.072	5.718
CP/CMD	5.703	5.334
EMUL/CMD (H)	8.198	8.325
EMUL/CMD	8.541	8.690
Processor Util.		
TOTAL (H)	543.02	539.71
TOTAL	542.00	539.00
UTIL/PROC (H)	90.50	89.95
UTIL/PROC	90.33	89.83
TOTAL EMUL (H)	311.97	319.95
TOTAL EMUL	325.00	334.00
MASTER TOTAL (H)	96.10	94.61
MASTER TOTAL	96.00	94.00
MASTER EMUL (H)	24.72	26.85
MASTER EMUL	26.00	29.00
TVR(H)	1.74	1.69
TVR	1.67	1.61
Storage		
NUCLEUS SIZE (V)	5172KB	2376KB
TRACE TABLE (V)	1200KB	1200KB
WKSET (V)	67	68
PGBLPGS	149 k	147 k
PGBLPGS/USER	13.8	13.6
FREEPGS	26932	28786
FREE UTIL	0.96	0.90
SHRPGS	2555	2515
Paging		
READS/SEC	2370	2540
WRITES/SEC	1814	1980
PAGE/CMD	10.995	11.761
PAGE IO RATE (V)	958.400	861.700
PAGE IO/CMD (V)	2.519	2.242
XSTOR IN/SEC	960	1318
XSTOR OUT/SEC	3141	3753
XSTOR/CMD	10.777	13.194
FAST CLR/CMD	5.784	6.297

RELEASE RUN ID	VM/ESA 1.1 264RA802	VM/ESA 2 265RA800
Environment		
REAL STORAGE	768MB	768MB
EXP. STORAGE	2048MB	2048MB
USERS	10800	10800
VTAMs	1	1
VSCSs	6	6
PROCESSORS	6	6
Queues		
DISPATCH LIST	236.98	235.24
ELIGIBLE LIST	0.00	0.02
I/O		
VIO RATE	3302	3313
VIO/CMD	8.677	8.620
RIO RATE (V)	1807	1698
RIO/CMD (V)	4.749	4.418
MDC READS	2130	2169
MDC WRITES	1036	1039
MDC MODS	855	854
MDC HIT RATIO	0.92	0.92
PRIVOPs		
PRIVOP/CMD	15.940	16.070
DIAG/CMD	22.521	24.088
DIAG 08/CMD	0.738	0.757
DIAG 14/CMD	0.018	0.023
DIAG 58/CMD	1.259	1.220
DIAG 98/CMD	0.265	0.268
DIAG A4/CMD	3.958	3.921
DIAG A8/CMD	1.887	1.897
DIAG 214/CMD	11.671	13.210
SIE/CMD	47.303	46.834
SIE INTCPT/CMD	32.166	30.442
FREE TOTL/CMD	86.722	88.465
VTAM Machines		
WKSET (V)	2531	2309
TOT CPU/CMD (V)	2.3667	2.3216
CP CPU/CMD (V)	1.2505	1.1847
VIRT CPU/CMD (V)	1.1162	1.1370
DIAG 98/CMD (V)	0.267	0.270
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 3. CMS-Intensive Minidisk Only Storage Constrained on the 9021-900.

9121-480 / Minidisk

Workload: FS7B0R

Hardware Configuration

Processor model: 9121-480
 Processors used: 2
 Storage
 Real: 192MB
 Expanded: 64MB (All used for MDC)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3390-2	3990-2	4	8	4	4			
3390-2	3990-3	4				10 R		
3390-2	3990-3	2	4	2	2		2 R	

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3088-02	1	NA	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVED	Other Options
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	400	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
Users	1870	User	3MB/XC	100	OFF	

Measurement Discussion

The following table summarizes the results of the migration measurements on the 9121-480 comparing VM/ESA Release 1.1 and VM/ESA Release 2 for the mini-disk workload. These results were similar to the 9021-900 minidisk migration results (see "9021-900 / Minidisk" on page 33 for further details).

The 9121-480 experienced a 0.5% increase in internal throughput rate (ITR (H)). This is from the combined effects of a 4.4% decrease in CP time per command (CP/CMD (H)) and a 1.7% increase in emulation time per command (EMUL/CMD

(H)). External response times (AVG LAST (T)) are equivalent within run variability. Refer to "Measurement Variability" on page 173 for further details.

There was a similar increase in storage requirements for VM/ESA Release 2 on the 9121-480 as observed on the 9021-900. There was a 4 page (5.5%) increase in average working set size (WKSET (V)) and an increase in the storage required by CP, shown as an increase in the FREEPGS and a decrease in the PGBLPGS.

RELEASE RUN ID	VM/ESA 1.1 L24R1876	VM/ESA 2 L25R187F
Environment		
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
USERS	1870	1870
VTAMs	1	1
VSCSs	0	0
PROCESSORS	2	2
Response Time		
TRIV INT	0.123	0.127
NONTRIV INT	0.523	0.539
TOT INT	0.384	0.395
TOT INT ADJ	0.351	0.357
AVG FIRST (T)	0.334	0.330
AVG LAST (T)	0.492	0.498
Throughput		
AVG THINK (T)	25.54	25.43
ETR	61.26	60.60
ETR (T)	66.98	66.99
ETR RATIO	0.915	0.905
ITR (H)	73.51	73.86
ITR	33.65	33.44
EMUL ITR	50.63	49.18
ITRR (H)	1.000	1.005
ITRR	1.000	0.994
Proc. Usage		
PBT/CMD (H)	27.209	27.078
PBT/CMD	27.172	27.021
CP/CMD (H)	9.794	9.364
CP/CMD	9.107	8.659
EMUL/CMD (H)	17.415	17.714
EMUL/CMD	18.065	18.362
Processor Util.		
TOTAL (H)	182.24	181.38
TOTAL	182.00	181.00
UTIL/PROC (H)	91.12	90.69
UTIL/PROC	91.00	90.50
TOTAL EMUL (H)	116.64	118.66
TOTAL EMUL	121.00	123.00
MASTER TOTAL (H)	90.97	90.48
MASTER TOTAL	91.00	90.00
MASTER EMUL (H)	52.74	53.86
MASTER EMUL	55.00	56.00
TVR(H)	1.56	1.53
TVR	1.50	1.47
Storage		
NUCLEUS SIZE (V)	2792KB	2304KB
TRACE TABLE (V)	400KB	400KB
WKSET (V)	73	77
PGBLPGS	39883	39665
PGBLPGS/USER	21.3	21.2
FREEPGS	4741	4876
FREE UTIL	0.97	0.95
SHRPGS	1109	1177

RELEASE RUN ID	VM/ESA 1.1 L24R1876	VM/ESA 2 L25R187F
Environment		
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
USERS	1870	1870
VTAMs	1	1
VSCSs	0	0
PROCESSORS	2	2
Paging		
READS/SEC	562	586
WRITES/SEC	370	389
PAGE/CMD	13.915	14.555
PAGE IO RATE (V)	184.200	177.300
PAGE IO/CMD (V)	2.750	2.647
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	5.748	6.404
Queues		
DISPATCH LIST	39.12	45.11
ELIGIBLE LIST	0.00	0.00
I/O		
VIO RATE	624	616
VIO/CMD	9.316	9.196
RIO RATE (V)	391	376
RIO/CMD (V)	5.838	5.613
MDC READS	363	367
MDC WRITES	182	181
MDC MODS	150	151
MDC HIT RATIO	0.91	0.92
PRIVOPs		
PRIVOP/CMD	17.610	17.507
DIAG/CMD	24.604	26.135
DIAG 08/CMD	0.746	0.776
DIAG 14/CMD	0.015	0.015
DIAG 58/CMD	1.254	1.224
DIAG 98/CMD	0.851	0.866
DIAG A4/CMD	3.912	3.852
DIAG A8/CMD	1.971	1.926
DIAG 214/CMD	11.601	13.197
SIE/CMD	52.374	51.877
SIE INTCPT/CMD	36.662	35.276
FREE TOTL/CMD	83.115	64.626
VTAM Machines		
WKSET (V)	421	415
TOT CPU/CMD (V)	3.7739	3.7736
CP CPU/CMD (V)	1.7086	1.6836
VIRT CPU/CMD (V)	2.0653	2.0900
DIAG 98/CMD (V)	0.852	0.872
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 4. Minidisk Migration from VM/ESA Release 1.1 to VM/ESA Release 2 on the 9121-480.

9121-480 / Minidisk / Storage-Constrained

Workload: FS7B0R

Hardware Configuration

Processor model: 9121-480
 Processors used: 2
 Storage
 Real: 128MB
 Expanded: None
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3390-2	3990-2	4	10	2	4			
3390-2	3990-3	4				10 R		
3390-2	3990-3	2	4	2	2			2 R

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3088-02	1	NA	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVED	Other Options
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	480	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	100	OFF	
Users	1820	User	3MB/XC	100	OFF	

Measurement Discussion

This section summarizes the results of the minidisk-only CMS-intensive measurements on the 9121-480 comparing VM/ESA Release 1.1 and VM/ESA Release 2 in a storage-constrained environment.

The 9121-480 storage-constrained environment experienced a 0.9% increase in the internal throughput rate (ITR (H)). This was due to the combined effects of a 4.5% decrease in CP time per command (CP/CMD (H)) and a 1.1% increase in emulation time per command (EMUL/CMD (H)).

External response time (AVG LAST (T)) decreased by 0.07 seconds (12.2%). A contributing factor to the response time improvement was the decrease in page I/O requests per command (PAGE IO/CMD (V)). This decreased, even though the paging rate (PAGE/CMD) increased, because of the DASD slot allocation enhancement (see "DASD Slot Allocation" on page 11).

When comparing these constrained measurements to the corresponding measurements done with 192MB real storage and 64MB expanded storage (see "9121-480 / Minidisk" on page 42):

- The paging per command (PAGE/CMD) was higher in the constrained measurements. This demonstrates the trend that paging increases as storage decreases.
- The constrained measurements did not have minidisk caching because there was no expanded storage. This caused more real I/O per command (RIO/CMD (V)).

RELEASE RUN ID	VM/ESA 1.1 L24R1821	VM/ESA 2 L25R1823
Environment		
REAL STORAGE	128MB	128MB
EXP. STORAGE	0	0
USERS	1820	1820
VTAMs	1	1
VSCSs	0	0
PROCESSORS	2	2
Response Time		
TRIV INT	0.139	0.140
NONTRIV INT	0.612	0.559
TOT INT	0.454	0.424
TOT INT ADJ	0.409	0.372
AVG FIRST (T)	0.344	0.297
AVG LAST (T)	0.541	0.474
Throughput		
AVG THINK (T)	25.48	25.37
ETR	58.61	57.57
ETR (T)	65.01	65.66
ETR RATIO	0.902	0.877
ITR (H)	72.54	73.20
ITR	32.70	32.12
EMUL ITR	49.32	47.40
ITRR (H)	1.000	1.009
ITRR	1.000	0.982
Proc. Usage		
PBT/CMD (H)	27.571	27.321
PBT/CMD	27.535	27.262
CP/CMD (H)	9.988	9.538
CP/CMD	9.230	8.833
EMUL/CMD (H)	17.583	17.783
EMUL/CMD	18.306	18.428
Processor Util.		
TOTAL (H)	179.23	179.39
TOTAL	179.00	179.00
UTIL/PROC (H)	89.62	89.69
UTIL/PROC	89.50	89.50
TOTAL EMUL (H)	114.30	116.76
TOTAL EMUL	119.00	121.00
MASTER TOTAL (H)	89.66	89.83
MASTER TOTAL	90.00	90.00
MASTER EMUL (H)	52.12	53.47
MASTER EMUL	54.00	56.00
TVR(H)	1.57	1.54
TVR	1.50	1.48

RELEASE RUN ID	VM/ESA 1.1 L24R1821	VM/ESA 2 L25R1823
Environment		
REAL STORAGE	128MB	128MB
EXP. STORAGE	0	0
USERS	1820	1820
VTAMs	1	1
VSCSs	0	0
PROCESSORS	2	2
Storage		
NUCLEUS SIZE (V)	2792KB	2304KB
TRACE TABLE (V)	400KB	400KB
WKSET (V)	70	73
PGBLPGS	23779	23600
PGBLPGS/USER	13.1	13.0
FREEPGS	4640	4787
FREE UTIL	0.94	0.96
SHRPGS	1037	1102
Paging		
READS/SEC	583	635
WRITES/SEC	396	429
PAGE/CMD	15.060	16.205
PAGE IO RATE (V)	204.100	196.600
PAGE IO/CMD (V)	3.140	2.994
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	5.830	6.381
Queues		
DISPATCH LIST	49.35	48.16
ELIGIBLE LIST	0.00	0.00
I/O		
VIO RATE	609	604
VIO/CMD	9.368	9.199
RIO RATE (V)	593	581
RIO/CMD (V)	9.122	8.849
MDC READS	0	0
MDC WRITES	0	0
MDC MODS	0	0
MDC HIT RATIO	0.00	0.00
PRIVOPs		
PRIVOP/CMD	17.623	17.325
DIAG/CMD	24.449	25.809
DIAG 08/CMD	0.754	0.746
DIAG 14/CMD	0.015	0.015
DIAG 58/CMD	1.246	1.249
DIAG 98/CMD	1.015	1.097
DIAG A4/CMD	3.861	3.808
DIAG A8/CMD	1.907	1.767
DIAG 214/CMD	11.537	13.022
SIE/CMD	55.548	55.057
SIE INTCPT/CMD	37.217	35.787
FREE TOTL/CMD	88.760	70.515
VTAM Machines		
WKSET (V)	485	486
TOT CPU/CMD (V)	3.9607	3.9937
CP CPU/CMD (V)	1.7858	1.7684
VIRT CPU/CMD (V)	2.1748	2.2253
DIAG 98/CMD (V)	1.014	1.100
Note:	T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM	

Table 5. Minidisk Storage Constrained from VM/ESA Release 1.1 to VM/ESA Release 2 on the 9121-480

9121-480 / SFS**Workload: FS7BMAXR****Hardware Configuration**

Processor model: 9121-480
 Processors used: 2
 Storage
 Real: 192MB
 Expanded: 64MB (all retained for MDC)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3390-2	3990-2	4	8	4	4			
3390-2	3990-3	4					10 R	
3390-2	3990-3	2	4 R	2 R	2 R			2 R

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3088-02	1	NA	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVE	Other Options
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	400	QUICKDSP ON
ROSERV1	1	SFS	32MB/XC	1500	OFF	QUICKDSP ON
RWSERV1	1	SFS	32MB/XC	1500	1000	QUICKDSP ON
RWSERV2	1	SFS	32MB/XC	1500	1000	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	100	OFF	
Users	1560	User	3MB/XC	100	OFF	

Measurement Discussion

The FS7BMAXR workload has all user data in SFS files. The read/write files are in FILECONTROL directories, and the read-only files are in DIRCONTROL directories that are in VM Data Spaces. The read-only files reside in a separate read-only file pool. The system files are on S and Y read-only minidisks.

The performance of VM/ESA Release 2 showed improvement over that of VM/ESA Release 1.1. The SFS workload showed lower response times, increased internal throughput rate, but an increase in paging.

There was a 10% improvement in both internal (TOT INT ADJ) and external (AVG LAST(T)) response times. In addition to the CP enhancements cited in previous measurements, this improvement was attributed to the performance enhancements made to the SFS checkpoint process and increasing the SFS file cache from 12KB to 20KB. Refer to "CMS File Cache for SFS" on page 146 for the performance results of this enhancement. The SFS checkpoint enhancements reduced the time to perform a checkpoint from 6.3 seconds to 3.8 seconds (40%). This is shown in the QUERY FILEPOOL STATUS data in Appendix A, "SFS Counter Data" on page 219.

There was a 0.5% improvement in internal throughput rate (ITR(H)) from VM/ESA Release 1.1 to VM/ESA Release 2. This improvement was due to previously cited CP enhancements, path length improvements in the SFS log manager, and an improved catalog insert algorithm that reduces the time spent looking for space to insert catalog rows for large files such as compiler work files.

Due to the increase in storage requirements for VM/ESA Release 2, there was a 19% increase in paging rate (PAGE/CMD). This increase was significantly more than the 5% observed in the minidisk-only environment (refer to "9121-480 / Minidisk" on page 42). In the minidisk-only environment, the format of the read-only minidisk directories was improved to reduce virtual storage and page references (see "Read-Only Minidisks" on page 15 for more details). The enhancement does not apply to the measured SFS environment because the only read-only minidisks are the S and Y minidisks, which are already in the improved format. This was verified by measuring a 35% SFS environment (not included in this report) that has the read-only data on minidisks. A lower increase in paging similar to the minidisk-only environment was observed.

RELEASE RUN ID	VM/ESA 1.1 L24M1566	VM/ESA 2 L25M1564
Environment		
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
USERS	1560	1560
VTAMs	1	1
VSCSs	0	0
PROCESSORS	2	2
Response Time		
TRIV INT	0.120	0.128
NONTRIV INT	0.677	0.608
TOT INT	0.498	0.454
TOT INT ADJ	0.442	0.394
AVG FIRST (T)	0.305	0.286
AVG LAST (T)	0.542	0.486
Throughput		
AVG THINK (T)	25.62	25.58
ETR	49.34	49.08
ETR (T)	55.60	56.49
ETR RATIO	0.887	0.869
ITR (H)	61.61	61.92
ITR	27.35	26.94
EMUL ITR	41.87	40.20
ITRR (H)	1.000	1.005
ITRR	1.000	0.985
Proc. Usage		
PBT/CMD (H)	32.464	32.301
PBT/CMD	32.375	32.217
CP/CMD (H)	11.963	11.397
CP/CMD	11.151	10.621
EMUL/CMD (H)	20.501	20.903
EMUL/CMD	21.224	21.596
Processor Util.		
TOTAL (H)	180.49	182.47
TOTAL	180.00	182.00
UTIL/PROC (H)	90.25	91.24
UTIL/PROC	90.00	91.00
TOTAL EMUL (H)	113.98	118.08
TOTAL EMUL	118.00	122.00
MASTER TOTAL (H)	90.19	91.04
MASTER TOTAL	90.00	91.00
MASTER EMUL (H)	52.73	54.50
MASTER EMUL	55.00	56.00
TVR(H)	1.58	1.55
TVR	1.53	1.49
Storage		
NUCLEUS SIZE (V)	2792KB	2304KB
TRACE TABLE (V)	400KB	400KB
WKSET (V)	64	76
PGBLPGS	40824	40629
PGBLPGS/USER	26.2	26.0
FREEPGS	4171	4276
FREE UTIL	0.91	0.90
SHRPGS	1311	1377
Paging		
READS/SEC	412	492
WRITES/SEC	278	340
PAGE/CMD	12.410	14.728
PAGE IO RATE (V)	132.700	140.700
PAGE IO/CMD (V)	2.387	2.491
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	5.863	6.709

RELEASE RUN ID	VM/ESA 1.1 L24M1566	VM/ESA 2 L25M1564
Environment		
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
USERS	1560	1560
VTAMs	1	1
VSCSs	0	0
PROCESSORS	2	2
Queues		
DISPATCH LIST	42.26	41.09
ELIGIBLE LIST	0.00	0.00
I/O		
VIO RATE	534	530
VIO/CMD	9.605	9.382
RIO RATE (V)	345	344
RIO/CMD (V)	6.205	6.089
MDC READS	259	266
MDC WRITES	114	118
MDC MODS	75	78
MDC HIT RATIO	0.86	0.86
PRIVOPs		
PRIVOP/CMD	29.533	24.666
DIAG/CMD	21.925	24.027
DIAG 08/CMD	0.791	0.814
DIAG 14/CMD	0.018	0.018
DIAG 58/CMD	1.259	1.239
DIAG 98/CMD	1.223	1.257
DIAG A4/CMD	2.068	2.053
DIAG A8/CMD	1.727	1.699
DIAG 214/CMD	10.828	12.763
SIE/CMD	64.912	60.806
SIE INTCPT/CMD	48.035	43.172
FREE TOTL/CMD	95.632	73.286
VTAM Machines		
WKSET (V)	421	407
TOT CPU/CMD (V)	4.1568	4.0714
CP CPU/CMD (V)	1.8586	1.7899
VIRT CPU/CMD (V)	2.2982	2.2816
DIAG 98/CMD (V)	1.223	1.259
SFS Servers		
WKSET (V)	1573	1611
TOT CPU/CMD (V)	4.1568	3.6977
CP CPU/CMD (V)	2.0384	1.7899
VIRT CPU/CMD (V)	2.1184	1.9079
FP REQ/CMD(Q)	1.392	1.270
IO/CMD (Q)	2.000	1.804
IO TIME/CMD (Q)	0.033	0.032
SFS TIME/CMD (Q)	0.083	0.054
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Query Filepool Counters, Unmarked=RTM		

Table 6. SFS migration from VM/ESA Release 1.1 to VM/ESA Release 2 on the 9121-480

9221-170 / Minidisk

Workload: FS7B0R

Hardware Configuration

Processor model: 9221-170
 Processors used: 1
 Storage
 Real: 48MB
 Expanded: 16MB (all retained for MDC)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3390-2	3990-3	1				8	R	
3390-2	3990-3	1	3	2	2			2 R
3390-2	3990-2	1	5	5	5			

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3088-02	1	NA	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB
 DSPSLICE: Set to 15 for VM/ESA Release 1.1, default for VM/ESA Release 2
 Preloaded segments: FORTRAN and DCF

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVED	Other Options
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	OFF	QUICKDSP ON IPOLL ON DELAY 0.2
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	100	OFF	
Users	290	User	3MB/XC	100	OFF	

Measurement Discussion

This section summarizes the results of the minidisk-only CMS-intensive measurements on the 9221-170 comparing VM/ESA Release 1.1 and VM/ESA Release 2.

Notice in the above configuration sections that the VM/ESA Release 1.1 and VM/ESA Release 2 measurements followed the 9221 tuning recommended in the

VM/ESA Release 1.1 Performance Report. This tuning concentrates on reducing the number of I/Os and SIE instructions to significantly improve internal throughput rate and response time for the 9221. In particular:

- Configure a portion of real storage as expanded storage and use it exclusively for minidisk cache. The result is a replacement of DASD I/Os with less processor-intensive minidisk cache reads.
- Set the dispatching minor time slice to three times the default for the VM/ESA Release 1.1 measurement. For VM/ESA Release 1.1, DSPSLICE was set to 15 milliseconds. An improvement was made in VM/ESA Release 2 to the way the default DSPSLICE is calculated, so it was not necessary to override the default for VM/ESA Release 2. For the VM/ESA Release 2 measurement, DSPSLICE defaulted to 16.158 ms. For more details, see “DSPSLICE Calculation Change” on page 13.
- Set the DELAY operand in the VTAM CTCA channel-attachment major node to 0.2. This reduces VTAM I/Os and the associated SIE interceptions.
- Set IPOLL ON for VTAM. This reduces the number of IUCV instructions and the associated SIE interceptions.
- Preload shared segments (FORTRAN and DCF) from an idle user. This prevents the shared segments’ page frames from becoming invalid when not in use to avoid page reads when the next user accesses the shared segments.

The 9221-170 experienced a 0.9% increase in internal throughput rate (ITR (H)) with improvements in CP time per command (CP/CMD (H)) offsetting a minor increase in emulation time per command (EMUL/CMD (H)). External response times (AVG LAST (T)) were similar, and internal response times (TOT INT ADJ) were also similar.

An additional measurement (not shown) was made on VM/ESA Release 1.1 that was like H14R0292, except DSPSLICE defaulted to 5 ms. The internal throughput rate (ITR (H)) was 3.7% higher when DSPSLICE was set to 15 ms. The larger DSPSLICE default in VM/ESA Release 2 may be helpful if the previous system defaulted to the smaller value.

The CP nucleus size (NUCLEUS SIZE (V)) increased. This was primarily due to an increase in the portion of the nucleus containing code. Systems with large amounts of real memory show a decrease in the nucleus size (see “9021-900 / Minidisk” on page 33 for more information).

RELEASE RUN ID	VM/ESA 1.1 H14R0292	VM/ESA 2 H15R0294
Environment		
REAL STORAGE	48MB	48MB
EXP. STORAGE	16MB	16MB
USERS	290	290
VTAMs	1	1
VSCSs	0	0
PROCESSORS	1	1
Response Time		
TRIV INT	0.164	0.153
NONTRIV INT	1.210	1.202
TOT INT	0.892	0.879
TOT INT ADJ	0.761	0.752
AVG FIRST (T)	0.316	0.310
AVG LAST (T)	0.818	0.824
Throughput		
AVG THINK (T)	28.23	28.32
ETR	8.61	8.65
ETR (T)	10.09	10.11
ETR RATIO	0.853	0.856
ITR (H)	11.44	11.55
ITR	9.77	9.89
EMUL ITR	15.12	14.99
ITRR (H)	1.000	1.009
ITRR	1.000	1.012
Proc. Usage		
PBT/CMD (H)	87.411	86.611
PBT/CMD	87.180	87.052
CP/CMD (H)	36.665	35.303
CP/CMD	30.711	29.677
EMUL/CMD (H)	50.745	51.307
EMUL/CMD	56.469	57.375
Processor Util.		
TOTAL (H)	88.23	87.55
TOTAL	88.00	88.00
TOTAL EMUL (H)	51.22	51.87
TOTAL EMUL	57.00	58.00
TVR(H)	1.72	1.69
TVR	1.54	1.52
Storage		
NUCLEUS SIZE (V)	2360KB	2448KB
TRACE TABLE (V)	200KB	200KB
WKSET (V)	89	80
PGBLPGS	9535	9454
PGBLPGS/USER	32.9	32.6
FREEPGS	900	896
FREE UTIL	0.89	0.90
SHRPGS	927	980
Paging		
READS/SEC	67	71
WRITES/SEC	53	56
PAGE/CMD	11.888	12.563
PAGE IO RATE (V)	27.700	25.600
PAGE IO/CMD (V)	2.744	2.532
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	5.944	6.430
Queues		
DISPATCH LIST	12.54	11.85
ELIGIBLE LIST	0.00	0.00

RELEASE RUN ID	VM/ESA 1.1 H14R0292	VM/ESA 2 H15R0294
Environment		
REAL STORAGE	48MB	48MB
EXP. STORAGE	16MB	16MB
USERS	290	290
VTAMs	1	1
VSCSs	0	0
PROCESSORS	1	1
I/O		
VIO RATE	107	109
VIO/CMD	10.600	10.783
RIO RATE (V)	72	69
RIO/CMD (V)	7.133	6.826
MDC READS	57	58
MDC WRITES	27	27
MDC MODS	22	22
MDC HIT RATIO	0.91	0.91
PRIVOPs		
PRIVOP/CMD	14.744	14.594
DIAG/CMD	26.449	28.448
DIAG 08/CMD	0.693	0.692
DIAG 14/CMD	0.000	0.000
DIAG 58/CMD	1.189	1.187
DIAG 98/CMD	2.180	2.176
DIAG A4/CMD	3.864	3.858
DIAG A8/CMD	1.882	2.077
DIAG 214/CMD	11.789	13.256
SIE/CMD	56.865	55.990
SIE INTCPT/CMD	40.943	39.753
FREE TOTL/CMD	91.936	72.214
VTAM Machines		
WKSET (V)	175	177
TOT CPU/CMD (V)	17.6122	17.2630
CP CPU/CMD (V)	7.8429	7.5404
VIRT CPU/CMD (V)	9.7693	9.7225
DIAG 98/CMD (V)	2.222	2.255
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 7. Minidisk migration from VM/ESA Release 1.1 to VM/ESA Release 2 on the 9221-170

CMS-Intensive Migration from VM/SP HPO Release 5

This section summarizes the results from a set of measurements that were obtained on the 9121-320 and 9121-480 processors to assess the performance effects of migrating a CMS-intensive workload from VM/SP HPO Release 5.0 to VM/ESA Release 2.

The 9121-320 results show an improvement in response time but a decrease in internal throughput rate (ITR) and users supported at 90% processor utilization. The response time improvement reflects the benefits of CMS write-behind, XA I/O architecture, and minidisk caching. These changes have an especially significant positive effect on response time because they reduce I/O wait time, which is often the largest component of system response time. The ITR decrease is primarily because:

VM/SP HPO runs especially well on uniprocessors. This is because 1) HPO's multiprocessor support can be excluded from CP and 2) VM/ESA's more efficient multiprocessor support does not apply.

The extensive new functions provided by VM/SP Release 6 CMS (especially 31-bit support and the Shared File System) caused a significant increase in processor usage by CMS guests even when these new functions are not used. Although improvements were made in subsequent releases, some of this impact remains.

VM/ESA Release 2 was measured in two ways on the 9121-320:

1. with the 256MB total storage configured the same as HPO (64MB central, 192MB expanded)
2. with the 256MB configured more optimally for VM/ESA (192MB central, 64MB expanded)

For the first case (storage configured the same), external response time improved by 33%, while ITR decreased by 9.1%. For the second case (storage configured optimally), external response time improved by 45%, while ITR decreased by 5.5%.

The 9121-480 results show VM/SP HPO Release 5.0 reaching a 16MB-line constraint at between 1200 and 1300 users, while VM/ESA Release 2 was able to support 1870 users at about 90% processor utilization with 0.5 second average external response time. ITR improved by 6.1% relative to the 1200 user VM/SP HPO Release 5.0 base measurement. This reflects the benefits of VM/ESA's more efficient multiprocessor support, which more than compensated for the increase in CMS processing requirements.

9121-320 / Minidisk

Workload: FS7B0R

Hardware Configuration

Processor model: 9121-320³
 Processors used: 1
 Storage: 256MB real plus expanded (see discussion and table)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SWAP	SPOOL	- Number of Volumes -		System
						TDSK	User	
3390-2	3990-3	2						2 R
3390-2	3990-2	2					10	
3390-2	3990-2	2	2	1	1	3		
3390-2	3990-3	2	2	1	1	3		
3380K	3880-3	2	2	2	1			1
3380K	3880-3	1						1
3380K	3880-3	2	2		1		2	1
3380K	3880-3	1						1
3380K	3880-3	1						2

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Note: The 4 DASD that were used as SWAP volumes in the HPO measurement were used as additional paging volumes in the VM/ESA measurements.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3088-02	1	NA	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

³ See "Hardware Used" on page 25 for an explanation of how this processor model was defined.

Virtual Machines:

<i>Virtual Machine</i>	<i>Number</i>	<i>Type</i>	<i>Machine Size/Mode</i>	<i>PRIORITY/ RELSHARE</i>	<i>RESERVED</i>	<i>Other Options</i>
For VM/SP HPO 5:						
VTAM	1	VTAM/VSCS	16MB/370	1	OFF	QDROP OFF FAVOR ON
SMART	1	VM/RTM	6MB/370	1	OFF	QDROP OFF FAVOR ON
Users	1030	User	3MB/370	64	OFF	
For VM/ESA Release 2:						
VTAM	1	VTAM/VSCS	16MB/370	10000	OFF	QUICKDSP ON
SMART	1	RTM/ESA	6MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
Users	950	User	3MB/370	100	OFF	

Measurement Discussion

The following three CMS-intensive measurements were obtained on a 256MB 9121-320 to examine the performance effects of migrating from VM/SP HPO Release 5.0 to VM/ESA Release 2 on a 9121 uniprocessor:

1. VM/SP HPO Release 5.0 — 64MB central, 192MB expanded, 1030 users

The 256MB total memory was configured optimally for HPO, which does not support more than 64MB of central storage.

Of the 192MB expanded storage, 180MB were allocated for swap paging and 12MB were allocated for demand paging — a 15:1 ratio. This ratio was selected based on an earlier study on the 3090-200J, where experiments showed this to be a good ratio for the FS7B0R workload. (Most production HPO systems are run with a ratio in the range of 1:1 to 3:1.)

2. VM/ESA Release 2 — 64MB central, 192MB expanded, 950 users

The 256MB was configured the same as for the HPO base run. The 192MB of expanded storage was used for paging and minidisk caching (MDC).

3. VM/ESA Release 2 — 192MB central, 64MB expanded, 950 users

The 256MB is configured optimally (or nearly so) for VM/ESA. As much storage as possible was configured as central storage, while leaving some expanded storage for exclusive use by MDC.

For each of these measurements, the number of users was selected to result in a processor utilization of approximately 90%.

The first two measurements compare VM/SP HPO Release 5.0 and VM/ESA Release 2 on identical hardware. Under these conditions, the results showed a 33% improvement in external response time (AVG LAST(T)) and a 9.1% decrease in processor capacity (ITR(H)).

The response time improvement reflects the benefits of CMS write-behind (introduced in VM/SP Release 6), XA I/O architecture, and minidisk caching. The 27% decrease in real I/O (RIO/CMD(V)) comes from the write-behind and minidisk caching improvements.

The ITR decrease is primarily because:

The 9121-320 is a uniprocessor. Multiprocessor support can be (and was) excluded from HPO, while it is always present with VM/ESA. This gives HPO an ITR advantage on uniprocessors that has been estimated to be in the

range of 4% to 6%. In addition to this, VM/ESA's multiprocessor support is more efficient than HPO's but this advantage does not apply to this uniprocessor environment.

The extensive new functions that were provided by VM/SP Release 6 CMS (especially 31-bit support and SFS capability) caused a significant increase in processor usage by CMS guests. Although improvements were made in subsequent releases, some of this impact remains. The 15% increase in virtual processor time (EMUL/CMD(H)) reflects this factor.

The ITR decrease was somewhat mitigated by the fact that HPO, being a 370-based operating system, had to be run in a 370 LPAR. For the VM/ESA Release 2 measurements, the hardware was run in basic mode (no LPAR). The HPO system was run in a dedicated LPAR, which normally adds negligible overhead. However, there is probably some overhead in this case because the LPAR has to emulate 370 architecture.

A comparison of measurements 2 and 3 shows that VM/ESA Release 2 performed significantly better when most of the 256MB of storage was configured as central storage and the remaining expanded storage was just used for minidisk caching. External response time improved by 19%, while ITR improved by 4%. The ITR improvement resulted from eliminating the overhead associated with moving pages between central and expanded storage. The response time improvement was mostly due to the ITR improvement, which is manifested as reduced processor time (PBT/CMD(H)) and reduced processor utilization (TOTAL(H)).

Measurements 1 and 3 represent a migration from VM/SP HPO Release 5.0 to VM/ESA Release 2 where an attempt is made to configure the 256MB of storage optimally for each operating system. For this comparison, external response time improved by 45%, while ITR decreased by 5.5%.

The measurement results are summarized in the following table. In the table, those metrics identified with an asterisk (*) are identical in meaning for both the HPO and VM/ESA measurements. Exercise caution when comparing HPO to VM/ESA using any of the remaining metrics. Although all these measures are closely analogous, their exact meanings are not identical due to differences between these two operating systems and differences in how the data are collected and reduced. The derivation of each metric, for both HPO and VM/ESA, is provided in the glossary.

RELEASE RUN ID	VM/SP HPO 5 L1HR1033	VM/ESA 2 L15R0951	VM/ESA 2 L15R0950
Environment			
REAL STORAGE	64MB	64MB	192MB
EXP. STORAGE	192MB	192MB	64MB
USERS	1030	950	950
VTAMs	1	1	1
VSCSs	0	0	0
PROCESSORS	1	1	1
Response Time			
TRIV INT	0.060	0.116	0.110
NONTRIV INT	1.020	0.610	0.485
TOT INT	0.260	0.471	0.380
TOT INT ADJ	0.223	0.395	0.319
AVG FIRST (T) *	0.470	0.221	0.203
AVG LAST (T) *	0.650	0.439	0.356
Throughput			
AVG THINK (T) *	25.59	25.36	25.46
ETR	31.43	28.80	28.73
ETR (T) *	36.67	34.36	34.20
ETR RATIO	0.857	0.838	0.840
ITR (H) *	41.08	37.33	38.84
ITR	35.20	31.29	32.62
EMUL ITR	58.53	47.38	48.17
ITRR (H) *	1.000	0.909	0.945
ITRR	1.000	0.889	0.927
Proc. Usage			
PBT/CMD (H) *	24.344	26.787	25.748
PBT/CMD	24.353	26.774	25.729
CP/CMD (H)	9.594	9.819	8.917
CP/CMD	9.709	9.022	8.186
EMUL/CMD (H)	14.750	16.968	16.831
EMUL/CMD	14.645	17.752	17.542
Processor Util.			
TOTAL (H) *	89.26	92.04	88.07
TOTAL	89.30	92.00	88.00
TOTAL EMUL (H)	54.09	58.30	57.57
TOTAL EMUL	53.70	61.00	60.00
TVR(H)	1.65	1.58	1.53
TVR	1.66	1.51	1.47
Storage			
NUCLEUS SIZE (V)	1003KB	2300KB	2300KB
TRACE TABLE (V)	200KB	200KB	200KB
WKSET (V)	63	65	82
PGBLPGS	12635	10684	43278
PGBLPGS/USER	12.3	11.2	45.6
FREEPGS	1997	2474	2467
FREE UTIL	na	0.92	0.93
SHRPGS	347	699	802
Paging			
READS/SEC	160	296	310
WRITES/SEC	107	182	179
PAGE/CMD	7.281	13.911	14.297
PAGE IO RATE (V)	78.454	77.700	75.900
PAGE IO/CMD (V)	2.536	2.261	2.219
XSTOR IN/SEC	177	115	0
XSTOR OUT/SEC	178	340	0
XSTOR/CMD	9.695	13.241	0.000
FAST CLR/CMD	na	6.402	6.052
Queues			
DISPATCH LIST	16.52	23.22	21.08
ELIGIBLE LIST	0.65	0.00	0.00

RELEASE RUN ID	VM/SP HPO 5 L1HR1033	VM/ESA 2 L15R0951	VM/ESA 2 L15R0950
Environment			
REAL STORAGE	64MB	64MB	192MB
EXP. STORAGE	192MB	192MB	64MB
USERS	1030	950	950
VTAMs	1	1	1
VSCSs	0	0	0
PROCESSORS	1	1	1
I/O			
VIO RATE	540	354	355
VIO/CMD	14.733	10.302	10.379
RIO RATE (V)	325	222	222
RIO/CMD (V)	8.863	6.461	6.491
MDSK/CMD	6.000	na	na
MDC READS	na	190	186
MDC WRITES	na	92	93
MDC MODS	na	76	77
MDC HIT RATIO	na	0.92	0.92
PRIVOPs			
PRIVOP/CMD (R)	26.752	17.616	17.519
DIAG/CMD (R)	14.603	27.785	27.668
DIAG 08/CMD	na	0.786	0.760
DIAG 14/CMD	na	0.000	0.000
DIAG 58/CMD	na	1.251	1.257
DIAG 98/CMD	na	2.008	2.047
DIAG A4/CMD	na	3.871	3.859
DIAG A8/CMD	na	1.833	1.930
DIAG 214/CMD	na	12.892	12.718
SIE/CMD	na	57.622	54.265
SIE INTCPT/CMD	na	38.607	38.528
FREE TOTL/CMD	na	74.327	74.409
VTAM Machines			
WKSET (V)	661.75	460	460
TOT CPU/CMD (V)	3.7256	4.3491	4.3207
CP CPU/CMD (V)	1.8926	1.8108	1.8030
VIRT CPU/CMD (V)	1.8330	2.5384	2.5177
DIAG 98/CMD (V)	0.893	2.010	2.058
Note: T=TPNS, V=VMPRF (ESA) or VMMAP (370), H=Hardware Monitor, R=RTM, Unmarked=RTM (ESA) or VMMAP (370), *=Identical meaning for HPO and VM/ESA runs.			

Table 8. CMS-Intensive Migration from VM/SP HPO Release 5 on the 9121-320

9121-480 / Minidisk

Workload: FS7B0R

Hardware Configuration

Processor model: 9121-480
Processors used: 2

For VM/SP HPO 5:

Storage

Real: 64MB
Expanded: 192MB (180MB SWAP/12MB PAGE)

For VM/ESA Release 2:

Storage

Real: 192MB
Expanded: 64MB (All used for MDC)

Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SWAP	SPOOL	- Number of Volumes -		
						TDSK	User	
For VM/SP HPO 5:								
3390-2	3990-3	2						2 R
3390-2	3990-2	2					10	
3390-2	3990-2	2	2	2		3		
3390-2	3990-3	2	2	2		3		
3380K	3880-3	2	2	2	1			1
3380K	3880-3	1						1
3380K	3880-3	2	2		1		2	1
3380K	3880-3	1						1
3380K	3880-3	1						2
For VM/ESA Release 2:								
3390-2	3990-2	4	8		4	4		
3390-2	3990-3	4					10 R	
3390-2	3990-3	2	4		2	2		2 R

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3088-02	1	NA	3.0MB

Software Configuration

Driver: TPNS
Think time distribution: Bactrian
CMS block size: 4KB

Virtual Machines:

<i>Virtual Machine</i>	<i>Number</i>	<i>Type</i>	<i>Machine Size/Mode</i>	<i>PRIORITY</i>	<i>RESERVED</i>	<i>Other Options</i>
For VM/HPO 5:						
VTAM	1	VTAM/VSCS	16MB/370	1	OFF	QDROP OFF FAVOR ON
SMART	1	VM/RTM	6MB/370	1	OFF	QDROP OFF FAVOR ON
Users	1200	User	3MB/370	100	OFF	
For VM/ESA Release 2:						
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	400	QUICKDSP ON
SMART	1	RTM/ESA	16MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
Users	1870	User	3MB/XC	100	OFF	

Measurement Discussion

VM/SP HPO Release 5.0 was measured on a 9121-480 and compared to a similar VM/ESA Release 2 measurement that was already available (run L25R187F, see "9121-480 / Minidisk" on page 42).

As shown in the configuration section, there are certain configuration differences between these two measurements. Although both measurements are with non-constrained I/O subsystems, the DASD layouts are different. In addition, VTAM was run in XA mode and the user virtual machines were run in XC mode for the VM/ESA Release 2 measurement.

For the HPO measurement, the expanded storage was used for swap paging (180MB) and demand paging (12MB). These allocations are the same as those made for the 9121-320 measurements (see "9121-320 / Minidisk" on page 57 for discussion).

The HPO measurement was run with 1200 users. This is close to the point where HPO generally cannot support more users due to the 16MB-line constraint. This was demonstrated by an additional HPO measurement (not shown) that was done with 1300 users. For that measurement, external response time (AVG LAST (T)) averaged 34 seconds. The average time per below-the-line core table scan (1/STEALSCANS) was only 0.2 seconds (versus 3.7 seconds for the 1200 user measurement), verifying that the poor response time was due to a severe 16MB-line constraint. Because HPO could not be run with acceptable performance at a higher number of users, the VM/SP HPO Release 5.0 measurement results show a low processor utilization of 62.5% (UTIL/PROC (H)).

The VM/ESA Release 2 measurement, having no such constraint, was run with a higher number of users (1870) so as to achieve approximately 90% processor utilization. Even though this VM/ESA Release 2 measurement had 56% more users and therefore ran at a much higher processor utilization, external response time was 0.50 seconds, only 0.12 seconds (33%) higher than the VM/SP HPO Release 5.0 base measurement.

Relative to the VM/SP HPO Release 5.0 base measurement, processor capacity (ITR (H)) improved by 6.1%. This is in contrast to the 5.5% ITR decrease observed for the corresponding 9121-320 measurements (see measurements 1 and 3, "9121-320 / Minidisk" on page 57). This improvement is because the 9121-480 is a multiprocessor. On the 320, HPO has a performance advantage because, with HPO, multiprocessor support can be excluded from CP. On the

480, HPO loses this advantage. In addition, VM/ESA's multiprocessor support is more efficient than HPO's multiprocessor support.

The VM/ESA Release 2 measurement was obtained with the users run in XC mode. As was demonstrated by measurements shown in the *VM/ESA Release 1.1 Performance Report*, this decreases ITR by about 2% relative to running the users in 370 mode.

For the purpose of making processor capacity comparisons, ITRs should be compared for measurements taken at similar processor utilizations. Because this was not the case for these two measurements, this can affect the accuracy of the observed 6.1% ITR improvement. However, these results are consistent with the 3090-200J VM/SP HPO Release 5.0 to VM/ESA Release 1.0 measurements used in Table 11 on page 73. Those results, when projected to VM/ESA Release 2, show a 2% ITR improvement when going from VM/SP HPO Release 5.0 to VM/ESA Release 2. That figure applies to an equal-memory scenario where both HPO and VM/ESA are run with 64MB of central storage. The results for measurements 2 and 3 shown in "9121-320 / Minidisk" on page 57 demonstrate that an additional 4% ITR improvement is reasonable when going from VM/ESA with 64MB/192MB (64MB central, 192MB expanded) to VM/ESA with 192MB/64MB.

The measurement results are summarized in the following table. In the table, those metrics identified with an asterisk (*) are identical in meaning for both the HPO and VM/ESA measurements. Exercise caution when comparing HPO to VM/ESA using any of the remaining metrics. Although all these measures are closely analogous, their exact meanings are not identical due to differences between these two operating systems and differences in how the data are collected and reduced. The derivation of each metric, for both HPO and VM/ESA, is provided in the glossary.

RELEASE RUN ID	VM/SP HPO 5 L2HR1204	VM/ESA 2 L25R187F
Environment		
REAL STORAGE	64MB	192MB
EXP. STORAGE	192MB	64MB
USERS	1200	1870
VTAMs	1	1
VSCSs	0	0
PROCESSORS	2	2
Response Time		
TRIV INT	0.050	0.127
NONTRIV INT	0.768	0.539
TOT INT	0.210	0.395
TOT INT ADJ	0.175	0.357
AVG FIRST (T) *	0.270	0.330
AVG LAST (T) *	0.374	0.498
Throughput		
AVG THINK (T) *	25.44	25.43
ETR	36.35	60.60
ETR (T) *	43.51	66.99
ETR RATIO	0.835	0.905
ITR (H) *	69.60	73.86
ITR	58.18	33.44
EMUL ITR	108.37	49.18
ITRR (H) *	1.000	1.061
ITRR	1.000	0.575
Proc. Usage		
PBT/CMD (H) *	28.736	27.078
PBT/CMD	28.716	27.021
CP/CMD (H)	13.207	9.364
CP/CMD	13.299	8.659
EMUL/CMD (H)	15.529	17.714
EMUL/CMD	15.416	18.362
Processor Util.		
TOTAL (H) *	125.04	181.38
TOTAL	124.95	181.00
UTIL/PROC (H) *	62.52	90.69
UTIL/PROC	62.48	90.50
TOTAL EMUL (H)	67.57	118.66
TOTAL EMUL	67.08	123.00
TVR(H)	1.85	1.53
TVR	1.86	1.47
Storage		
NUCLEUS SIZE (V)	1040KB	2304KB
TRACE TABLE (V)	200KB	400KB
WKSET (V)	62	77
PGBLPGS	12623	39665
PGBLPGS/USER	10.5	21.2
FREEPGS	2277	4876
FREE UTIL	na	0.95
SHRPGS	325	1177
Paging		
READS/SEC	238	586
WRITES/SEC	174	389
PAGE/CMD	9.453	14.555
PAGE IO RATE (V)	112.802	177.300
PAGE IO/CMD (V)	2.551	2.647
XSTOR IN/SEC	188	0
XSTOR OUT/SEC	189	0
XSTOR/CMD	8.658	0.000
FAST CLR/CMD	na	6.404
Queues		
DISPATCH LIST	16.79	45.11
ELIGIBLE LIST	0.02	0.00

RELEASE RUN ID	VM/SP HPO 5 L2HR1204	VM/ESA 2 L25R187F
Environment		
REAL STORAGE	64MB	192MB
EXP. STORAGE	192MB	64MB
USERS	1200	1870
VTAMs	1	1
VSCSs	0	0
PROCESSORS	2	2
I/O		
VIO RATE	650	616
VIO/CMD	14.946	9.196
RIO RATE (V)	410	376
RIO/CMD (V)	9.423	5.613
MDSK/CMD	6.021	na
MDC READS	na	367
MDC WRITES	na	181
MDC MODS	na	151
MDC HIT RATIO	na	0.92
PRIVOPs		
PRIVOP/CMD (R)	28.509	17.507
DIAG/CMD (R)	14.785	26.135
DIAG 08/CMD	na	0.776
DIAG 14/CMD	na	0.015
DIAG 58/CMD	na	1.224
DIAG 98/CMD	na	0.866
DIAG A4/CMD	na	3.852
DIAG A8/CMD	na	1.926
DIAG 214/CMD	na	13.197
SIE/CMD	na	51.877
SIE INTCPT/CMD	na	35.276
FREE TOTL/CMD	na	64.626
VTAM Machines		
WKSET (V)	783	415
TOT CPU/CMD (V)	4.5710	3.7736
CP CPU/CMD (V)	2.5198	1.6836
VIRT CPU/CMD (V)	2.0513	2.0900
DIAG 98/CMD (V)	1.042	0.872
Note: T=TPNS, V=VMPRF (ESA) or VM MAP (370), H=Hardware Monitor, R=RTM, Unmarked=RTM (ESA) or VM MAP (370), *=Identical meaning for HPO and VM/ESA runs		

Table 9. CMS-Intensive Migration from VM/SP HPO Release 5 on the 9121-480

CMS-Intensive Migration from VM/SP Release 5

This section presents the results of four CMS-intensive measurements that were made on the 9221-120 to examine the performance effects of migrating from VM/SP Release 5. The results show the following:

- When going from VM/SP Release 5 to VM/ESA Release 1.0 (370 Feature), both with 16MB of real storage, ITR decreased by 8% and external response time increased by 0.3 seconds (25%).
- When going from VM/SP Release 5 with 16MB to VM/ESA Release 1.0 (370 Feature) with 32MB, ITR decreased by 1.5% and external response time was equivalent.
- When going from VM/SP Release 5 with 16MB to VM/ESA Release 2 with 32MB of real storage and 16MB of expanded storage (all used for minidisk caching), ITR decreased by 17% and external response time increased by 0.7 seconds (60%).

For the CMS-intensive environment, the performance effects of migrating from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 2 are especially processor-dependent. As discussed in “9221-170 / Equal Number of Users / With MDC” on page 126, results on the 9221-170 show a 4% ITR decrease and a 29% external response time improvement.

See “VSE/ESA Guest Migration from VM/SP Release 5” on page 100 for an analogous set of results for the VSE guest environment (PACE workload) on the 9221-170 and 9221-120 processors. Those results show, for both processor models, an ITR decrease of about 2% when going from VM/SP Release 5 to VM/ESA Release 1.0 (370 Feature) and an ITR decrease of about 15% when going from VM/SP Release 5 to VM/ESA Release 2.

Workload: FS7B0R

Hardware Configuration

Processor model: 9221-120
 Processors used: 1

For VM/SP 5 and VM/ESA Release 1.0 (370 Feature):

Storage
 Real: 16MB-32MB (see table)
 Expanded: None

For VM/ESA Release 2:

Storage
 Real: 32MB
 Expanded: 16MB (all retained for MDC)

Tape: 3480 (MONITOR)

DASD:

<i>Type of DASD</i>	<i>Control Unit</i>	<i>Number of Paths</i>	<i>PAGE</i>	<i>SPOOL</i>	<i>- Number of Volumes -</i>			<i>System</i>
					<i>TDSK</i>	<i>User</i>	<i>Server</i>	
9336-10	6310	1		1	1	1		1
9336-20	6310	1	1		1	1		1
9336-10	6310	1	1		1	1		1
9336-20	6310	1	1	1		2		

Communications:

<i>Control Unit</i>	<i>Number</i>	<i>Speed</i>
3088-02	1	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB (Page-aligned for FBA devices)

Virtual Machines:

<i>Virtual Machine</i>	<i>Number</i>	<i>Type</i>	<i>Machine Size/Mode</i>	<i>PRIORITY/ RELSHARE</i>	<i>RESERVED</i>	<i>Other Options</i>
For VM/SP 5 and VM/ESA Release 1.0 (370 Feature):						
VTAM	1	VTAM/VSCS	16MB/370	1	OFF	QDROP OFF FAVOR DELAY=0.1
SMART	1	VM/RTM	3MB/370	1	OFF	QDROP OFF FAVOR
Users	90	User	3MB/370	64	OFF	
For VM/ESA Release 2:						
VTAM	1	VTAM/VSCS	16MB/370	10000	OFF	QUICKDSP ON IPOLL ON DELAY=0.3
SMART	1	RTM/ESA	4MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	48MB/XA	100	OFF	
Users	90	User	3MB/370	100	OFF	

Measurement Discussion

Four minidisk-only CMS-intensive measurements were obtained on a 9221-120 processor. All measurements were run with 90 users and differed only in storage sizes, VM release, and release-dependent tuning options. The tuning options used in these measurements were the same as for "9221-170 / Minidisk" on page 121. The only difference was that the VM/ESA Release 2 measurement shown here used a VTAM delay of 0.3 seconds (compared to 0.2 seconds for the 9221-170 measurements) to further reduce processor utilization.

The first two measurements shown in the table provide a comparison between VM/SP 5 and VM/ESA Release 1.0 (370 Feature), both run with 16MB of real storage. The results show that external response time (AVG LAST (T)) increased nearly 0.3 seconds (25%) in the VM/ESA Release 1.0 (370 Feature) measurement and internal throughput (ITR (H)) decreased by 8%. The primary causes were an increase in path length and working set in VM/ESA 1.0 CMS. These increases, which were not fully recovered with performance improvements made in VM/ESA Release 1.0 (370 Feature), were mainly due to the 31-bit CMS support and SFS support put into VM/SP 6. The *VM/ESA Release 1.0 Performance Report* includes measurement results that compare VM/SP 6 and VM/ESA Release 1.0 (370 Feature) performance.

The third measurement shown in this table was run with 32MB of real storage to show the possible storage constraint relief (VM/SP 5 only supports up to 16MB of real storage). The results of this measurement show that the extra real storage reduced the paging rate per command (PAGE/CMD) by 50% compared to the 16MB VM/ESA Release 1.0 (370 Feature) measurement. The decreased paging rate led to a reduction in CP processor time per command (CP/CMD (H)), while external response time decreased 0.3 seconds (20%) and internal throughput increased by 7%. A comparison between this measurement and the VM/SP 5 measurement shows external response time was equivalent and that internal throughput rate was 1.5% less. Due to the reduced paging rate, the CP processor time per command (CP/CMD (H)) was lower in this measurement than in VM/SP 5. However, the virtual processor time per command (EMUL/CMD (H)) was larger in this measurement due to the 31-bit CMS support and SFS support in VM/ESA Release 1.0 (370 Feature) CMS.

The VM/ESA Release 2 measurement shows that external response time was approximately 0.7 seconds (60%) larger than both the VM/ESA Release 1.0 (370

Feature) measurement with 32MB of real storage and the VM/SP 5 measurement. The internal throughput for the VM/ESA Release 2 measurement was approximately 16% less than the 32MB VM/ESA Release 1.0 (370 Feature) measurement and 17% less than the VM/SP 5 measurement. Several factors contributed to these results. There was a decrease in processor speed obtained when running the 9221-120 in ESA mode and an increased path length for various CP services, such as paging, in VM/ESA. Because there was only one path to each DASD volume for these measurements, there was no performance improvement achieved due to the XA I/O architecture.

As shown in "9221-170 / Minidisk," VM/ESA Release 2 exhibited an external response time that was 29% less than VM/ESA Release 1.0 (370 Feature) on the 9221-170 processor. The internal throughput rate for VM/ESA Release 2 was 4.1% less than VM/ESA Release 1.0 (370 Feature). Because the 9221-170 is a faster processor, less processor time was required for certain CP functions such as page frame reordering. Finally, the decrease in processor speed going from 370 to ESA mode was not as great on the 9221-170 as it was on the 9221-120.

The measurement results are summarized in the following table. In the table, those metrics identified with an asterisk (*) are identical in meaning for both the VM/SP, VM/ESA Release 1.0 (370 Feature), and VM/ESA measurements. Exercise caution when comparing VM/SP or VM/ESA Release 1.0 (370 Feature) to VM/ESA using any of the remaining metrics. Although all these measures are closely analogous, their exact meanings are not identical due to differences between these two operating systems and differences in how the data are collected and reduced. The derivation of each metric, for both VM/SP, VM/ESA Release 1.0 (370 Feature), and VM/ESA, is provided in the glossary.

RELEASE RUN ID	VM/SP 5	VM/ESA 1.0 370		VM/ESA 2
	H1SR0091	H17R0090	H17R0091	H15R0091
Environment				
REAL STORAGE	16MB	16MB	32MB	32MB
EXP. STORAGE	0MB	0MB	0MB	16MB
USERS	90	90	90	90
VTAMs	1	1	1	1
VSCSs	0	0	0	0
PROCESSORS	1	1	1	1
Response Time				
TRIV INT	0.200	0.240	0.150	0.329
NONTRIV INT	3.444	3.781	3.350	3.091
TOT INT	0.810	0.970	0.770	2.231
TOT INT ADJ	0.749	0.900	0.714	1.786
AVG FIRST (T) *	0.656	0.852	0.673	0.593
AVG LAST (T) *	1.171	1.467	1.176	1.880
Throughput				
AVG THINK (T) *	29.39	29.36	29.03	29.13
ETR	2.87	2.84	2.90	2.47
ETR (T) *	3.10	3.06	3.13	3.09
ETR RATIO	0.925	0.928	0.927	0.800
ITR (H) *	4.21	3.88	4.15	3.48
ITR	3.92	3.61	3.86	2.79
EMUL ITR	6.65	6.05	6.20	4.39
ITRR (H) *	1.000	0.920	0.985	0.826
ITRR	1.000	0.921	0.986	0.712

RELEASE RUN ID	VM/SP 5	VM/ESA 1.0 370		VM/ESA 2
	H1SR0091	H17R0090	H17R0091	H15R0091
Environment				
REAL STORAGE	16MB	16MB	32MB	32MB
EXP. STORAGE	0MB	0MB	0MB	16MB
USERS	90	90	90	90
VTAMs	1	1	1	1
VSCSs	0	0	0	0
PROCESSORS	1	1	1	1
Proc. Usage				
PBT/CMD (H) *	237.259	258.022	240.786	287.237
PBT/CMD	236.138	257.176	240.152	288.435
CP/CMD (H)	109.645	116.997	101.169	124.708
CP/CMD	97.055	103.772	90.728	106.948
EMUL/CMD (H)	127.613	141.025	139.617	162.529
EMUL/CMD	139.083	153.404	149.456	181.487
Processor Util.				
TOTAL (H) *	73.56	78.97	75.32	88.63
TOTAL	73.21	78.71	75.12	89.00
TOTAL EMUL (H)	39.56	43.16	43.67	50.15
TOTAL EMUL	43.12	46.95	46.75	56.00
TVR(H)	1.86	1.83	1.72	1.77
TVR	1.70	1.68	1.61	1.59
Storage				
NUCLEUS SIZE (V)	475KB	514KB	578KB	2460KB
TRACE TABLE (V)	232KB	256KB	320KB	464KB
WKSET (V)	64	78	103	79
PGBLPGS	3519	3403	7467	5979
PGBLPGS/USER	39.1	37.8	83.0	66.4
FREELPGS	400	501	501	323
FREE UTIL	na	na	na	0.82
SHRPGS	223	301	275	670
Paging				
READS/SEC	20	22	12	23
WRITES/SEC	9	11	4	19
PAGE/CMD	9.457	10.511	5.237	13.612
PAGE IO RATE (V)	22.212	25.111	9.388	8.400
PAGE IO/CMD (V)	7.164	8.205	3.001	2.722
XSTOR IN/SEC	na	na	na	0
XSTOR OUT/SEC	na	na	na	0
XSTOR/CMD	na	na	na	0.000
FAST CLR/CMD	na	na	na	6.158
Queues				
DISPATCH LIST	2.60	3.15	2.63	8.17
ELIGIBLE LIST	0.02	0.02	0.02	0.00
I/O				
VIO RATE	54	51	52	37
VIO/CMD	17.463	16.657	16.669	11.991
RIO RATE (H,V)	57	57	42	26
RIO/CMD (H,V)	18.463	18.685	13.457	8.426
MDSK/CMD	6.128	5.555	5.435	na
MDC READS	na	na	na	17
MDC WRITES	na	na	na	6
MDC MODS	na	na	na	5
MDC HIT RATIO	na	na	na	0.93

RELEASE RUN ID	VM/SP 5	VM/ESA 1.0 370		VM/ESA 2
	H1SR0091	H17R0090	H17R0091	H15R0091
Environment				
REAL STORAGE	16MB	16MB	32MB	32MB
EXP. STORAGE	0MB	0MB	0MB	16MB
USERS	90	90	90	90
VTAMs	1	1	1	1
VSCSs	0	0	0	0
PROCESSORS	1	1	1	1
PRIVOPs				
PRIVOP/CMD (R)	24.693	11.359	11.494	14.953
DIAG/CMD (R)	17.119	18.326	18.325	29.924
DIAG 08/CMD	na	na	na	0.648
DIAG 14/CMD	na	na	na	0.000
DIAG 58/CMD	na	na	na	0.972
DIAG 98/CMD	na	na	na	3.565
DIAG A4/CMD	na	na	na	3.565
DIAG A8/CMD	na	na	na	1.620
DIAG 214/CMD	na	na	na	12.315
SIE/CMD	na	na	na	58.335
SIE INTCPT/CMD	na	na	na	41.418
FREE TOTL/CMD	na	na	na	81.993
VTAM Machines				
WKSET (V)	196.25	417.25	454.25	117
TOT CPU/CMD (V)	47.1394	47.4181	46.2664	66.2573
CP CPU/CMD (V)	19.2271	19.5376	19.1223	28.0273
VIRT CPU/CMD (V)	27.9123	27.8805	27.1441	38.2300
DIAG 98/CMD (V)	3.522	3.422	3.434	3.841
Note: T=TPNS, V=VMPRF (ESA) or VM MAP (370), H=Hardware Monitor, R=RTM, Unmarked=RTM (ESA) or VM MAP (370), (H,V)=Hardware Monitor (370) or VMPRF (ESA), *=Identical meaning for VM/SP, VM/ESA Release 1.0 (370 Feature), and VM/ESA runs.				

Table 10. CMS-Intensive Migration from VM/SP Release 5 on the 9221-120

CMS-Intensive Migration Summary

A large body of performance information for the CMS-intensive environment was collected over the last several releases of VM. This section summarizes the internal throughput rate (ITR) data from those measurements to show, for CMS-intensive workloads, the approximate changes in processing capacity that may occur when migrating from one VM release to another. As such, this section can serve as one source of migration planning information.

The performance relationships shown here are limited to the CMS-intensive environment. Other types of VM usage may show different relationships. Furthermore, any one measure such as ITR cannot provide a complete picture of the performance differences between VM releases. The VM performance reports from which the ITR ratios (ITRRs) were extracted can serve as a good source of additional performance information. Those reports are listed on page 78.

Table 11 summarizes the ITR relationships that were observed for the CMS-intensive environment for a number of VM release-to-release transitions:

Source	Target	Case	ITRR	ITRR Derivation	Notes
VM/SP 5	VM/ESA 1.0 (370) VM/ESA 2	9221-120	0.92 0.83	R5 R5*R13a	1,5 1,2,6
VM/SP 6	VM/ESA 1.0 (370) VM/ESA 2	9221-120	1.07 0.96	R6 R6*R13a	5 2,6
VM/ESA 1.0 (370)	VM/ESA 2	9221-120 9221-170	0.90 0.96	R13a R13b*R11	2,6 5,6
VM/SP HPO 5	VM/ESA 2	UP, -4381 MP, -4381	0.91 1.02	RHa RHb*R1E*R11	4,5 3,4,5
VM/XA 2.0	VM/ESA 2		1.11	R20*R21*R1E*R11	
VM/XA 2.1	VM/ESA 2		1.09	R21*R1E*R11	
VM/ESA 1.0 (ESA)	VM/ESA 2		1.05	R1E*R11	
VM/ESA 1.1	VM/ESA 2		1.01	R11	

Table 11. Approximate VM Relative Capacity: CMS-Intensive Environment

Explanation of columns:

- Case** The set of conditions for which the stated ITRR approximately applies. When not specified, no large variations in ITRR were found among the cases that were measured. However, smaller variations are typically seen. These ITRR variations are shown in "Derivation and Supporting Data" on page 75.
- ITRR** The target ITR divided by the source ITR. A number greater than 1.00 indicates an improvement in processor capacity.
- ITRR Derivation** Shows how the ITRR was derived. See "Derivation and Supporting Data" on page 75 for discussion.

Notes:

1. The VM/SP 5 system is assumed to include APAR VM30315, the performance SPE that adds segment protection and 4KB key support. Other measurements have shown that VM/SP 5 ITR is 4% to 6% lower without this APAR.

2. This includes an increase of central storage from 16MB to 32MB to compensate for VM/ESA's larger storage requirements. The VM/ESA case also includes 16MB of expanded storage for minidisk caching.
3. The VM/SP HPO 5 to VM/ESA Release 1.0 (ESA Feature) portion of the derivation was done with reduced think time to avoid a 16MB-line real storage constraint in the HPO case. In cases where the base HPO system is 16MB-line constrained, migration to VM/ESA will yield additional performance benefits by eliminating this constraint.
4. The ESA-capable 4381 models provide less processing capacity when run in ESA mode as compared to 370 mode. Therefore, expect a less favorable ITR ratio when migrating a 4381 configuration from VM/SP HPO 5 to VM/ESA 2.
5. The target VM system supports a larger real memory size than the stated migration source and this potential benefit is not reflected in the stated ITR ratios. Migrations from memory-constrained environments will yield additional ITRR and other performance benefits when the target configuration has additional real storage.

A VM/SP example: The stated VM/SP 5 to VM/ESA Release 1.0 (370 Feature) ITRR of 0.92 is based on measurements done on a system with 16MB of real memory. However, VM/ESA Release 1.0 (370 Feature) supports up to 64MB of real memory (but subject to the 16MB-line constraint). When VM/SP 5 with 16MB was compared to VM/ESA Release 1.0 (370 Feature) with 32MB, an ITRR of 0.98 was observed. See "CMS-Intensive Migration from VM/SP Release 5" on page 67 for details.

A VM/SP HPO example: The stated VM/SP HPO 5 to VM/ESA 2 ITRR of 0.91 (for uniprocessors) is based on measurements done on a 9121-320 system with its 256MB of storage configured as 64MB of real storage and 192MB of expanded storage (64MB/192MB). The 9121-320 had to be configured that way because 64MB is the maximum real storage supported by HPO. When HPO5 (64MB/192MB) was compared to VM/ESA 2 (192MB/64MB), an ITRR of 0.95 was observed. See "9121-320 / Minidisk" on page 57 for details.

6. These results apply to the case where the following recommended tuning is done for the target system:
 - Configure 16MB as expanded storage for minidisk caching.
 - On VM/ESA systems before VM/ESA Release 2, set DSPSLICE to three times the default. On VM/ESA Release 2, use the default value.
 - For the 9221-120, set the VTAM DELAY operand in the VTAM CTCA channel-attachment major node to 0.3 seconds. For the 9221-170, set the VTAM delay to 0.2 seconds.
 - Set IPOLL ON for VTAM.
 - Preload the key shared segments.

See "9221-170 / Minidisk" on page 52 for more information on these tuning items. The purpose of this tuning is to configure VM/ESA for use on ESA-mode 9221 processors. If this tuning is not done, lower ITR ratios will be experienced. For example, for the FS7B0R CMS-intensive workload, going from VM/ESA Release 1.0 (370 Feature) to VM/ESA 1.1 resulted in an ITRR of 0.95 with the above tuning and an ITRR of 0.86 without it. This comparison is shown in the *VM/ESA Release 1.1 Performance Report*.

This table only shows relative performance in terms of ITR ratios (processor capacity). It does not provide, for example, any response time information. An

improved ITR tends to result in better response times and vice versa. However, exceptions occur. An especially noteworthy exception is the migration from 370-based VM releases to VM/ESA. In such migrations, response times have frequently been observed to improve significantly, even in the face of an ITR decrease. One pair of measurements, for example, showed a 30% improvement in response time, even though ITR decreased by 5%. When this occurs, factors such as XA I/O architecture and minidisk caching outweigh the adverse effects of increased processor usage. These factors have an especially significant positive impact on response time because they reduce I/O wait time, which is often the largest component of system response time.

Keep in mind that in an actual migration to a new VM release, other factors (such as hardware, licensed product release levels, and workload) are often changed in the same time frame. It is not unusual for the performance effects from upgrading VM to be outweighed by the performance effects from these additional changes.

These VM ITRR estimates can be used in conjunction with the appropriate hardware ITRR figures to estimate the overall performance change that would result from migrating both hardware and VM. For example, suppose that the new processor's ITR is 1.30 times that of the current system and suppose that the migration also includes an upgrade from VM/XA 2.1 to VM/ESA Release 2. From the above table, the estimated ITRR for migrating from VM/XA 2.1 to VM/ESA 2 is 1.09. Therefore, the estimated overall increase in system capacity is $1.30 \times 1.09 = 1.42$.

The above table represents CMS-intensive performance for the case where all files are on minidisks. The release-to-release ITR ratios for shared file system (SFS) usage are very similar to the ones shown here. SFS release-to-release measurement results are provided in the reports listed on page 78.

Derivation and Supporting Data

This section explains how the ITR ratios shown above were derived.

The derivation column in Table 11 on page 73 shows how the stated ITR ratio was calculated. For example, the ITRR of 1.05 for migrating from VM/ESA 1.0 to VM/ESA 2 was calculated by multiplying the average ITRR for migrating from VM/ESA 1.0 to VM/ESA 1.1 (R1E) by the average ITRR for migrating from VM/ESA 1.1 to VM/ESA 2 (R11): $1.04 \times 1.01 = 1.05$. R1E was calculated by averaging the ITRRs for VM measurement pairs 11 through 14 (see Table 12 on page 77). Likewise, R11 was calculated by averaging the ITRRs for VM measurement pairs 15 through 17.

Except where noted, any given measurement pair represents two measurements where the only difference is the VM release. As such, all the performance results obtained for one of the measurements in the pair can validly be compared to the corresponding results for the other measurement.

By contrast, there are often substantial environmental differences between unpaired measurements. Factors such as number of users, workload, processor model, and I/O configuration will often be different. This greatly limits the kinds of valid inferences that can be drawn when trying to compare data across two or more measurement pairs. For example, response times are very sensitive to a

number of specific environmental factors and therefore should only be compared within a set of controlled, comparable measurements.

For this reason, the above table only covers ITR ratios. Experience has shown that ITR ratios are fairly resistant to changes in the measurement environment. Consequently, combining the ITR ratios observed for individual release transitions (as explained above) provides a reasonably good estimate of the ITR ratio that would result for a migration that spans all those releases.

The ITR ratios shown in Table 11 on page 73 are based on the following pairs of measurements:

Pair Number	Source Run ID	Target Run ID	Processor	Memory	Proc. Util.	Base Pg/cmd	ITR Ratio	Symbol
VM/SP 5 to VM/ESA 1.0 (370 Feature): Reference 3								
1	H1SR0091	H17R0090	9221-120	16MB	80	9	0.92	(R5)
VM/SP 6 to VM/ESA 1.0 (370 Feature): FS7B0 Workload; Reference 1								
2	EC4295	EC7603	4381-13	16MB	70	15	1.069	
3	EC4295	EC7603	4381-13	16MB	80	20	1.075	
avg							1.07	(R6)
VM/ESA 1.0 (370 Feature) to VM/ESA 2, 9221-120: Reference 3								
4	H17R0090	H15R0091	9221-120	16MB, 32MB	80	11	0.90	(R13a)
VM/ESA 1.0 (370 Feature) to VM/ESA 1.1, 9221-170: Reference 2								
5	H17R0281	H14R0287	9221-170	64MB	80	7	0.95	(R13b)
VM/SP HPO 5 to VM/ESA 2: Reference 3								
6	L1HR1033	L15R0951	9121-320	64MB/192MB	90	6	0.91	(RHa)
VM/SP HPO 5 to VM/ESA 1.0 (ESA Feature): Reference 1								
7	Y25R1141	Y23R1143	3090-200J	64MB/512MB	90	22	0.97	(RHb)
VM/XA 2.0 to VM/XA 2.1: Reference 1								
8	Y62R5401	Y6\$R5401	3090-600J	512MB/2GB	90	15	1.02	(R20)
VM/XA 2.1 to VM/ESA 1.0 (ESA Feature): FS7B0R Workload; Reference 1								
9	Y2\$R2001	Y23R2001	3090-200J	256MB/2GB	90	11	1.064	
10	Y6\$R5401	Y63R5405	3090-600J	512MB/2GB	90	12	1.029	
avg							1.04	(R21)
VM/ESA 1.0 (ESA Feature) to VM/ESA 1.1: Reference 2								
11	Y63R5866	Y64R5865	9021-720	512MB/2GB	90	13	1.059	
12	L23R1770	L24R1770	9121-480	192MB/64MB	90	14	1.032	
13	L13R0911	L14R0910	9121-320	192MB/64MB	90	12	1.045	
14	H13R0280	H14R0287	9221-170	48M/16MB	80	11	1.043	
avg							1.04	(R1E)
VM/ESA 1.1 to VM/ESA 2: Reference 3								
15	264RB424	265RB426	9021-900	1GB/4GB	90	16	1.018	
16	L24R1876	L25R187E	9121-480	192MB/64MB	90	14	1.003	
17	H14R0292	H15R0294	9221-170	48MB/16MB	90	12	1.009	
avg							1.01	(R11)

Table 12. Derivation and Supporting Data: VM Measurement Pairs

Explanation of columns:

- Memory** The amount of real storage and (when applicable) expanded storage in the measured configuration.
- Proc. Util.** Approximate processor utilization. The number of users is adjusted so that the source case runs at or near the stated utilization. The target case is then run with the same number of users.
- Base Pg/cmd** The average number of paging operations per command measured for the source case. This value gives an indication of how real-memory-constrained the environment is. For configurations

with expanded storage used for paging, this value includes expanded storage PGIN and PGOUT operations in addition to DASD page I/Os.

Symbol The symbol used to represent this release transition in Table 11 on page 73.

The FS7B0R workload (CMS-intensive, minidisks, remote users simulated by TPNS) was used for all comparisons except those involving VM/SP 6. For those comparisons, the FS7B0 workload was used (CMS-intensive, minidisks, local users simulated by the full screen internal driver (FSID) tool).

The results in this table illustrate that the release-to-release ITR ratios can and do vary to some extent from one measured environment to another.

A complete description of the measurement configuration and results for each of the VM measurement pairs shown above can be found in the following reports:

1. *VM/ESA Release 1.0 Performance Report, ZZ05-0469*
2. *VM/ESA Release 1.1 Performance Report, GG66-3236*
3. *VM/ESA Release 2 Performance Report (this document)*

Table 12 on page 77 refers to this list to show where each measurement pair is documented.

OfficeVision Migration from VM/ESA Release 1.1

The following sections document the performance measurements made in an OfficeVision* (OV/VM) environment on a 9121-480 and a 9021-720. Each environment demonstrated a reduction in processor usage and response time when migrating from VM/ESA Release 1.1 to VM/ESA Release 2.

The results of migrating from VM/ESA Release 1.1 to VM/ESA Release 2 and from OV/VM Release 1.1 to OV/VM Release 1.2 are summarized in the following graphs.

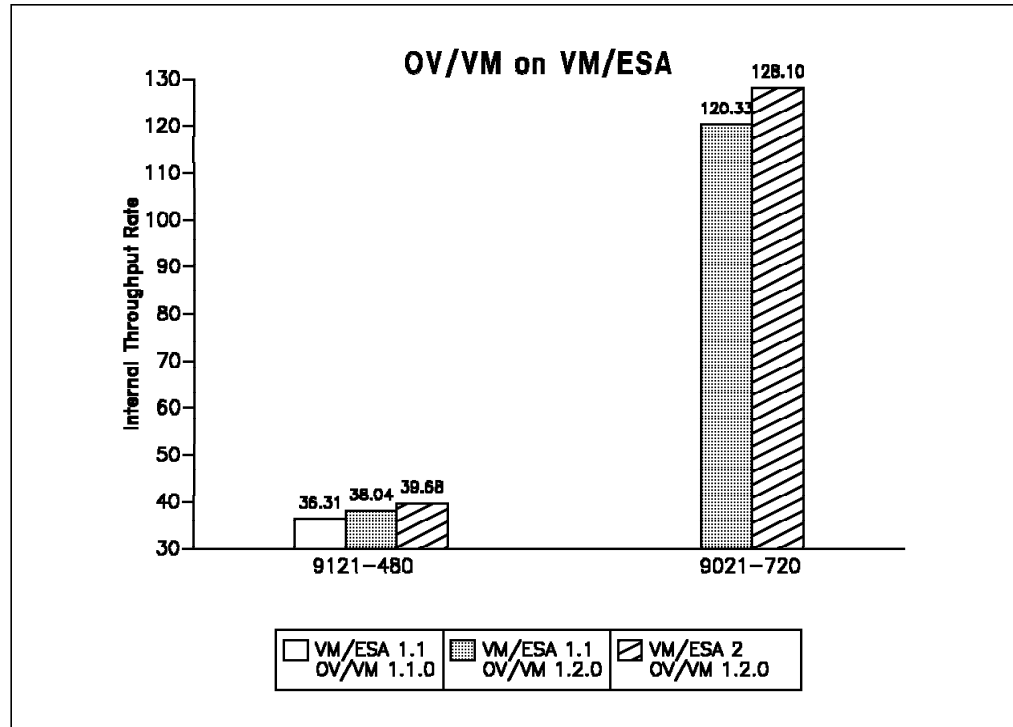


Figure 3. Internal Throughput Rate for OfficeVision/VM

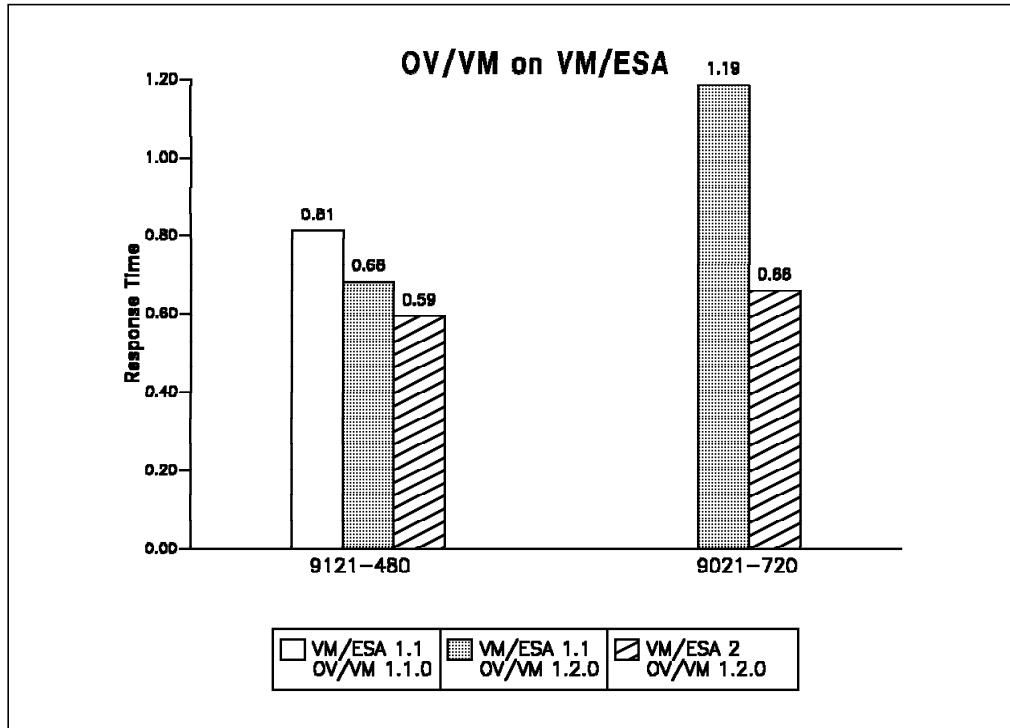


Figure 4. External Response Time for OfficeVision/VM

9121-480

Workload: IOB V2.1

Hardware Configuration

Processor model: 9121-480
 Processors used: 2
 Storage
 Real: 192MB
 Expanded: 64MB (32MB for MDC)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3390-2	3990-2	4	8	4	4			
3390-2	3990-3	2	4	2	2		1 R	2 R
3390-2	3990-3	4				10 R	5 R	

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3088-02	1	NA	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: IOB
 CMS block size: 4KB

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVE	Other Options
VTAMXA	1	VTAM/VSCS	64MB/XC	10000	400	QUICKDSP ON
PRODBM	1	OV/VM	16MB/XA	10000	300	QUICKDSP ON
PROCAL	1	OV/VM	16MB/XA	10000	600	QUICKDSP ON
PROMAIL	1	OV/VM	16MB/XA	10000	OFF	QUICKDSP ON
PROMBXnn	8	OV/VM	16MB/XA	10000	OFF	QUICKDSP ON
PROMBX	1	OV/VM	16MB/XA	10000	OFF	QUICKDSP ON IBCENTRL = Y
WRITER	1	CP MONITOR	24MB/XA	10000	OFF	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
VMCF	1	MONITOR	4MB/XA	10000	OFF	QUICKDSP ON
PRNTEAT1	1	WORKLOAD	2MB/XA	10000	OFF	QUICKDSP ON
Ennnn	8	WORKLOAD	2MB/XA	10000	OFF	QUICKDSP ON
Users	2000/2200	User	2MB/XA	100	OFF	

Note: IBCENTRL = Y is an OV/VM option causing the users' inbaskets to reside in the mail box machines and not on the users' A-disks.

Measurement Discussion

This section documents the migration data collected for an OV/VM environment on a 9121-480. The base starting point was 2000 users running OV/VM Release 1.1.0 on VM/ESA Release 1.1 (see column 1 in Table 13 on page 83). With the hardware configuration available, adequate performance was achieved at about 93% processor utilization (UTIL/PROC (H)), with an external response time (AVG LAST (T)) of 0.81 seconds.

Using VM/ESA Release 1.1 with OV/VM Release 1.1.0 as the base (column 1) and VM/ESA Release 1.1 with OV/VM Release 1.2.0 (column 2), the effects of migrating to the new level of OV/VM can be evaluated. The internal throughput rate on OV/VM Release 1.2.0 improved by 4.8% and external response time improved by 0.13 seconds (16.0%). Information on the enhancements made to OfficeVision/VM R2 SL200 can be found in *OfficeVision/VM Performance comparison R1 SL101 to R2 SL200 on ES9021-500 and 4381* and *OfficeVision/VM Performance comparison OV/VM R1 SL101 to OfficeVision/VM R2 SL200 and Current-OV/VM Direct Connect performance study on ES/9021-580*.

Next, the level of VM was changed to VM/ESA Release 2 (column 3). Many improvements were implemented in this release, as described in "Changes That Affect Performance" on page 5. Of these enhancements, "CP Free Storage" on page 6, "IUCV Storage Management" on page 8, and "Reduced Master Processor Usage" on page 9 contributed the most to the improvement in ITR (H). This migration to VM/ESA Release 2 yielded an improvement in internal throughput rate (ITR (H)) of 4.3% and in external response time of 0.09 seconds (12.9%).

There is some overlap in the improvements made in VM/ESA Release 2 and OV/VM Release 1.2.0; that is, their respective gains are not totally additive. By comparing column 1 to column 3, the combined effects of migrating both VM/ESA and OV/VM can be observed. For this comparison, the internal throughput rate improved 9.3% and external response time improved by 0.22 seconds (26.9%).

With the improvements in both VM/ESA and OV/VM, the processor utilization decreased significantly, indicating that the system could have supported more users. Accordingly, another measurement (column 4) was made with 10% more users (2200) than the base starting point (column 1). This resulted in a similar processor utilization as the base measurement, with an improved internal throughput rate (9.1%) and a slightly higher external response time difference (0.14 seconds, 17.7%). This comparison indicates that the user capacity was improved by almost 10% for this environment.

VM RELEASE OV/VM RELEASE RUN ID	VM/ESA 1.1 OV/VM 1.1.0 L24V200C	VM/ESA 1.1 OV/VM 1.2.0 L24V200D	VM/ESA 2 OV/VM 1.2.0 L25V200B	VM/ESA 2 OV/VM 1.2.0 L25V2201
Environment				
REAL STORAGE	192MB	192MB	192MB	192MB
EXP. STORAGE	64MB	64MB	64MB	64MB
USERS	2000	2000	2000	2200
VTAMs	1	1	1	1
VSCSs	0	0	0	0
PROCESSORS	2	2	2	2
Response Time				
TRIV INT	0.082	0.077	0.067	0.098
NONTRIV INT	0.834	0.670	0.574	0.910
TOT INT	0.800	0.643	0.551	0.865
TOT INT ADJ	0.747	0.603	0.518	0.817
AVG FIRST (T)	0.571	0.481	0.428	0.640
AVG LAST (T)	0.812	0.682	0.594	0.956
Throughput				
AVG THINK (T)	46.26	46.28	46.24	45.92
ETR	31.55	31.71	31.81	35.18
ETR (T)	33.81	33.84	33.84	37.23
ETR RATIO	0.933	0.937	0.940	0.945
ITR (H)	36.31	38.04	39.68	39.62
ITR	16.95	17.85	18.68	18.74
EMUL ITR	32.05	33.61	33.68	33.91
ITRR (H)	1.000	1.047	1.093	1.091
ITRR	1.000	1.053	1.102	1.106
Proc. Usage				
PBT/CMD (H)	55.077	52.581	50.401	50.482
PBT/CMD	55.015	52.606	50.240	50.492
CP/CMD (H)	27.885	26.598	24.208	24.388
CP/CMD	26.028	24.825	22.460	22.560
EMUL/CMD (H)	27.191	25.983	26.194	26.094
EMUL/CMD	28.986	27.781	27.780	27.932
Processor Util.				
TOTAL (H)	186.21	177.92	170.54	187.96
TOTAL	186.00	178.00	170.00	188.00
UTIL/PROC (H)	93.11	88.96	85.27	93.98
UTIL/PROC	93.00	89.00	85.00	94.00
TOTAL EMUL (H)	91.93	87.92	88.63	97.16
TOTAL EMUL	98.00	94.00	94.00	104.00
MASTER TOTAL (H)	93.81	89.84	86.20	94.43
MASTER TOTAL	94.00	90.00	86.00	94.00
MASTER EMUL (H)	34.97	34.19	35.48	38.84
MASTER EMUL	38.00	37.00	38.00	42.00
TVR(H)	2.03	2.02	1.92	1.93
TVR	1.90	1.89	1.81	1.81
Storage				
NUCLEUS SIZE (V)	2792KB	2792KB	2304KB	2304KB
TRACE TABLE (V)	400KB	400KB	400KB	400KB
WKSET (V)	78	78	82	81
PGBLPGS	40033	39978	39893	39174
PGBLPGS/USER	20.0	20.0	19.9	17.8
FREEPGS	5166	5205	5219	5727
FREE UTIL	0.99	0.98	0.98	0.98
SHRPGS	1196	1259	1322	1340
Paging				
READS/SEC	507	502	553	648
WRITES/SEC	455	455	506	590
PAGE/CMD	28.454	28.283	31.297	33.250
PAGE IO RATE (V)	207.000	193.700	161.300	191.300
PAGE IO/CMD (V)	6.123	5.725	4.767	5.138
XSTOR IN/SEC	79	81	80	90
XSTOR OUT/SEC	273	280	280	271
XSTOR/CMD	10.411	10.669	10.639	9.696
FAST CLR/CMD	17.244	17.112	16.668	16.813

VM RELEASE OV/VM RELEASE RUN ID	VM/ESA 1.1 OV/VM 1.1.0 L24V200C	VM/ESA 1.1 OV/VM 1.2.0 L24V200D	VM/ESA 2 OV/VM 1.2.0 L25V200B	VM/ESA 2 OV/VM 1.2.0 L25V2201
Environment				
REAL STORAGE	192MB	192MB	192MB	192MB
EXP. STORAGE	64MB	64MB	64MB	64MB
USERS	2000	2000	2000	2200
VTAMs	1	1	1	1
VSCSs	0	0	0	0
PROCESSORS	2	2	2	2
Queues				
DISPATCH LIST	42.96	39.85	35.42	51.08
ELIGIBLE LIST	0.00	0.00	0.00	0.00
I/O				
VIO RATE	797	805	790	862
VIO/CMD	23.573	23.791	23.347	23.151
RIO RATE (V)	505	495	458	512
RIO/CMD (V)	14.937	14.629	13.535	13.751
MDC READS	412	421	391	432
MDC WRITES	348	348	346	386
MDC MODS	299	300	299	332
MDC HIT RATIO	0.89	0.89	0.89	0.88
PRIVOPs				
PRIVOP/CMD	20.286	20.281	20.789	20.830
DIAG/CMD	88.562	88.605	88.325	87.758
DIAG 08/CMD	12.541	10.669	10.698	10.609
DIAG 14/CMD	1.124	1.271	1.241	1.182
DIAG 58/CMD	2.041	2.039	2.039	2.041
DIAG 98/CMD	1.420	1.478	1.537	1.289
DIAG A4/CMD	11.180	11.290	10.846	10.904
DIAG A8/CMD	5.945	5.911	5.911	5.935
DIAG 214/CMD	31.737	34.371	34.311	34.646
SIE/CMD	129.403	125.604	121.552	122.040
SIE INTCP/CMD	91.876	89.179	86.302	85.428
FREE TOTL/CMD	263.597	245.386	208.084	208.629
VTAM Machines				
WKSET (V)	409	414	401	406
TOT CPU/CMD (V)	5.0118	5.0242	5.0274	4.8467
CP CPU/CMD (V)	2.2676	2.2822	2.2420	2.1609
VIRT CPU/CMD (V)	2.7442	2.7419	2.7855	2.6857
DIAG 98/CMD (V)	1.433	1.490	1.539	1.312
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

Table 13. OfficeVision Migration of VM/ESA and OV/VM on a 9121-480

9021-720

Workload: IOB V2.1

Hardware Configuration

Processor model: 9021-720
 Processors used: 6
 Storage
 Real: 512MB
 Expanded: 2048MB (128MB for MDC)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3390-2	3990-2	4	5	5	5			
3390-2	3990-3	4				10 W	4 W	
3390-2	3990-3	4				8 W	4 W	
3390-2	3990-3	4	3	2	2			2 R
3390-2	3390-2	4	8	4	4			
3390-2	3390-3	4				10 W	4 W	
3390-2	3390-3	4	4	2	2	4 W		

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3088-02	2	NA	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: IOB
 CMS block size: 4KB

Virtual Machines:

<i>Virtual Machine</i>	<i>Number</i>	<i>Type</i>	<i>Machine Size/Mode</i>	<i>RELSHARE</i>	<i>RESERVED</i>	<i>Other Options</i>
VTAMXA	1	VTAM	64MB/XC	10000	100	QUICKDSP ON
VSCSn	4	VSCS	64MB/XC	10000	1024	QUICKDSP ON DEFINITE RESPONSE
PRODBM	1	OV/VM	16MB/XA	10000	400	QUICKDSP ON
PROCAL	1	OV/VM	16MB/XA	10000	2048	QUICKDSP ON
PROMAIL	1	OV/VM	16MB/XA	10000	OFF	QUICKDSP ON
PROMBXnn	20	OV/VM	16MB/XA	10000	OFF	QUICKDSP ON IBCENTRL = Y
PROMBX	1	OV/VM	16MB/XA	10000	OFF	QUICKDSP ON IBCENTRL = Y
WRITER	1	CP MONITOR	24MB/XA	10000	OFF	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
VMCF	1	MONITOR	4MB/XA	10000	OFF	QUICKDSP ON
PRNTEAT1	1	WORKLOAD	2MB/XA	10000	OFF	QUICKDSP ON
Ennnn	20	WORKLOAD	2MB/XA	10000	OFF	QUICKDSP ON
Users	6200/6600	User	2MB/XA	100	OFF	

Note: *IBCENTRL = Y* is an OV/VM option causing the users' inbaskets to reside in the mail box machines and not on the users' A-disks.

Note: APAR VM52691 was included in the VM/ESA Release 1.1 system used for these measurements. This allowed for the support of more than 4096 users using *MSG.

Measurement Discussion

This section documents the migration data collected for an OV/VM environment on a 9021-720. The base starting point was 6200 users running OV/VM Release 1.2.0 on VM/ESA Release 1.1. (see column 1 in Table 14 on page 87). With the hardware configuration available, adequate performance was achieved at about 87% processor utilization (UTIL/PROC (H)), with an external response time (AVG LAST (T)) of 1.19 seconds. The level of VM was then changed to VM/ESA Release 2 (column 2), which yielded an improvement in internal throughput rate (ITR (H)) of 6.5% and in external response time of 0.53 seconds (44.4%).

This improvement is larger than seen on the 9121-480 due to improvements made that were directed to relieve constraints on the high-end processors. See "Reduced Master Processor Usage" on page 9 for a description of the improvements to relieve master processor constraints, which are evident in this environment.

One indicator of a master processor constraint is the amount of master processor emulation time that is consumed on the master compared to the other processors in the configuration. As more master processor work is required, the amount of master emulation that can take place is reduced. For the base VM/ESA Release 1.1 measurement, master emulation (MASTER EMUL (H)) was 2.68%. With the new enhancements made to VM/ESA Release 2, master emulation rose to 10.43%.

With the improvements in VM/ESA Release 2, the processor utilization decreased significantly, indicating that the system could now support more users. Accordingly, another measurement (column 3) was made with 400 more users (6600) than the base starting point (column 1). This resulted in a similar processor utilization as the base measurement, with an internal throughput rate improvement (4.8%) and an external response time improvement (0.34 seconds, 28.9%). The master processor emulation was down slightly from column 2, but still higher than that of the base VM/ESA Release 1.1 measurement (column 1).

An improvement that was implemented late in the development cycle, which is available as APAR VM54161, was also measured. (See "IUCV Storage Management" on page 8 for more information on this APAR.) These results are in column 4 and can be compared to the column 3 measurement which was obtained before the APAR was available. This comparison shows an internal throughput rate improvement of 1.3% and an external response time improvement of 0.08 seconds (9.6%).

RELEASE APAR VM54161 RUN ID	VM/ESA 1.1 NA 064V6204	VM/ESA 2 NO 065V6201	VM/ESA 2 NO 065V6601	VM/ESA 2 YES 065V6602
Environment				
REAL STORAGE	512MB	512MB	512MB	512MB
EXP. STORAGE	2048MB	2048MB	2048MB	2048MB
USERS	6200	6200	6600	6600
VTAMs	1	1	1	1
VSCSs	4	4	4	4
PROCESSORS	6	6	6	6
Response Time				
TRIV INT	0.110	0.061	0.070	0.066
NONTRIV INT	0.872	0.394	0.481	0.451
TOT INT	0.670	0.371	0.412	0.409
TOT INT ADJ	0.860	0.359	0.454	0.416
AVG FIRST (T)	0.866	0.416	0.553	0.484
AVG LAST (T)	1.185	0.659	0.843	0.762
Throughput				
AVG THINK (T)	46.00	46.24	46.18	46.08
ETR	134.39	101.42	122.77	113.26
ETR (T)	104.76	104.84	111.51	111.29
ETR RATIO	1.283	0.967	1.101	1.018
ITR (H)	120.33	128.10	126.05	127.75
ITR	25.72	20.66	23.15	21.68
EMUL ITR	50.10	37.54	42.70	39.50
ITRR (H)	1.000	1.065	1.048	1.062
ITRR	1.000	0.803	0.900	0.843
Proc. Usage				
PBT/CMD (H)	49.864	46.840	47.602	46.966
PBT/CMD	49.924	46.835	47.620	46.996
CP/CMD (H)	26.435	23.144	23.902	23.290
CP/CMD	24.341	21.081	21.792	21.207
EMUL/CMD (H)	23.429	23.696	23.700	23.676
EMUL/CMD	25.582	25.755	25.828	25.789
Processor Util.				
TOTAL (H)	522.37	491.05	530.80	522.67
TOTAL	523.00	491.00	531.00	523.00
UTIL/PROC (H)	87.06	81.84	88.47	87.11
UTIL/PROC	87.17	81.83	88.50	87.17
TOTAL EMUL (H)	245.44	248.42	264.28	263.48
TOTAL EMUL	268.00	270.00	288.00	287.00
MASTER TOTAL (H)	99.33	93.77	97.76	96.81
MASTER TOTAL	99.00	94.00	98.00	97.00
MASTER EMUL (H)	2.68	10.43	8.39	9.08
MASTER EMUL	3.00	12.00	9.30	10.00
TVR(H)	2.13	1.98	2.01	1.98
TVR	1.95	1.82	1.84	1.82
Storage				
NUCLEUS SIZE (V)	4072KB	2304KB	2304KB	2304KB
TRACE TABLE (V)	1200KB	1200KB	1200KB	1200KB
WKSET (V)	66	74	72	72
PGBLPGS	106 k	106 k	104 k	105 k
PGBLPGS/USER	17.1	17.1	15.8	15.9
FREEPGS	15454	15961	16813	16424
FREE UTIL	0.96	0.95	0.96	0.97
SHRPGS	1787	1849	1904	1856

RELEASE APAR VM54161 RUN ID	VM/ESA 1.1 NA 064V6204	VM/ESA 2 NO 065V6201	VM/ESA 2 NO 065V6601	VM/ESA 2 YES 065V6602
Environment				
REAL STORAGE	512MB	512MB	512MB	512MB
EXP. STORAGE	2048MB	2048MB	2048MB	2048MB
USERS	6200	6200	6600	6600
VTAMs	1	1	1	1
VSCSs	4	4	4	4
PROCESSORS	6	6	6	6
Paging				
READS/SEC	204	192	310	294
WRITES/SEC	240	200	306	290
PAGE/CMD	4.238	3.739	5.524	5.248
PAGE IO RATE (V)	92.500	63.800	103.000	98.000
PAGE IO/CMD (V)	0.883	0.609	0.924	0.881
XSTOR IN/SEC	1337	1473	1632	1601
XSTOR OUT/SEC	1729	1796	2086	2035
XSTOR/CMD	29.267	31.182	33.343	32.673
FAST CLR/CMD	16.944	15.996	16.071	16.031
Queues				
DISPATCH LIST	148.45	91.27	121.55	105.24
ELIGIBLE LIST	0.00	0.00	0.02	0.00
I/O				
VIO RATE	2498	2366	2497	2493
VIO/CMD	23.845	22.569	22.393	22.402
RIO RATE (V)	1058	944	1025	1017
RIO/CMD (V)	10.099	9.005	9.192	9.139
MDC READS	1300	1198	1273	1273
MDC WRITES	1099	1074	1144	1144
MDC MODS	958	931	978	989
MDC HIT RATIO	0.90	0.89	0.88	0.89
PRIVOPs				
PRIVOP/CMD	25.469	27.901	26.254	26.846
DIAG/CMD	86.848	83.525	82.934	83.359
DIAG 08/CMD	11.312	10.321	10.214	10.262
DIAG 14/CMD	1.384	0.982	0.960	0.970
DIAG 58/CMD	2.052	2.051	2.054	2.049
DIAG 98/CMD	0.907	1.183	0.969	1.024
DIAG A4/CMD	11.398	10.750	10.833	10.765
DIAG A8/CMD	6.243	5.971	5.955	5.976
DIAG 214/CMD	34.393	33.624	33.558	33.796
SIE/CMD	133.639	124.003	125.551	125.802
SIE INTCPT/CMD	92.211	85.562	84.119	84.287
FREE TOTL/CMD	257.733	209.852	206.263	242.618
VTAM Machines				
WKSET (V)	3432	3136	3308	3284
TOT CPU/CMD (V)	5.2231	5.6247	5.3429	5.3407
CP CPU/CMD (V)	2.8817	3.0189	2.9030	2.8473
VIRT CPU/CMD (V)	2.3414	2.6058	2.4399	2.4933
DIAG 98/CMD (V)	0.905	1.181	0.971	1.027
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

Table 14. OfficeVision Migration of VM/ESA on a 9021-720

MVS/ESA Guest Migration from VM/ESA Release 1.1

Workload: Internal batch workload

Hardware Configuration

Processor model: Physically partitioned 9021-900
 Processors used: 1 or 2
 Storage
 Real: 512MB
 Expanded: 1024MB (256MB dedicated to guest; balance for paging)
 Tape: 3480 (MONITOR)

DASD:

<i>Type of DASD</i>	<i>Control Unit</i>	<i>Number of Paths</i>	<i>- Number of Volumes -</i>			<i>VM Sys.</i>
			<i>PAGE</i>	<i>SPOOL</i>	<i>MVS Sys.</i>	
3380K	3990-3	4			144	
3380K	3880-3	6			3	
3380K	3880-3	2			5	
3380K	3880-3	2	1	1		1

Software Configuration

MVS version: 4.2.2

Virtual Machines:

<i>Virtual Machine</i>	<i>Number</i>	<i>Type</i>	<i>Machine Size/Mode</i>	<i>RELSHARE</i>	<i>RESERVED</i>	<i>Other Options</i>
MVS	1	MVS V=R	128MB/ESA	100	OFF	IOASSIST ON NOTRANS ON
SMART	1	RTM	17MB/370	100	OFF	
WRITER	1	CP MONITOR	1MB/XA	100	OFF	

Measurement Discussion

VM/ESA Release 2 included no code specifically designed to improve the performance of V=R MVS guests, so no change was expected in MVS V=R performance. To verify this, measurements were made of an MVS 4.2.2 V=R system running an internal MVS batch workload. This workload and methodology details are described in the MVS "Workload Description" on page 241. The measurements all featured MVS 4.2.2 running on 1 virtual processor with 128MB of central storage and 256MB of expanded storage. For each release of VM, one measurement was taken of MVS running with a dedicated processor (that is, two physical processors were online to VM) and one with MVS and VM sharing a single physical processor. The measurements were performed on a physically partitioned 9021-900.

Table 15 on page 90 presents the results of these measurements. There is no significant difference in ITR between the two VM releases in the case of the dedicated processor. The VM/ESA Release 2 measurement using a single shared processor shows a small but possibly significant increase in ITR over VM/ESA Release 1.1.

The higher-than-native ETRs deserve comment. Although the ITRs could be measured fairly consistently, there was distinct variability in run time—and thus ETR—when the workload was run on VM. Especially with dedicated processors, the run times on VM were often shorter than native run times. Also, the average number of batch address spaces active during the VM runs (4.5) was slightly higher than during the native run (4.4), indicating that the jobs were slightly less spread out in time when running on VM. Because there were only 32 jobs running over a period of roughly 20 minutes, variation in the order in which jobs started seems to have had an effect on the actual run time.

GUEST TYPE VM RELEASE RUN ID	native NA 22NMVSX1	V=R Ded. VM/ESA 1.1 224MVSD2	V=R Ded. VM/ESA 2 225MVSD4	V=R Shared VM/ESA 1.1 224MVSU2	V=R Shared VM/ESA 2 225MVSU3
Configuration					
Num. Processors	1	1 dedicated	1 dedicated	1 shared	1 shared
Real Storage	128MB	128MB	128MB	128MB	128MB
Expanded Storage	256MB	256MB	256MB	256MB	256MB
Throughput					
Int Throughput (ITR)	0.033579	0.032900	0.032897	0.03233	0.032852
ITR % of Native	100	97.98	97.97	96.28	97.83
Ext Throughput (ETR)	0.027004	0.0272805	0.027515	0.02686	0.0270042
ETR % of Native	100	101.02	101.89	99.50	100.00
Processor Data					
Elapsed Seconds	1185	1173	1163	1191	1185
Processor Busy %	80.42	82.72	83.44	NA	NA
Processor Busy % (V)	NA	0.02	0.2	83.1	82.2
Processor Seconds	953.0	972.7	972.7	989.7	974.1
Batch Data					
Num. of Initiators	5	5	5	5	5
Num. of Batch Jobs	32	32	32	32	32
Chan Path/DASD Data					
No. of DASD Paths	43	43	43	43	43
Avg. Ch Path Busy	4.75	4.80	4.77	4.69	4.63
High Ch Path Busy	27.90	26.45	26.59	26.95	25.93
I/O Interrupt Rate	168.4	170.4	171.2	176.90	167.0
I/O Interrupts/Tran	6236	6246	6222	6584	6184
Paging					
Total: In+Out	0.00	0.00	0.00	0.00	0.00
NSW/NVIO: Total	0.00	0.00	0.00	0.00	0.00
VIO: Total	0.00	0.00	0.00	0.00	0.00
Main Storage Data					
Storage Size	128MB	128MB	128MB	128MB	128MB
Ext. Storage Size	256MB	256MB	256MB	256MB	256MB
Total Frames					
SQA - Avg	3457	1290	3041	1291	3041
LPA - Avg	1952	1913	1924	1943	1922
CSA - Avg	1222	1162	1159	1170	1156
LSQA - Avg	1572	1512	1527	1524	1529
Priv Area - Avg	18860	21087	20865	20949	21242
Unused - Avg	1838	4754	3201	4842	2826
Total - Avg	30976	32768	32768	32768	32768
Fixed Frames					
SQA - Avg	3298	1183	2934	1184	2934
LPA - Avg	45	44	42	44	42
CSA - Avg	1	1	1	1	1
LSQA - Avg	1314	1143	1143	1141	1146
Priv Area - Avg	715	748	729	730	725
Below 16MB - Avg	191	108	140	106	155
Nucleus	2070	1046	1046	1046	1046
Tot Fixed - Avg	7445	4167	5896	4147	5897

GUEST TYPE VM RELEASE RUN ID	native NA 22NMV SX1	V=R Ded. VM/ESA 1.1 224MVSD2	V=R Ded. VM/ESA 2 225MVSD4	V=R Shared VM/ESA 1.1 224MVSU2	V=R Shared VM/ESA 2 225MVSU3
Exp. Storage Frames					
SQA - Avg	0	0	0	0	0
LPA - Avg	0	0	0	0	0
CSA - Avg	3	3	3	3	2
LSQA - Avg	3235	0	0	0	0
Priv Area - Avg	21917	10411	10386	10384	10440
Unused - Avg	43527	55104	55151	55132	55123
Total - Avg	65536	65536	65536	65536	65536
Note: V=VMPRF, Unmarked=RMF					

Table 15. MVS Batch Measurement Data: MVS 4.2.2 Guests on VM/ESA Release 1.1 and VM/ESA Release 2

VSE/ESA Guest Migration from VM/ESA Release 1.1

VM/ESA Release 2 included no code specifically designed to improve the performance of VSE guests. Little or no change was expected in the performance of the VSE guest environments when compared to VM/ESA Release 1.1. As shown in Figure 5, VM/ESA Release 2 performed as expected.

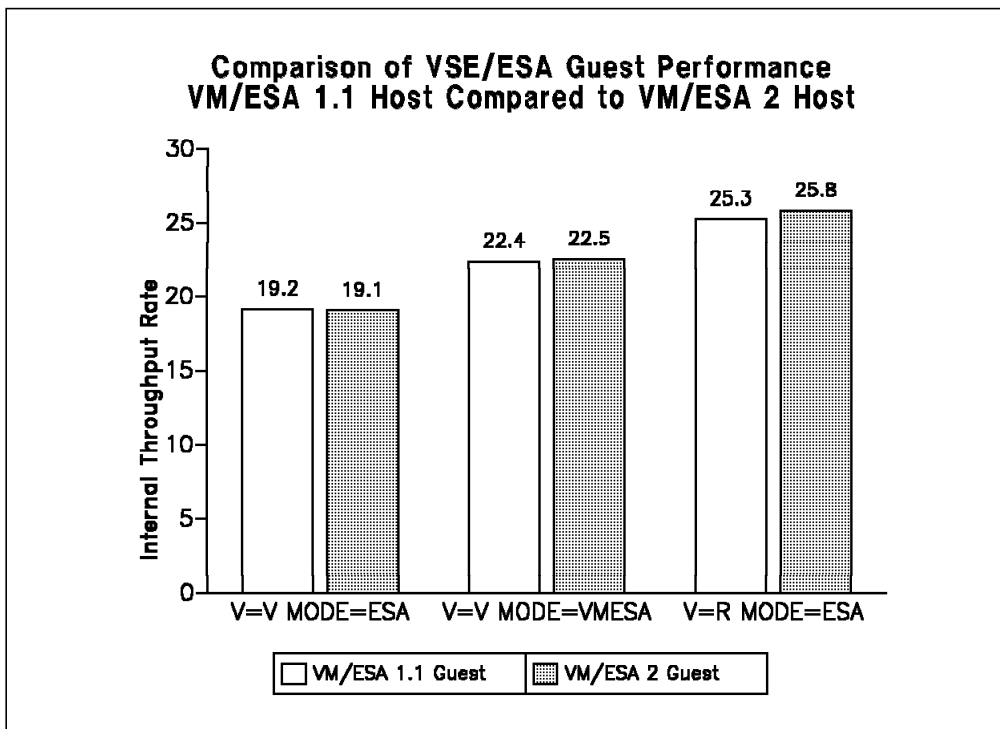


Figure 5. Internal Throughput for VSE Guest Batch Workload. VSE guest environments compared between VM/ESA Release 1.1 and VM/ESA Release 2 on a 9121-320 processor.

The following configuration was used in measurements discussed in this section:

Workload: PACEX8

Hardware Configuration

Processor model: 9121-320⁴
 Processors used: 1
 Storage
 Real: 192MB
 Expanded: 64MB (not used)
 Tape: 3480 (MONITOR)

⁴ See "Hardware Used" on page 25 for an explanation of how this processor model was defined.

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			VM Sys.
					TDSK	VSAM	VSE Sys.	
3390-2	3990-2	2						1
3380-A	3880-03	2					2	
3380-A	3880-03	2				10		

Note: All guest DASD was dedicated to the guest (as opposed to used as mini-disks).

Software Configuration

VSE version: 1.2.0
 VSE supervisor mode: ESA, VMESA as appropriate

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVED	Other Options
For the V=V measurements:						
VSE	1	VSE V=V	16MB/ESA	100	OFF	
For the V=R measurements:						
VSEVR	1	VSE V=R	64MB/ESA	100	OFF	IOASSIST ON NOTRANS ON
SMART	1	RTM	16MB/370	100	OFF	
WRITER	1	CP MONITOR	2MB/XA	100	OFF	

Note: CCW fast path (APAR VM51012) was included in the VM/ESA Release 1.1 systems used for these measurements.

Measurement Discussion

Some general observations can be made about the measurements performed to analyze the three VSE guest scenarios described later:

- The internal throughput rates (ITR (H)) are the same, within measurement variation, for each pair of VM/ESA Release 1.1 and VM/ESA Release 2 measurements.
- For each comparison, the VM/ESA Release 2 case showed a decreased elapsed time (Elapsed Time (H)). Because the workload volume is fixed, an increase in processor utilization (TOTAL (H)) occurred to complete the same volume of work in less elapsed time.
- No paging occurred during any of the measurements. This is typical for this workload run in isolation (that is, when no other major demands are placed on the system other than the guest and its workload).

9121-320 / V=V / MODE=ESA

Table 16 shows the behavior of VSE/ESA 1.2.0 running as a V=V MODE=ESA guest under VM/ESA Release 1.1 and VM/ESA Release 2. The ITRs (ITR (H)) are the same, within measurement variation. The elapsed time for VM/ESA Release 2 is much shorter, giving a higher external throughput rate (ETR (H)). However, it was found that the elapsed time can vary to a large extent for this workload. The processor utilization (TOTAL (H)) varies as well, tending to cancel the effect of the elapsed time, providing a stable internal throughput rate.

RELEASE RUN ID	VM/ESA 1.1 LB48VEX2	VM/ESA 2 LB58VEX4
Environment		
IML MODE	ESA	ESA
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
VM MODE	ESA	ESA
VM SIZE	16MB	16MB
GUEST SETTING	V = V	V = V
VSE SUPERVISOR	ESA	ESA
PROCESSORS	1	1
Throughput (Min)		
Elapsed Time (H)	618.0	555.0
ETR (H)	5.44	6.05
ITR (H)	19.16	19.06
ITR	18.75	18.35
ITRR (H)	1.000	0.995
ITRR	1.000	0.979
Proc. Usage (Sec)		
PBT/CMD (H)	3.132	3.149
PBT/CMD	3.200	3.271
CP/CMD (H)	1.159	1.176
CP/CMD	1.104	1.189
EMUL/CMD (H)	1.973	1.973
EMUL/CMD	2.097	2.081
Processor Util.		
TOTAL (H)	28.38	31.77
TOTAL	29.00	33.00
TOTAL EMUL (H)	17.88	19.91
TOTAL EMUL	19.00	21.00
TVR(H)	1.59	1.60
TVR	1.53	1.57
Storage		
NUCLEUS SIZE (V)	2680KB	2456KB
TRACE TABLE (V)	1008KB	3036KB
VSE WKSET (V)	1354	1371
PGBLPGS	14899	46667
FREEPGS	76	87
FREE UTIL	0.50	0.45
SHRPGS	869	869
Paging		
PAGE/CMD	0.000	0.000
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	88.286	99.107
I/O		
VIO RATE	300.000	327.000
VIO/CMD	3310.714	3240.804
RIO RATE (V)	299.000	343.000
RIO/CMD (V)	3299.679	3399.375
DASD IO TOTAL (V)	178157	163819
DASD IO RATE (V)	296.93	341.29
DASD IO/CMD (V)	3276.82	3382.42

RELEASE RUN ID	VM/ESA 1.1 LB48VEX2	VM/ESA 2 LB58VEX4
Environment		
IML MODE	ESA	ESA
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
VM MODE	ESA	ESA
VM SIZE	16MB	16MB
GUEST SETTING	V=V	V=V
VSE SUPERVISOR	ESA	ESA
PROCESSORS	1	1
PRIVOPs		
PRIVOP/CMD	3319.809	3248.580
DIAG/CMD	217.613	221.600
SIE/CMD	10958.464	10743.214
SIE INTCPT/CMD	9862.618	9668.893
FREE TOTL/CMD	3862.500	6838.393
Note: V=VMPRF, H=Hardware Monitor, Unmarked=RTM. Per command (/CMD), for this workload, means per batch job.		

Table 16. VSE Guest Comparison: V=V MODE=ESA on a 9121-320

9121-320 / V=V / MODE=VMESA

Table 17 shows the results of measuring the performance of VSE/ESA 1.2.0 as a V=V MODE=VMESA guest of VM/ESA Release 1.1 and VM/ESA Release 2. As in the previous case, the internal throughput rate (ITR (H)) was the same within measurement tolerance. There also was an analogous decrease in elapsed time (Elapsed Time (H)) and increase in processor utilization (TOTAL (H)) when moving from VM/ESA Release 1.1 to VM/ESA Release 2.

RELEASE RUN ID	VM/ESA 1.1 LB48VEM2	VM/ESA 2 LB58VEM3
Environment		
IML MODE	ESA	ESA
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
VM MODE	ESA	ESA
VM SIZE	16MB	16MB
GUEST SETTING	V = V	V = V
VSE SUPERVISOR	VMESA	VMESA
PROCESSORS	1	1
Throughput (Min)		
Elapsed Time (H)	634.0	555.0
ETR (H)	5.30	6.05
ITR (H)	22.39	22.54
ITR	22.08	22.42
ITRR (H)	1.000	1.007
ITRR	1.000	1.015
Proc. Usage (Sec)		
PBT/CMD (H)	2.680	2.662
PBT/CMD	2.717	2.676
CP/CMD (H)	1.101	1.102
CP/CMD	1.019	0.991
EMUL/CMD (H)	1.578	1.560
EMUL/CMD	1.698	1.685
Processor Util.		
TOTAL (H)	23.67	26.86
TOTAL	24.00	27.00
TOTAL EMUL (H)	13.94	15.74
TOTAL EMUL	15.00	17.00
TVR(H)	1.70	1.71
TVR	1.60	1.59
Storage		
NUCLEUS SIZE (V)	4220KB	2456KB
TRACE TABLE (V)	3036KB	3036KB
VSE WKSET (V)	1219	1242
PGBLPGS	46441	46668
FREEPGS	79	87
FREE UTIL	0.48	0.45
SHRPGS	852	869
Paging		
PAGE/CMD	0.000	0.000
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	90.571	89.196
I/O		
VIO RATE	291.000	326.000
VIO/CMD	3294.536	3230.893
RIO RATE (V)	296.000	320.000
RIO/CMD (V)	3351.143	3171.429
DASD IO TOTAL (V)	176166	171983
DASD IO RATE (V)	293.61	318.49
DASD IO/CMD (V)	3324.08	3156.43

RELEASE RUN ID	VM/ESA 1.1 LB48VEM2	VM/ESA 2 LB58VEM3
Environment		
IML MODE	ESA	ESA
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
VM MODE	ESA	ESA
VM SIZE	16MB	16MB
GUEST SETTING	V=V	V=V
VSE SUPERVISOR	VMESA	VMESA
PROCESSORS	1	1
PRIVOPs		
PRIVOP/CMD	3298.353	3227.710
DIAG/CMD	42.265	39.224
SIE/CMD	10540.250	10307.143
SIE INTCPT/CMD	9697.030	9482.571
FREE TOTL/CMD	3917.214	6709.554
Note: V=VMPRF, H=Hardware Monitor, Unmarked=RTM. Per command (/CMD), for this workload, means per batch job.		

Table 17. VSE Guest Comparison: V=V MODE=VMESA on a 9121-320

9121-320 / V=R / MODE=ESA

Table 18 shows the measurement comparison for VSE/ESA 1.2.0 V=R MODE=ESA guests of VM/ESA Release 1.1 and VM/ESA Release 2. As the V=V data showed previously, the V=R case yielded similar ITRs (ITR (H)), decreased elapsed time (Elapsed Time (H)), and increased processor utilization (TOTAL (H)). As the DASD used for these measurements were in IOASSIST, the VIO RATE and RIO RATE (H,V) for these measurements were very small compared to the V=V measurements. Because the I/O was being assisted, the VM monitoring tools did not see the true rates. (See "I/O Assist for Guests" on page 150 for a detailed discussion.)

RELEASE RUN ID	VM/ESA 1.1 LB48REX4	VM/ESA 2 LB58REX2
Environment		
IML MODE	ESA	ESA
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
VM MODE	ESA	ESA
VM SIZE	64MB	64MB
GUEST SETTING	V=R	V=R
VSE SUPERVISOR	ESA	ESA
PROCESSORS	1	1
Throughput (Min)		
Elapsed Time (H)	638.0	560.0
ETR (H)	5.27	6.00
ITR (H)	25.27	25.83
ITR	25.08	26.09
ITRR (H)	1.000	1.022
ITRR	1.000	1.040
Proc. Usage (Sec)		
PBT/CMD (H)	2.374	2.323
PBT/CMD	2.392	2.300
CP/CMD (H)	0.477	0.451
CP/CMD	0.456	0.400
EMUL/CMD (H)	1.897	1.872
EMUL/CMD	1.937	1.900
Processor Util.		
TOTAL (H)	20.84	23.23
TOTAL	21.00	23.00
TOTAL EMUL (H)	16.65	18.72
TOTAL EMUL	17.00	19.00
TVR(H)	1.25	1.24
TVR	1.24	1.21
Storage		
NUCLEUS SIZE (V)	2680KB	2452KB
TRACE TABLE (V)	1008KB	1996KB
VSE WKSET (V)	0	0
PGBLPGS	14909	30292
FREEPGS	77	79
FREE UTIL	0.46	0.45
SHRPGS	869	851
Paging		
PAGE/CMD	0.000	0.000
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	0.000	0.000
I/O		
VIO RATE	5.000 *	5.000 *
VIO/CMD	56.964 *	50.000 *
RIO RATE (V)	2.000 *	2.000 *
RIO/CMD (V)	22.786 *	20.000 *
DASD IO TOTAL (V)	175936	171731
DASD IO RATE (V)	293.23	318.02
DASD IO/CMD (V)	3340.69	3180.20

RELEASE RUN ID	VM/ESA 1.1 LB48REX4	VM/ESA 2 LB58REX2
Environment		
IML MODE	ESA	ESA
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
VM MODE	ESA	ESA
VM SIZE	64MB	64MB
GUEST SETTING	V = R	V = R
VSE SUPERVISOR	ESA	ESA
PROCESSORS	1	1
PRIVOPs		
PRIVOP/CMD	59.728	56.800
DIAG/CMD	224.084	216.891
SIE/CMD	3383.679	3150.000
SIE INTCPT/CMD	2977.637	2772.000
FREE TOTL/CMD	763.321	630.000
<p>Note: V=VMPRF, H=Hardware Monitor, Unmarked=RTM. Per command (/CMD), for this workload, means per batch job. I/O-related data marked with an asterisk(*) is as reported, but does not include assisted I/O for dedicated devices.</p>		

Table 18. VSE Guest Comparison: V=R MODE=ESA on a 9121-320

VSE/ESA Guest Migration from VM/SP Release 5

This section shows a comparison of VSE V=V guest performance between VM/SP 5, VM/ESA Release 1.0 (370 Feature), and VM/ESA Release 2 on both 9221-170 and 9221-120 processors. Three PACEX4 measurements, one for each VM release stated above, were obtained on a 9221-170. Each measurement had 16MB of real storage, no expanded storage, and an identical I/O hardware configuration. Another set of three measurements was obtained on a 9221-120 with the same hardware configuration described above. In all of the measurements, the storage size was sufficient to eliminate any paging. The external and internal throughputs for these measurements are shown in the following two graphs.

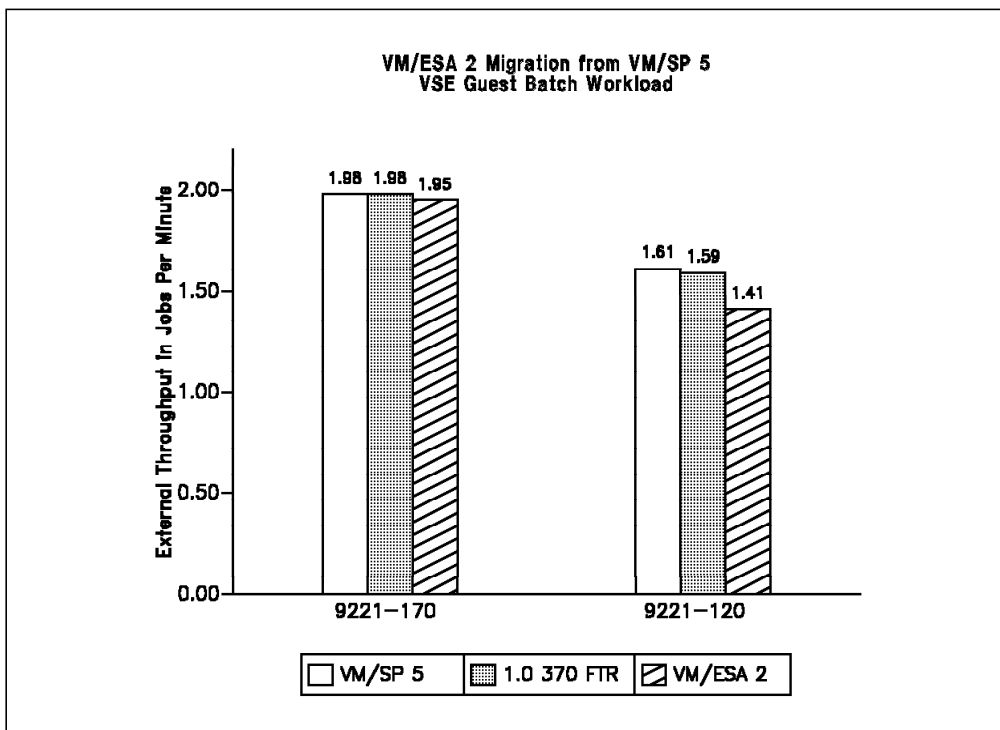


Figure 6. External Throughput for VSE Guest Batch Workload, Migration from VM/SP 5




Figure 7. Internal Throughput for VSE Guest Batch Workload, Migration from VM/SP 5

The graphs show that, on the 9221-170 processor, external throughput was nearly equivalent for VM/SP 5, VM/ESA Release 1.0 (370 Feature), and VM/ESA Release 2. The internal throughput for VM/ESA Release 1.0 (370 Feature) was 1.9% less than VM/SP 5, primarily due to an increase in CP and emulation path length. The internal throughput for VM/ESA Release 2 was 13% less than VM/ESA Release 1.0 (370 Feature) and 15% less than VM/SP 5. The reduced processor speed on the 9221 processor when run in ESA mode accounts for most of the internal throughput decrease in VM/ESA. In fact, the CP path length was smaller (due to fast path CCW translation) in VM/ESA Release 2 than in either VM/SP 5 or VM/ESA Release 1.0 (370 Feature), yet the reduced processor speed still resulted in a higher CP processor time per job (CP/CMD (H)) and a lower internal throughput.

With the 9221-120 processor, the external throughput of the VM/ESA Release 2 measurement was 11.3% less than VM/ESA Release 1.0 (370 Feature) and 12.4% less than VM/SP 5. The primary contributing factor to the reduction in external throughput was that the VM/ESA Release 2 measurement was processor-constrained. This measurement had an average processor utilization (TOTAL (H)) of 90%, with many instances when the processor utilization was over 99%. As with the 9221-170 processor, the reduced processor speed when the 9221 was run in ESA mode resulted in a VM/ESA Release 2 internal throughput that was 13% less than VM/ESA Release 1.0 (370 Feature) and 15% less than VM/SP 5.

9221-170

Workload: PACEX4

Hardware Configuration

Processor model: 9221-170
 Processors used: 1
 Storage
 Real: 16MB
 Expanded: None
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			
					TDSK	VSAM	VSE Sys.	VM Sys.
9336-10	6310	1		1	1			1
9336-20	6310	1	1		1			1
9336-10	6310	1	1		1			
9336-20	6310	1				3	1	
9336-10	6310	1	1	1				
9336-20	6310	1				3	1	

Software Configuration

VSE version: 1.2.0
 VSE supervisor mode: 370

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	PRTY/RELSHARE	RESERVED	Other Options
For VM/SP 5 and VM/ESA Release 1.0 (370 Feature):						
VSE	1	VSE V=V	16MB/370	64	OFF	BMX ECMODE
SMART	1	VM/RTM	3MB/370	1	OFF	QDROP OFF FAVOR
For VM/ESA Release 2:						
VSE	1	VSE V=V	16MB/370	100	OFF	
SMART	1	RTM/ESA	4MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	6MB/XA	100	OFF	

Measurement Discussion

Three PACEX4 measurements were obtained. The first measurement was run on VM/SP 5, the second measurement was run on VM/ESA Release 1.0 (370 Feature), and the third measurement was run on VM/ESA Release 2. Each measurement had an identical hardware configuration. 16MB of real storage was used in all measurements. This was large enough to contain the VSE guest, so no paging occurred in any of the measurements.

The results indicated that external throughput (ETR (H)) was roughly equivalent in all measurements: 1.98 jobs per minute in VM/SP 5 and VM/ESA Release 1.0 (370 Feature), and 1.95 jobs per minute (1.5% less) in VM/ESA Release 2. This difference is believed to be run variation.

The internal throughput rate (ITR (H)) was 1.9% less (6.08 jobs per processor minute) in VM/ESA Release 1.0 (370 Feature) than it was in VM/SP 5, which had an ITR of 6.20 jobs per processor minute. The primary contributing factor to the decrease in internal throughput was a greater path length in CP and emulation, resulting in a higher processor busy time per job (PBT/CMD (H)).

The internal throughput was 13% less in VM/ESA Release 2 (5.27 jobs per processor minute) than it was in VM/ESA Release 1.0 (370 Feature), and it was 15% less than VM/SP 5. The primary contributing factor to the decrease in internal throughput is the reduced processor speed when the 9221 is run in ESA mode. With fast path CCW translation in VM/ESA Release 2, the CP path length per job was lower in VM/ESA Release 2 than it was in either VM/SP 5 or VM/ESA Release 1.0 (370 Feature). However, the loss in processor speed still resulted in a higher CP processor time per job (CP/CMD (H)) in VM/ESA Release 2.

The measurement results are summarized in the following table. In the table, those metrics identified with an asterisk (*) are identical in meaning for the VM/SP, VM/ESA Release 1.0 (370 Feature), and VM/ESA measurements. Exercise caution when comparing VM/SP or VM/ESA Release 1.0 (370 Feature) to VM/ESA using any of the remaining metrics. Although all these measures are closely analogous, their exact meanings are not identical due to differences between these operating systems and differences in how the data are collected and reduced. The derivation of each metric, for VM/SP, VM/ESA Release 1.0 (370 Feature), and VM/ESA, is provided in the glossary.

RELEASE RUN ID	VM/SP 5 HBS4V331	VM/ESA 1.0 370 HB74V33D	VM/ESA 2 HB54V33H
Environment			
IML MODE	370	370	ESA
REAL STORAGE	16MB	16MB	16MB
EXP. STORAGE	0MB	0MB	0MB
VM MODE	370	370	370
VM SIZE	16MB	16MB	16MB
GUEST SETTING	V = V	V = V	V = V
VSE SUPERVISOR	370	370	370
PROCESSORS	1	1	1
Throughput (Min)			
Elapsed Time (H) *	848.4	849.1	859.7
ETR (H) *	1.98	1.98	1.95
ITR (H) *	6.20	6.08	5.27
ITR	6.19	6.10	5.28
ITRR (H) *	1.000	0.982	0.851
ITRR	1.000	0.986	0.853
Proc Usage (Sec)			
PBT/CMD (H) *	9.681	9.863	11.380
PBT/CMD	9.687	9.828	11.360
CP/CMD (H)	4.366	4.480	4.984
CP/CMD	4.027	4.139	4.298
EMUL/CMD (H)	5.315	5.384	6.395
EMUL/CMD	5.660	5.689	7.062
Processor Util.			
TOTAL (H) *	31.95	32.53	37.06
TOTAL	31.97	32.41	37.00
TOTAL EMUL (H)	17.54	17.75	20.83
TOTAL EMUL	18.68	18.76	23.00
TVR(H)	1.82	1.83	1.78
TVR	1.71	1.73	1.61
Storage			
NUCLEUS SIZE (V)	475KB	514KB	2460KB
TRACE TABLE (V)	232KB	256KB	256KB
VSE WKSET (V)	1167	1169.75	1155
PGBLPGS	3519	3403	3199
FREEPGS	400	501	73
FREE UTIL	na	na	0.49
SHRPGS	86	66	312
Paging			
PAGE/CMD	0.000	0.000	0.000
XSTOR/CMD	na	na	0.000
FAST CLR/CMD	na	na	92.110
I/O			
VIO RATE	112.150	112.090	110.000
VIO/CMD	3398.175	3399.040	3377.356
RIO RATE (H)	111.364	111.290	110.527
RIO/CMD (H)	3374.357	3374.786	3393.539
DASD IO TOTAL (V)	94068	94016	93464
DASD IO RATE (V)	111.99	111.92	111.27
DASD IO/CMD (V)	3393.20	3394.00	3416.25
PRIVOPs			
PRIVOP/CMD (R)	5502.557	5358.335	3399.467
DIAG/CMD (R)	221.363	237.574	288.517
SIE/CMD	na	na	7798.622
SIE INTCPT/CMD	na	na	6862.787
FREE TOTL/CMD	na	na	7368.777
Note: V=VMMAP(370)/VMPrf(ESA), H=Hardware Monitor, R=RTM, Unmarked=VMMAP(370)/RTM(ESA), *=Identical meaning for VM/SP, VM/ESA Release 1.0 (370 Feature), and VM/ESA runs.			

Table 19. VSE Guest Migration from VM/SP 5 on 9221-170

9221-120

Workload: PACEX4

Hardware Configuration

Processor model: 9221-120
 Processors used: 1
 Storage
 Real: 16MB
 Expanded: None
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			
					TDSK	VSAM	VSE Sys.	VM Sys.
9336-10	6310	1		1	1			1
9336-20	6310	1	1		1			1
9336-10	6310	1	1		1			
9336-20	6310	1				3	1	
9336-10	6310	1	1	1				
9336-20	6310	1				3	1	

Software Configuration

VSE version: 1.2.0
 VSE supervisor mode: 370

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	PRTY/RELSHARE	RESERVED	Other Options
For VM/SP 5 and VM/ESA Release 1.0 (370 Feature):						
VSE	1	VSE V=V	16MB/370	64	OFF	BMX ECMODE
SMART	1	VM/RTM	3MB/370	1	OFF	QDROP OFF FAVOR
For VM/ESA Release 2:						
VSE	1	VSE V=V	16MB/370	100	OFF	
SMART	1	RTM/ESA	4MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	6MB/XA	100	OFF	

Measurement Discussion

Three PACEX4 measurements were obtained. The first measurement was run on VM/SP 5, the second measurement was run on VM/ESA Release 1.0 (370 Feature), and the third measurement was run on VM/ESA Release 2. Each measurement had an identical hardware configuration. 16MB of real storage was used in all measurements. This was large enough to contain the VSE guest, so no paging occurred in any of the measurements.

The results indicated that external throughput (ETR (H)) was roughly equivalent in VM/SP 5 and VM/ESA Release 1.0 (370 Feature) (1.61 and 1.59 jobs per processor minute respectively). However, in VM/ESA Release 2, the external throughput was 11.3% less than VM/ESA Release 1.0 (370 Feature) and 12.4%

less than VM/SP 5. The reason for the decrease in external throughput was that VM/ESA Release 2 was more processor-constrained than VM/SP 5 or VM/ESA Release 1.0 (370 Feature). For VM/ESA Release 2, the overall processor utilization (TOTAL (H)) was 89.7% compared to 87.8% and 87.1%, respectively, for VM/SP 5 and VM/ESA Release 1.0 (370 Feature). Because of this, there were more periods when the processor utilization was over 99% in the VM/ESA Release 2 measurement. This resulted in periods of especially high processor contention, which reduced external throughput.

The internal throughput (ITR (H)) was 2.2% less (1.81 jobs per processor minute) in VM/ESA Release 1.0 (370 Feature) than it was in VM/SP 5, which had an ITR of 1.85 jobs per processor minute. The internal throughput was 13% less in VM/ESA Release 2 (1.57 jobs per processor minute) than it was in VM/ESA Release 1.0 (370 Feature), and it was 15% less than VM/SP 5. The differences in internal throughput between VM/SP 5, VM/ESA Release 1.0 (370 Feature), and VM/ESA Release 2 on the 9221-120 are due to the same contributing factors as the 9221-170. The internal throughput decrease between VM/SP 5 and VM/ESA Release 1.0 (370 Feature) was due to CP and guest path length increases. The reduction in internal throughput between VM/ESA Release 2 and the 370 releases is due to the reduced processor speed when running the 9221 processor in ESA mode.

The measurement results are summarized in the following table. In the table, those metrics identified with an asterisk (*) are identical in meaning for the VM/SP, VM/ESA Release 1.0 (370 Feature), and VM/ESA measurements. Exercise caution when comparing VM/SP or VM/ESA Release 1.0 (370 Feature) to VM/ESA using any of the remaining metrics. Although all these measures are closely analogous, their exact meanings are not identical due to differences between these operating systems and differences in how the data are collected and reduced. The derivation of each metric, for VM/SP, VM/ESA Release 1.0 (370 Feature), and VM/ESA, is provided in the glossary.

RELEASE RUN ID	VM/SP 5 HBS4V330	VM/ESA 1.0 370 HB74V330	VM/ESA 2 HB54V33E
Environment			
IML MODE	370	370	ESA
REAL STORAGE	16MB	16MB	16MB
EXP. STORAGE	0MB	0MB	0MB
VM MODE	370	370	370
VM SIZE	16MB	16MB	16MB
GUEST SETTING	V = V	V = V	V = V
VSE SUPERVISOR	370	370	370
PROCESSORS	1	1	1
Throughput (Min)			
Elapsed Time (H) *	1040.4	1057.7	1190.3
ETR (H) *	1.61	1.59	1.41
ITR (H) *	1.85	1.81	1.57
ITR	1.84	1.81	1.55
ITRR (H) *	1.000	0.976	0.848
ITRR	1.000	0.981	0.842
Proc Usage (Sec)			
PBT/CMD (H) *	32.351	33.152	38.135
PBT/CMD	32.579	33.203	38.684
CP/CMD (H)	15.991	16.581	17.319
CP/CMD	13.090	13.568	13.603
EMUL/CMD (H)	16.360	16.572	20.816
EMUL/CMD	19.492	19.635	25.081
Processor Util.			
TOTAL (H) *	87.07	87.77	89.71
TOTAL	87.68	87.90	91.00
TOTAL EMUL (H)	44.03	43.87	48.97
TOTAL EMUL	52.46	51.98	59.00
TVR(H)	1.98	2.00	1.83
TVR	1.67	1.69	1.54
Storage			
NUCLEUS SIZE (V)	475KB	514KB	2460KB
TRACE TABLE (V)	232KB	256KB	208KB
VSE WKSET (V)	1154.5	1145.25	1024
PGBLPGS	3519	3403	2400
FREEPGS	400	501	73
FREE UTIL	na	na	0.49
SHRPGS	117	23	115
Paging			
PAGE/CMD	0.000	0.000	0.000
XSTOR/CMD	na	na	0.000
FAST CLR/CMD	na	na	85.021
I/O			
VIO RATE	91.650	89.920	80.000
VIO/CMD	3405.374	3396.578	3400.832
RIO RATE (H)	90.873	89.419	80.558
RIO/CMD (H)	3376.500	3377.643	3424.536
DASD IO TOTAL (V)	93335	94268	92518
DASD IO RATE (V)	91.50	89.78	81.16
DASD IO/CMD (V)	3399.98	3391.25	3449.98
PRIVOPs			
PRIVOP/CMD (R)	6441.896	6455.805	3428.186
DIAG/CMD (R)	224.493	198.257	327.666
SIE/CMD	na	na	8417.058
SIE INTCPT/CMD	na	na	5976.111
FREE TOTL/CMD	na	na	7736.892
Note: V=VMMAP(370)/VMPrf(ESA), H=Hardware Monitor, R=RTM, Unmarked=VMMAP(370)/RTM(ESA), *=Identical meaning for VM/SP, VM/ESA Release 1.0 (370 Feature), and VM/ESA runs.			

Table 20. VSE Guest Migration from VM/SP 5 on 9221-120

VMSES/E Migration from VM/ESA Release 1.1

A number of requirements were addressed in VM Service Enhancements Staged/Extended (VMSES/E) in VM/ESA Release 2. Most of the new function in this release was added to the build function and the Software Inventory Management (VMFSIM) function, some of which caused a degradation in response time for the build command (VMFBLD) when using the STATUS option. However, this same new function did yield performance benefits when executing the build command with the SERVICED option.

Three primary VMSES/E tools that help with the servicing of products were measured:

- VMFREC EXEC receives the raw materials from a service tape and places them into the raw materials database.
- VMFAPPLY EXEC defines new maintenance levels based on the contents of the raw materials database.
- VMFBLD EXEC uses the defined maintenance levels to select the correct level and build the running product.

Overall, for the measurements reported here, the process of receiving and applying CMS service, and building CMS on VMSES/E in VM/ESA Release 2 improved when compared to VMSES/E in VM/ESA Release 1.1. On the 9121-480 configuration and the 9021-720 configuration, the sum of the total elapsed time for the VMSES/E in VM/ESA Release 2 receive, apply, and build commands improved 23% and 24%, respectively, when compared to VMSES/E in VM/ESA Release 1.1.

A number of changes were made to the build function for VMSES/E in VM/ESA Release 2. These changes meant an increased cost in the up-front processing of VMFBLD and a decreased cost in the part handlers function. Hence, VMSES/E in VM/ESA Release 2 experiences slower build execution with the STATUS option and isolated builds, but much improved performance when there are many build lists or objects to process. This improvement comes from the removal of function from a major loop for the part handlers.

Finally, for VMSES/E in VM/ESA Release 2, the savings in terms of new automation are a result of the new part handlers, especially VMFBDMLB—exploited heavily by ESA with all of its MACLIBs. VMFBDMLB saves the manual research of determining which MACLIBs need to be rebuilt. This was made easier in VMSES/E in VM/ESA Release 1.1, but was still manual. It also builds the MACLIBs more quickly because it does it selectively (the old method rebuilt the MACLIB every time). Additional savings of manual tasks are a result of processing build requisites. In the past, the service guide identified any order dependencies. Now they are automatically identified and managed. Because of these improvements, the total time (including manual steps) required to complete these tasks can be significantly reduced.

The measurements described in the following sections are provided to demonstrate the performance impact of these changes on VMSES/E in VM/ESA Release 2.

9121-480 / Receive, Apply and Build Service

Hardware Configuration

Processor model:	9121-480
Processors used:	2
Storage	
Real:	192MB
Expanded:	64MB (not used)
Tape:	3480 (one service tape for the receive command)
DASD:	one 3390-3 with 3990-03 cached control unit

Additional Information

Installation 1: Installed VM/ESA Release 1.1 using VM/ESA Release 1.1 *Installation Guide*.

Installation 2: Installed VM/ESA Release 2 using VM/ESA Release 2 *Installation Guide*.

Segments: default CMS segments used (for example, CMS and VMLIB); additional spool space was defined to fit the default segments onto one volume.

Measurement Discussion

All measurements were performed on dedicated, first-level systems with only one user logged on (the MAINT user ID). The objective of these measurements was to show that enhancements to VMSES/E in VM/ESA Release 2 did not degrade performance when compared to VMSES/E in VM/ESA Release 1.1 in an established service environment. That is, all Software Inventory Management (SIM) tables had been previously initialized using the same Recommended Service Upgrade (RSU) tape with both releases of VMSES/E. The purpose of initializing SIM was to remove the one-time costs associated with setting up SIM.

Once initialized, a Corrective (COR) service tape containing CMS service was loaded onto the system. The performance test system used for these measurements was set up so that the COR tape would be compatible with both VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2. Hence, the VMSES/E function of both releases worked on exactly the same service and the same raw materials database.

The CMS service from the COR tape was received. VMFREC was used to receive a total of 1728 CMS parts from seven tape files. Next, the apply function (VMFAPPLY) was used to process 206 PTFs. The build function (VMFBLD) with the STATUS option was invoked and identified 149 build requirements. Finally, 16 build lists were processed after executing the VMFBLD command with the SERVICED option.

The methodology described in this section applies to both VMSES/E in VM/ESA Release 2 and VMSES/E in VM/ESA Release 1.1. Performance data were collected before and after each command execution to determine total response time and the total amount of resources used by the execution of the command. The performance data were generated by the CP QUERY TIME command and CP IND USER command. No intermediate steps were necessary that required human intervention (for example, entering data, pressing a function key, or

mounting a tape). Hence, the performance data reported were derived from uninterrupted execution of the command.

The following performance indicators were used and can be found in the tables below:

Total Time (seconds): the total elapsed time for the command. This is computed by taking the difference between the start and stop time. More specifically, it is the time after the ENTER key is pressed (the command had already been typed) until the READY message is received.

Total CPU (seconds): the difference in TOTCPU for the user before and after command execution.

Virtual CPU (seconds): the difference in VIRTCPU for the user before and after command execution.

I/O (number): the difference in non-spoiled I/O requests by the user before and after command execution.

Resident Pages growth (pages): the increase in the user's real storage pages after command completion (as compared to the start of command execution).

Working Set Size growth (pages): the increase in the user's working set size after command completion (as compared to the start of command execution).

Two performance factors were not included in the results: 1) the time taken to investigate the necessary steps to invoke the function and 2) the time to manually error check the correctness of the information or the results. (The successful completion of each service command was checked after the command finished.)

Workload: Receive

Command: VMFREC PPF ESA CMS

Scenario Details: 1728 parts received from 7 tape files.

Release	VMSES/E 1.1	VMSES/E 2
Total Time (Q)	608	589
Total CPU (Q)	230	202
Virtual CPU (Q)	211	184
I/O (I)	32898	32955
Resident Pages Growth (I)	314	247
Working Set Size Growth (I)	573	445
Note: Q=CP QUERY TIME, I=CP INDICATE USER.		

Table 21. VMFREC Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9121-480

Workload: Apply

Command: VMFAPPLY PPF ESA CMS

Scenario Details: 206 PTFs after receiving parts from COR tape.

Release	VMSES/E 1.1	VMSES/E 2
Total Time (Q)	612	367
Total CPU (Q)	393	304
Virtual CPU (Q)	382	296
I/O (I)	17936	20510
Resident Pages Growth (I)	136	133
Working Set Size Growth (I)	390	327
Note: Q=CP QUERY TIME, I=CP INDICATE USER.		

Table 22. VMFAPPLY Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9121-480

Workload: Build with STATUS Option

Command: VMFBLD PPF ESA CMS (STATUS

Scenario Details: 149 build requirements identified.

Release	VMSES/E 1.1	VMSES/E 2
Total Time (Q)	97	149
Total CPU (Q)	64	136
Virtual CPU (Q)	62	135
I/O (I)	4014	2677
Resident Pages Growth (I)	212	857
Working Set Size Growth (I)	310	1068
Note: Q=CP QUERY TIME, I=CP INDICATE USER.		

Table 23. VMFBLD (STATUS Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9121-480

Workload: Build with SERVICED Option

Command: VMFBLD PPF ESA CMS (SERVICED

Scenario Details: 16 build lists processed; 149 objects built.

Release	VMSES/E 1.1	VMSES/E 2
Total Time (Q)	1572	1106
Total CPU (Q)	836	622
Virtual CPU (Q)	792	605
I/O (I)	90989	35383
Resident Pages Growth (I)	17	951
Working Set Size Growth (I)	629	889
Note: Q=CP QUERY TIME, I=CP INDICATE USER.		

Table 24. VMFBLD (SERVICED Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9121-480

9021-720 / Receive, Apply and Build Service

Hardware Configuration

Processor model: 9021-720
 Processors used: 6
 Storage
 Real: 512MB
 Expanded: 2 GB (not used)
 Tape: 3480 (one service tape for the receive command)

 DASD: one 3390-2 with 3990-2 non-cached control unit

Additional Information

Installation 1: Installed VM/ESA Release 1.1 using VM/ESA Release 1.1 *Installation Guide*.

Installation 2: Installed VM/ESA Release 2 using VM/ESA Release 2 *Installation Guide*.

Segments: default CMS segments used (for example, CMS and VMLIB); additional spool space defined in order to fit default segments onto one volume.

Measurement Discussion

See “9121-480 / Receive, Apply and Build Service” on page 109 for a discussion of the measurement methodology and performance indicators.

The 9021-720 results shown here and the 9121-480 results show similar trends.

Workload: Receive

Command: VMFREC PPF ESA CMS

Scenario Details: 1728 parts received from 7 tape files.

<i>Release</i>	<i>VMSES/E 1.1</i>	<i>VMSES/E 2</i>
Total Time (Q)	582	508
Total CPU (Q)	262	189
Virtual CPU (Q)	246	174
I/O (I)	33004	32956
Resident Pages Growth (I)	283	252
Working Set Size Growth (I)	470	217
Note: Q=CP QUERY TIME, I=CP INDICATE USER.		

Table 25. VMFREC Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9021-720

Workload: Apply

Command: VMFAPPLY PPF ESA CMS

Scenario Details: 206 PTFs after receiving parts from COR tape.

Release	VMSES/E 1.1	VMSES/E 2
Total Time (Q)	481	343
Total CPU (Q)	412	288
Virtual CPU (Q)	403	282
I/O (I)	24662	20300
Resident Pages Growth (I)	98	133
Working Set Size Growth (I)	191	182
Note: Q=CP QUERY TIME, I=CP INDICATE USER.		

Table 26. VMFAPPLY Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9021-720

Workload: Build with Status Option

Command: VMFBLD PPF ESA CMS (STATUS

Scenario Details: 149 build requirements identified.

Release	VMSES/E 1.1	VMSES/E 2
Total Time (Q)	66	142
Total CPU (Q)	55	130
Virtual CPU (Q)	53	129
I/O (I)	3794	2678
Resident Pages Growth (I)	113	772
Working Set Size Growth (I)	98	860
Note: Q=CP QUERY TIME, I=CP INDICATE USER.		

Table 27. VMFBLD (STATUS Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9021-720

Workload: Build with Serviced Option

Command: VMFBLD PPF ESA CMS (SERVICED

Scenario Details: 16 build lists processed; 149 objects built.

Release	VMSES/E 1.1	VMSES/E 2
Total Time (Q)	1535	1037
Total CPU (Q)	1101	600
Virtual CPU (Q)	1073	585
I/O (I)	86697	35263
Resident Pages Growth (I)	212	893
Working Set Size Growth (I)	777	872
Note: Q=CP QUERY TIME, I=CP INDICATE USER.		

Table 28. VMFBLD (SERVICED Measurement Data: VMSES/E in VM/ESA Release 1.1 and VMSES/E in VM/ESA Release 2 on the 9021-720

New Functional Enhancements

Support for fixed block architecture (FBA) DASD is introduced in VM/ESA Release 2 for the first time since the VM/SP releases. Several measurements were made to evaluate the performance characteristics of FBA support in the CMS-intensive and VSE guest environments. The results are presented in "FBA DASD Support" on page 116.

CP configurability is also introduced in VM/ESA Release 2. Evaluations of WARM START IPL and re-IPL performance were made and are discussed in "CP Configurability" on page 138.

DIAGNOSE code X'250', also new to VM/ESA Release 2, performs synchronous or asynchronous block I/O for a virtual machine's primary address space or for a VM Data Space. A summary and discussion of several test cases is presented in "DIAGNOSE Code X'250'" on page 140.

Finally, the extended CMS file system interfaces, providing a single high-level language CSL interface to minidisk or SFS file systems, is discussed in "Extended CMS File System Interfaces" on page 143.

FBA DASD Support

With VM/ESA Release 2, support is provided for FBA DASD. This support includes the extension of fast path CCW translation and minidisk caching to FBA devices. CMS-intensive and VSE guest migration (from VM/ESA Release 1.0 (370 Feature)) measurements were obtained on a variety of processors and DASD types. In addition to measurements made on FBA DASD, measurements were also made on Count-Key-Data (CKD) devices to provide a migration comparison between FBA and CKD. The first two sections presented here show the performance results for the CMS-intensive and VSE guest workloads, respectively. The last section contains important information related to the new FBA support that affects performance. This last section includes system installation considerations as well as other performance related information.

CMS-Intensive Results

Three pairs of measurements were made on the 9221-170 processor with 64MB total storage. The first pair was a set of VM/ESA Release 1.0 (370 Feature) CMS-intensive minidisk-only (FS7B0R) measurements on 3380 (CKD) and 9336 (FBA) DASD, both run with 290 users to achieve an 80% processor utilization. The second pair was a set of VM/ESA Release 2 FS7B0R measurements on 3380 and 9336 DASD, with 64MB of real storage and no expanded storage, both run with 270 users to achieve an 80% processor utilization. The last pair was a set of VM/ESA Release 2 FS7B0R measurements on 3380 and 9336 DASD, both run with 290 users with 48MB real storage and 16MB expanded storage dedicated exclusively to minidisk cache. The 9336 and 3380 systems were constructed to be identical (as much as possible) to one another. Each 3380 volume had a corresponding 9336 volume with the same data on it, and each system had the same number of channel paths to the DASD volumes. The external response time (AVG LAST (T)) and internal throughput rate (ITR (H)) are shown in the following two graphs.




Figure 8. External Response Time for FS7B0R FBA and CKD Measurements

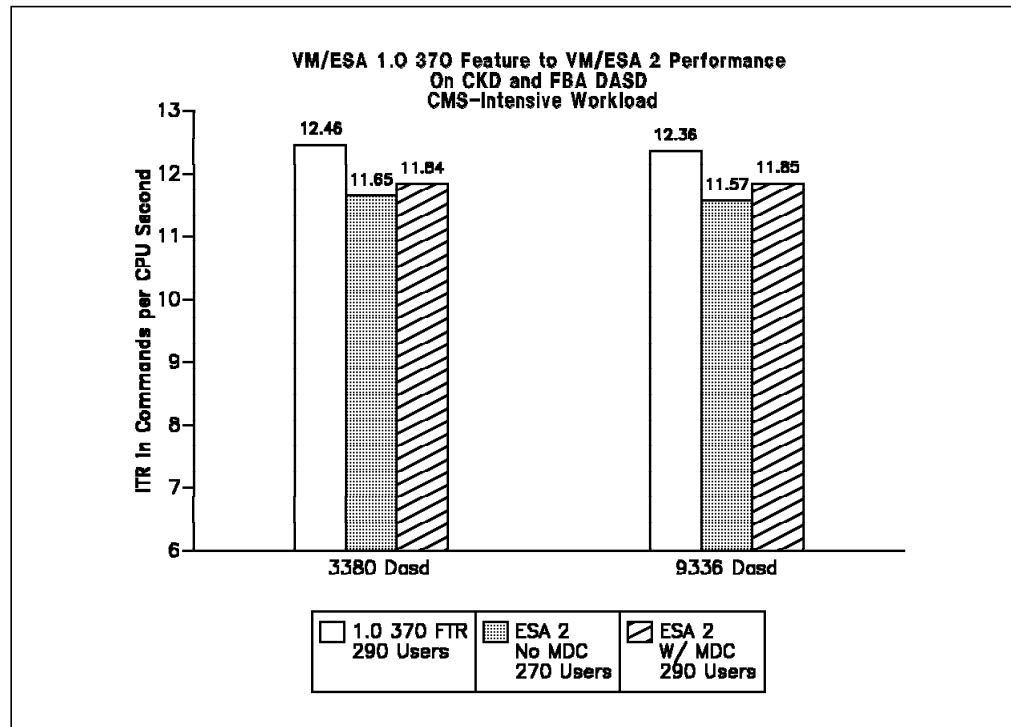


Figure 9. Internal Throughput for FS7B0R FBA and CKD Measurements

The results of the measurements indicated that, for the CMS-intensive workload, the VM/ESA Release 2 FBA measurements exhibited a greater external response time improvement relative to VM/ESA Release 1.0 (370 Feature) than the CKD. In both the equal utilization (no MDC) comparisons (see Table 29 on page 124) and equal user (with MDC) comparisons (see Table 30 on page 129) VM/ESA Release 2 CKD response time was 21% less than VM/ESA Release 1.0 (370 Feature). In contrast, the equal user (with MDC) VM/ESA Release 2 FBA measurement (see Table 30 on page 129) was 29% better in external response time. This was expected, because the elimination of slower FBA I/O through the use of minidisk caching provides a greater possible response time improvement. In addition, the equal utilization VM/ESA Release 2 FBA measurement was 32% better in external response time than VM/ESA Release 1.0 (370 Feature). The response times were greater in the FBA minidisk cache measurement than in the non-minidisk cache measurement because 20 more users were run, which resulted in a higher processor utilization.

The VM/ESA Release 2 FBA measurements also showed a smaller decrease in internal throughput than the CKD. For the equal user measurements (with MDC), the VM/ESA Release 2 CKD internal throughput was 5.0% less than VM/ESA Release 1.0 (370 Feature). For the corresponding FBA measurements, internal throughput was only 4.1% less. For the equal utilization measurements (no MDC), VM/ESA Release 2 CKD internal throughput was 6.5% less than VM/ESA Release 1.0 (370 Feature). The corresponding FBA measurements showed an internal throughput degradation of 6.4%.

Thus for the minidisk-only CMS-intensive workload, migration from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 2 on FBA DASD provided better relative performance than did the same migration from CKD DASD.

VSE Guest Results

Two pairs of VSE V=V guest measurements were obtained on the 9221-170 processor with 64MB total storage (see Table 31 on page 134), and another pair of measurements was obtained on the 9121-320 processor with 64MB total storage (see Table 32 on page 137). The first pair was a set of VM/ESA Release 1.0 (370 Feature) VSE guest PACEX4 measurements with 3380 (CKD) and 9336 (FBA) DASD run on a 9221-170 with 64MB of storage. The second pair was a set of VM/ESA Release 2 VSE guest PACEX4 measurements with 3380 and 9336 DASD run on a 9221-170 with 64MB of real storage and no expanded storage. The last pair was a set of VM/ESA Release 2 VSE guest PACEX4 measurements on 3380 and 3370 (FBA) DASD run on a 9121-320 with 64MB of real storage and no expanded storage. The FBA and CKD systems were constructed to be identical (as much as possible) to one another. Each 3380 volume had a corresponding 9336 or 3370 volume with the same data on it, and each system had the same number of channel paths to the DASD volumes. The external throughput rate (ETR (H)) and internal throughput (ITR (H)) are shown in the following two graphs.




Figure 10. External Throughput for VSE V=V Guest Batch FBA and CKD Measurements

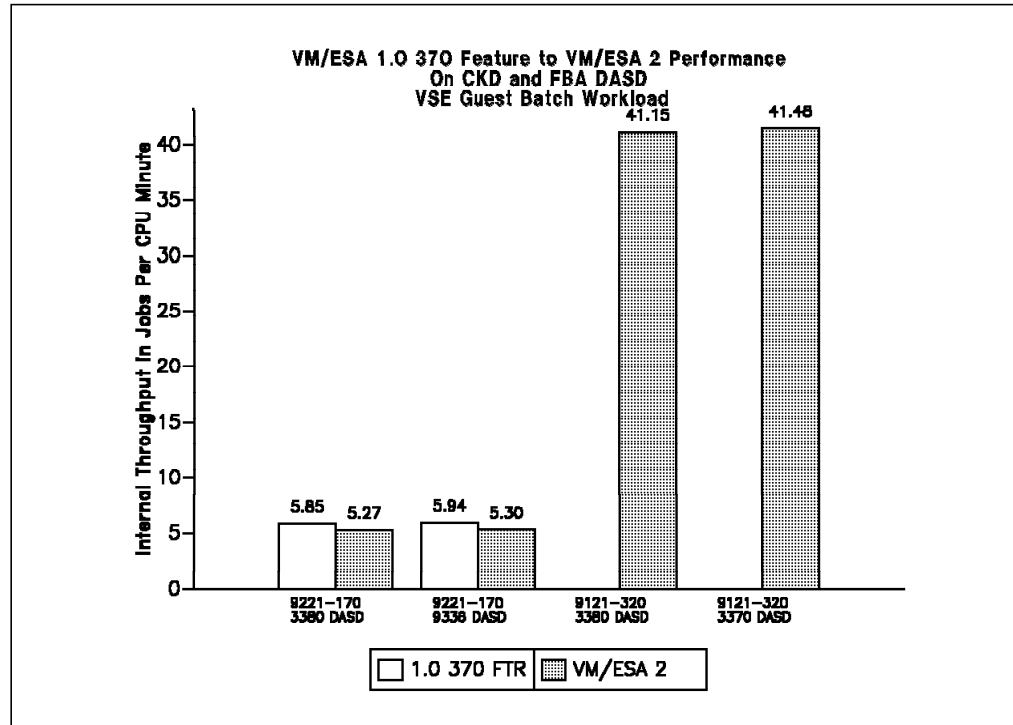


Figure 11. Internal Throughput for VSE V=V Guest Batch FBA and CKD Measurements

As shown here, the results of the measurements indicated that, for the VSE Guest batch workload, the VM/ESA Release 2 FBA measurements exhibited a slightly smaller external throughput decrease relative to VM/ESA Release 1.0 (370 Feature) than the CKD on the 9221-170 processor. On the 9221-170, the VM/ESA Release 2 CKD measurement was 1.3% slower in external throughput than VM/ESA Release 1.0 (370 Feature). The corresponding FBA measurements showed that external throughput was 0.5% less. Note that the external throughput was much lower for the FBA measurements. This was expected for the 9336 DASD because it is much slower than the 3380 DASD, and the PACEX4 workload is very I/O intensive.

The VM/ESA Release 2 FBA measurements on the 9221-170 show a slightly larger degradation in internal throughput than the corresponding CKD measurements. In the VM/ESA Release 2 CKD measurement, internal throughput was 9.9% less than VM/ESA Release 1.0 (370 Feature). The corresponding FBA measurements on the 9221-170 showed a 10.8% degradation. The primary reason for the extra degradation in the FBA case is the longer run time in the FBA measurements and the corresponding extra overhead in entering guest enabled wait state in VM/ESA. For a more thorough discussion, refer to "9221-170 / VSE Guest" on page 131.

The 9121-320 measurements show that the internal throughput for the VM/ESA Release 2 FBA measurement was slightly larger (0.8%) than the corresponding CKD measurement. This is slightly better than the results of the 9221-170 measurements. However, the external throughput gap observed between the FBA and CKD measurements was not nearly as big as the 9221-170, because the 3370 devices are faster than the 9336.

Important Notes on FBA Support Affecting Performance

This section briefly highlights important installation and FBA support information. As explained in "CMS-Intensive Results," because the FBA devices are in general slower than CKD devices, minidisk caching provides a greater benefit in response time reduction. To use minidisk caching with FBA minidisks, it is important for users to format their CMS minidisks with a block size of 4KB and allocate CMS minidisks on a 4KB page boundary on the real DASD volume. Also, for integrated adapter/control unit DASD, including all models of the 9332, 9335, and 9336 DASD, the hardware updates only the sub-channel measurement block I/O request count. When migrating from VM/SP or VM/ESA Release 1.0 (370 Feature), it is important to have page and spool areas on separate DASD volumes. A more complete discussion of these issues is included in "FBA DASD Considerations" on page 19.

9221-170 / Minidisk

A number of measurements were made to evaluate the performance characteristics of FBA support in the CMS-intensive environment. The following presents the results of these measurements for environments with and without the use of minidisk caching.

9221-170 / Equal Processor Utilization / No MDC

Workload: FS7B0R

Hardware Configuration

Processor model: 9221-170
 Processors used: 1
 Storage
 Real: 64MB
 Expanded: None
 Tape: 3480 (MONITOR)

DASD for 9336 Measurements:

<i>Type of DASD</i>	<i>Control Unit</i>	<i>Number of Paths</i>	<i>PAGE</i>	<i>SPOOL</i>	<i>- Number of Volumes -</i>			<i>System</i>
					<i>TDSK</i>	<i>User</i>	<i>Server</i>	
9336-10	6310	1		1	1	1		1
9336-20	6310	1	1		1	1		1
9336-10	6310	1	1		1	1		1
9336-20	6310	1	1	1		2		

DASD for 3380 Measurements:

<i>Type of DASD</i>	<i>Control Unit</i>	<i>Number of Paths</i>	<i>PAGE</i>	<i>SPOOL</i>	<i>- Number of Volumes -</i>			<i>System</i>
					<i>TDSK</i>	<i>User</i>	<i>Server</i>	
3380-A	3880-03	1	1		1	1		
3380-A	3880-03	1	1	1		2		
3380-D	3880-03	1		1	1	1		1
3380-A	3880-03	1	1		1	1		1

Communications:

<i>Control Unit</i>	<i>Number</i>	<i>Speed</i>
3088-02	1	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB (Page aligned for FBA devices)

Virtual Machines:

<i>Virtual Machine</i>	<i>Number</i>	<i>Type</i>	<i>Machine Size/Mode</i>	<i>PRIORITY/ RELSHARE</i>	<i>RESERVED</i>	<i>Other Options</i>
For VM/ESA Release 1.0 (370 Feature):						
VTAM	1	VTAM/VSCS	16MB/370	1	OFF	QDROP OFF FAVOR DELAY=0.1
SMART	1	VM/RTM	3MB/370	1	OFF	QDROP OFF FAVOR
Users	290	User	3MB/370	64	OFF	
For VM/ESA Release 2						
VTAM	1	VTAM/VSCS	16MB/370	10000	OFF	QUICKDSP ON IPOLL ON DELAY=0.2
SMART	1	RTM/ESA	4MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	6MB/XA	100	OFF	
Users (ESA)	270	User	3MB/370	100	OFF	

Measurement Discussion

Two pairs of measurements were obtained for analysis. The first pair of measurements compared VM/ESA Release 2 migration from the VM/ESA Release 1.0 (370 Feature) on 3380 (CKD) DASD. The second pair of measurements provided the same comparison except that 9336 (FBA) DASD were used. The hardware was configured in such a way that the 3380 and 9336 systems were essentially identical to one another. Each 3380 volume had a corresponding 9336 volume with the same data on it, and each system had the same number of channel paths to each volume. This way, a comparison of 3380 versus 9336 performance can also be made. The number of users for each measurement was chosen to provide a processor utilization of approximately 80%. For VM/ESA Release 1.0 (370 Feature), the number of users that provided an 80% utilization was 290. For VM/ESA Release 2, 270 users provided an 80% utilization. All measurements had 64MB of real storage, and there was no expanded storage used in the VM/ESA Release 2 measurements.

Various tuning options were used for the VM/ESA Release 2 measurements to provide better performance. The VTAM DELAY operand in the VTAM CTCA channel-attachment major node was set to 0.2 seconds, because this provided the best response times and the lowest processor utilization. For the 370 Feature, VTAM delay was optimal at 0.1 seconds. Also, SET IPOLL ON was issued for the VTAM server machine to reduce unassisted privileged operations. The FORTRAN and DCF saved segments were preloaded before the measurement began by issuing the SEGMENT LOAD command from an idle user. Because the 9221-170 had a low number of concurrent users, these shared segments were likely to be invalidated and would have caused page reads when the next user loaded the shared segment.

The measurements indicated that when migrating from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 2 on 3380 DASD, the internal throughput (ITR (H)) decreased 6.5% from 12.46 commands per processor second to 11.65. The equivalent measurements on 9336 DASD showed that internal throughput decreased 6.4% from 12.36 to 11.57. Thus, for the minidisk-only CMS-intensive workload, migration from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 2 on FBA DASD provided similar relative performance in terms of internal throughput to the same migration on CKD DASD. The primary contributing factors to the decrease in internal throughput for both pairs of measurements were the reduced processor speed when the 9221 is run in ESA mode, and an

increase in CMS path length between VM/ESA Release 1.0 (370 Feature) and VM/ESA Release 2.

Note that response times were higher in the 9336 measurements. This is to be expected, because the 9336 devices are slower than the 3380 devices. Still, the relative response time difference between VM/ESA Release 1.0 (370 Feature) and VM/ESA Release 2 was 32% lower for the 9336 measurements and 21% lower for the 3380 measurements. In addition, although the paging rates (PAGE/CMD) were higher in both VM/ESA Release 2 measurements, the corresponding real I/O rate to the paging DASD (PAGE IO/CMD (V)) was much smaller than the corresponding 370 Feature measurements due to the block paging used in VM/ESA.

The measurement results are summarized in the following table. In the table, those metrics identified with an asterisk (*) are identical in meaning for both the VM/ESA Release 1.0 (370 Feature) and VM/ESA measurements. Exercise caution when comparing VM/ESA Release 1.0 (370 Feature) to VM/ESA using any of the remaining metrics. Although all these measures are closely analogous, their exact meanings are not identical due to differences between these two operating systems and differences in how the data are collected and reduced. The derivation of each metric, for both VM/ESA Release 1.0 (370 Feature) and VM/ESA, is provided in the glossary.

DASD TYPE	3380 (CKD)		9336 (FBA)	
	VM/ESA 1.0 (370) H17R0295	VM/ESA 2 H15R0271	VM/ESA 1.0 (370) H17R0296	VM/ESA 2 H15R0272
Environment				
REAL STORAGE	64MB	64MB	64MB	64MB
EXP. STORAGE	0MB	0MB	0MB	0MB
USERS	290	270	290	270
VTAMs	1	1	1	1
VSCSs	0	0	0	0
PROCESSORS	1	1	1	1
Response Time				
TRIV INT	0.210	0.175	0.280	0.210
NONTRIV INT	3.026	1.037	3.735	1.143
TOT INT	0.500	0.769	0.620	0.844
TOT INT ADJ	0.405	0.659	0.497	0.723
AVG FIRST (T) *	0.550	0.289	0.613	0.329
AVG LAST (T) *	0.890	0.704	1.139	0.774
Throughput				
AVG THINK (T) *	28.04	28.05	28.23	28.22
ETR	8.16	8.14	7.96	8.08
ETR (T) *	10.07	9.50	9.92	9.43
ETR RATIO	0.811	0.856	0.802	0.857
ITR (H) *	12.46	11.65	12.36	11.57
ITR	10.11	9.98	9.93	9.93
EMUL ITR	16.75	15.46	16.52	15.31
ITRR (H) *	1.000	0.935	0.992	0.929
ITRR	1.000	0.987	0.983	0.982
Proc. Usage				
PBT/CMD (H) *	80.243	85.859	80.890	86.421
PBT/CMD	80.212	86.281	80.767	86.990
CP/CMD (H)	34.275	36.700	35.461	37.108
CP/CMD	31.811	30.514	32.216	30.765
EMUL/CMD (H)	45.968	49.158	45.429	49.313
EMUL/CMD	48.401	55.767	48.551	56.225
Processor Util.				
TOTAL (H) *	80.77	81.60	80.27	81.46
TOTAL	80.74	82.00	80.15	82.00
TOTAL EMUL (H)	46.27	46.72	45.08	46.48
TOTAL EMUL	48.72	53.00	48.18	53.00
TVR(H)	1.75	1.75	1.78	1.75
TVR	1.66	1.55	1.66	1.55
Storage				
NUCLEUS SIZE (V)	706KB	2460KB	706KB	2460KB
TRACE TABLE (V)	444KB	976KB	444KB	976KB
WKSET (V)	77	79	77	79
PGBLPGS	14508	13309	14508	13310
PGBLPGS/USER	50.0	49.3	50.0	49.3
FREEPGS	1201	771	1201	770
FREE UTIL	na	0.89	na	0.89
SHRPGS	387	849	406	848
Paging				
READS/SEC	48	62	47	62
WRITES/SEC	24	50	23	49
PAGE/CMD	7.096	11.785	7.093	11.775
PAGE IO RATE (V)	52.180	22.000	51.587	22.400
PAGE IO/CMD (V)	5.184	2.315	5.198	2.376
XSTOR IN/SEC	na	0	na	0
XSTOR OUT/SEC	na	0	na	0
XSTOR/CMD	na	0.000	na	0.000
FAST CLR/CMD	na	6.313	na	6.365

DASD TYPE	3380 (CKD)		9336 (FBA)	
	VM/ESA 1.0 (370) H17R0295	VM/ESA 2 H15R0271	VM/ESA 1.0 (370) H17R0296	VM/ESA 2 H15R0272
Environment				
REAL STORAGE	64MB	64MB	64MB	64MB
EXP. STORAGE	0MB	0MB	0MB	0MB
USERS	290	270	290	270
VTAMs	1	1	1	1
VSCSs	0	0	0	0
PROCESSORS	1	1	1	1
Queues				
DISPATCH LIST	4.37	10.28	5.90	11.34
ELIGIBLE LIST	0.02	0.00	0.02	0.00
I/O				
VIO RATE	158	106	154	103
VIO/CMD	15.666	11.153	15.567	10.927
RIO RATE (H,V)	141	91	138	90
RIO/CMD (H,V)	13.967	9.575	13.945	9.548
MDSK/CMD	5.563	na	5.442	na
MDC READS	na	0	na	0
MDC WRITES	na	0	na	0
MDC MODS	na	0	na	0
MDC HIT RATIO	na	0.00	na	0.00
PRIVOPs				
PRIVOP/CMD (R)	9.589	14.595	9.657	14.612
DIAG/CMD (R)	17.244	28.498	17.165	28.421
DIAG 08/CMD	na	0.631	na	0.637
DIAG 14/CMD	na	0.000	na	0.000
DIAG 58/CMD	na	1.157	na	1.167
DIAG 98/CMD	na	2.525	na	2.546
DIAG A4/CMD	na	3.893	na	3.819
DIAG A8/CMD	na	2.104	na	2.016
DIAG 214/CMD	na	12.626	na	12.730
SIE/CMD	na	54.820	na	54.528
SIE INTCPT/CMD	na	37.826	na	37.624
FREE TOTL/CMD	na	83.334	na	83.383
VTAM Machines				
WKSET (V)	534.5	204	531.5	208
TOT CPU/CMD (V)	11.8387	17.3906	13.1255	17.5925
CP CPU/CMD (V)	4.3050	7.3946	5.3631	7.4555
VIRT CPU/CMD (V)	7.5337	9.9959	7.7624	10.1371
DIAG 98/CMD (V)	2.221	2.570	2.255	2.577
Note: T=TPNS, V=VMPRF (ESA) or VMMAP (370), H=Hardware Monitor, R=RTM, Unmarked=RTM (ESA) or VMMAP (370), (H,V)=Hardware Monitor (370) or VMPRF (ESA), *=Identical meaning for VM/ESA Release 1.0 (370 Feature) and VM/ESA runs.				

Table 29. 9221-170 VM/ESA Release 2 CMS Migration on FBA and CKD Systems

9221-170 / Equal Number of Users / With MDC

Workload: FS7B0R

Hardware Configuration

Processor model: 9221-170
Processors used: 1

For VM/ESA Release 1.0 (370 Feature):

Storage
Real: 64MB
Expanded: None

For VM/ESA Release 2:

Storage
Real: 48MB
Expanded: 16MB (all retained for MDC)

Tape: 3480 (MONITOR)

DASD for 9336 Measurements:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
9336-10	6310	1		1	1		1	
9336-20	6310	1	1		1		1	
9336-10	6310	1	1		1		1	
9336-20	6310	1	1	1		2		

DASD for 3380 Measurements:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3380-A	3880-03	1	1		1	1		
3380-A	3880-03	1	1	1		2		
3380-D	3880-03	1		1	1		1	
3380-A	3880-03	1	1		1	1	1	

Communications:

Control Unit	Number	Speed
3088-02	1	3.0MB

Software Configuration

Driver: TPNS
Think time distribution: Bactrian
CMS block size: 4KB (Page aligned for FBA devices)

Virtual Machines:

<i>Virtual Machine</i>	<i>Number</i>	<i>Type</i>	<i>Machine Size/Mode</i>	<i>PRIORITY/ RELSHARE</i>	<i>RESERVED</i>	<i>Other Options</i>
For VM/ESA Release 1.0 (370 Feature):						
VTAM	1	VTAM/VSCS	16MB/370	1	OFF	QDROP OFF PRIORITY 1 FAVOR DELAY=0.1
SMART	1	VM/RTM	3MB/370	1	OFF	QDROP OFF FAVOR
Users	290	User	3MB/370	64	OFF	
For VM/ESA Release 2						
VTAM (ESA)	1	VTAM/VSCS	16MB/370	10000	OFF	QUICKDSP ON IPOLL ON DELAY=0.2
SMART	1	RTM/ESA	4MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	64MB/XA	100	OFF	
Users (ESA)	290	User	3MB/370	100	OFF	

Measurement Discussion

As in "9221-170 / Equal Processor Utilization / No MDC" on page 121, two pairs of measurements were obtained. The first pair showed migration from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 2 on 3380 (CKD) DASD, and the second pair showed the same comparison for 9336 (FBA) DASD. The hardware and software setup and workload were exactly the same as in "9221-170 / Equal Processor Utilization / No MDC," except that the VM/ESA Release 2 measurements used 48MB of real storage, and 16MB of expanded storage and an equal number of users (290) was run in all measurements. The same tuning options were also used except that for VM/ESA Release 2, the 16MB of expanded storage was used exclusively for minidisk caching to reduce response times.

The 3380 measurements showed that when migrating from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 2, internal throughput (ITR (H)) decreased 5.0% from 12.46 commands per processor second to 11.84. External response times (AVG LAST (T)) decreased 21% from 0.890 seconds to 0.701. The 9336 measurements showed that internal throughput decreased 4.1% from 12.36 to 11.85, and external response time decreased 29% from 1.139 to 0.808. Thus, with the use of minidisk caching, for FBA systems there was a slightly better relative performance difference when migrating from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 2 than for CKD systems. The primary contributing factors to the decrease in internal throughput for both pairs of measurements were the reduced processor speed when the 9221 was run in ESA mode, and an increase in CMS path length between VM/ESA Release 1.0 (370 Feature) and VM/ESA Release 2.

Note that the 9336 measurements showed a better response time improvement for VM/ESA Release 2 than the 3380 equivalents. This was expected. In both the 3380 and 9336 VM/ESA Release 2 measurements, the minidisk cache hit ratio (MDC HIT RATIO) was 92%. Because the 9336 I/O is slower, the elimination of real I/Os to 9336 DASD improved overall response times by a greater percentage. Refer to "FBA DASD Considerations" on page 19 for special system installation guidelines to make minidisk caching work for FBA minidisks.

The measurement results are summarized in the following table. In the table, those metrics identified with an asterisk (*) are identical in meaning for both the VM/ESA Release 1.0 (370 Feature) and VM/ESA measurements. Exercise caution

when comparing VM/ESA Release 1.0 (370 Feature) to VM/ESA using any of the remaining metrics. Although all these measures are closely analogous, their exact meanings are not identical due to differences between these two operating systems and differences in how the data are collected and reduced. The derivation of each metric, for both VM/ESA Release 1.0 (370 Feature) and VM/ESA, is provided in the glossary.

DASD Type	3380 (CKD)		9336 (FBA)	
	VM/ESA 1.0 (370) H17R0295	VM/ESA 2 H15R0295	VM/ESA 1.0 (370) H17R0296	VM/ESA 2 H15R0296
Environment				
REAL STORAGE	64MB	48MB	64MB	48MB
EXP. STORAGE	0MB	16MB	0MB	16MB
USERS	290	290	290	290
VTAMs	1	1	1	1
VSCSs	0	0	0	0
PROCESSORS	1	1	1	1
Response Time				
TRIV INT	0.210	0.154	0.280	0.182
NONTRIV INT	3.026	1.038	3.735	1.186
TOT INT	0.500	0.757	0.620	0.864
TOT INT ADJ	0.405	0.652	0.497	0.744
AVG FIRST (T) *	0.550	0.299	0.613	0.351
AVG LAST (T) *	0.890	0.701	1.139	0.808
Throughput				
AVG THINK (T) *	28.04	28.28	28.23	28.21
ETR	8.16	8.77	7.96	8.76
ETR (T) *	10.07	10.18	9.92	10.18
ETR RATIO	0.811	0.862	0.802	0.861
ITR (H) *	12.46	11.84	12.36	11.85
ITR	10.11	10.20	9.93	10.21
EMUL ITR	16.75	15.75	16.52	15.73
ITRR (H) *	1.000	0.950	0.992	0.951
ITRR	1.000	1.009	0.983	1.010
Proc. Usage				
PBT/CMD (H) *	80.243	84.490	80.890	84.356
PBT/CMD	80.212	84.493	80.767	84.489
CP/CMD (H)	34.275	35.489	35.461	35.560
CP/CMD	31.811	29.474	32.216	29.473
EMUL/CMD (H)	45.968	49.001	45.429	48.796
EMUL/CMD	48.401	55.019	48.551	55.016
Processor Util.				
TOTAL (H) *	80.77	86.00	80.27	85.86
TOTAL	80.74	86.00	80.15	86.00
TOTAL EMUL (H)	46.27	49.87	45.08	49.67
TOTAL EMUL	48.72	56.00	48.18	56.00
TVR(H)	1.75	1.72	1.78	1.73
TVR	1.66	1.54	1.66	1.54
Storage				
NUCLEUS SIZE (V)	706KB	2460KB	706KB	2460KB
TRACE TABLE (V)	444KB	720KB	444KB	720KB
WKSET (V)	77	78	77	78
PGBLPGS	14508	9195	14508	9186
PGBLPGS/USER	50.0	31.7	50.0	31.7
FREEPGS	1201	819	1201	830
FREE UTIL	na	0.90	na	0.90
SHRPGS	387	826	406	836
Paging				
READS/SEC	48	72	47	73
WRITES/SEC	24	58	23	58
PAGE/CMD	7.096	12.772	7.093	12.870
PAGE IO RATE (V)	52.180	25.600	51.587	26.800
PAGE IO/CMD (V)	5.184	2.515	5.198	2.633
XSTOR IN/SEC	na	0	na	0
XSTOR OUT/SEC	na	0	na	0
XSTOR/CMD	na	0.000	na	0.000
FAST CLR/CMD	na	6.288	na	6.288

DASD Type	3380 (CKD)		9336 (FBA)	
	VM/ESA 1.0 (370) H17R0295	VM/ESA 2 H15R0295	VM/ESA 1.0 (370) H17R0296	VM/ESA 2 H15R0296
Environment				
REAL STORAGE	64MB	48MB	64MB	48MB
EXP. STORAGE	0MB	16MB	0MB	16MB
USERS	290	290	290	290
VTAMs	1	1	1	1
VSCSs	0	0	0	0
PROCESSORS	1	1	1	1
Queues				
DISPATCH LIST	4.37	11.27	5.90	12.54
ELIGIBLE LIST	0.02	0.00	0.02	0.00
I/O				
VIO RATE	158	112	154	110
VIO/CMD	15.666	11.004	15.567	10.807
RIO RATE (H,V)	141	71	138	70
RIO/CMD (H,V)	13.967	6.976	13.945	6.877
MDSK/CMD	5.563	na	5.442	na
MDC READS	na	59	na	58
MDC WRITES	na	27	na	22
MDC MODS	na	22	na	17
MDC HIT RATIO	na	0.92	na	0.92
PRIVOPs				
PRIVOP/CMD (R)	9.589	14.546	9.657	14.584
DIAG/CMD (R)	17.244	28.498	17.165	28.240
DIAG 08/CMD	na	0.688	na	0.688
DIAG 14/CMD	na	0.000	na	0.000
DIAG 58/CMD	na	1.179	na	1.179
DIAG 98/CMD	na	2.456	na	2.358
DIAG A4/CMD	na	3.930	na	3.733
DIAG A8/CMD	na	2.161	na	2.063
DIAG 214/CMD	na	12.772	na	12.673
SIE/CMD	na	53.447	na	53.149
SIE INTCPT/CMD	na	37.413	na	37.204
FREE TOTL/CMD	na	78.402	na	78.005
VTAM Machines				
WKSET (V)	534.5	198	531.5	207
TOT CPU/CMD (V)	11.8387	16.9750	13.1255	17.0014
CP CPU/CMD (V)	4.3050	7.2048	5.3631	7.2044
VIRT CPU/CMD (V)	7.5337	9.7702	7.7624	9.7970
DIAG 98/CMD (V)	2.221	2.456	2.255	2.431
Note: T=TPNS, V=VMPRF (ESA) or VMMAP (370), H=Hardware Monitor, R=RTM, Unmarked=RTM (ESA) or VMMAP (370), (H,V)=Hardware Monitor (370) or VMPRF (ESA), *=Identical meaning for VM/ESA Release 1.0 (370 Feature) and VM/ESA runs.				

Table 30. VM/ESA Release 2 CMS Migration on FBA and CKD Systems with MDC on 9221-170

9221-170 / VSE Guest

Workload: PACEX4

Hardware Configuration

Processor model: 9221-170
 Processors used: 1
 Storage
 Real: 64MB
 Expanded: None
 Tape: 3480 (MONITOR)

DASD for 9336 Measurements:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			
					TDSK	VSAM	VSE Sys.	VM Sys.
9336-10	6310	1		1	1			1
9336-20	6310	1	1		1			1
9336-10	6310	1	1		1			
9336-20	6310	1				3	1	
9336-10	6310	1	1	1				
9336-20	6310	1				3	1	

DASD for 3380 Measurements:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			
					TDSK	VSAM	VSE Sys.	VM Sys.
3380-A	3880-03	1	1		1			
3380-A	3880-03	1	1	1				
3380-A	3880-03	1				3	1	
3380-D	3880-03	1		1	1	3	1	1
3380-A	3880-03	1	1		1			1

Software Configuration

VSE version: 1.2.0
 VSE supervisor mode: 370

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	PRTY/ RELSHARE	RESERVED	Other Options
For VM/ESA Release 1.0 (370 Feature):						
VSE	1	VSE V=V	16MB/370	64	OFF	BMX ECMODE
SMART	1	VM/RTM	3MB/370	1	OFF	QDROP OFF FAVOR
For VM/ESA Release 2:						
VSE	1	VSE V=V	16MB/370	100	OFF	
SMART	1	RTM/ESA	4MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	6MB/XA	100	OFF	

Measurement Discussion

Two pairs of VSE V=V guest measurements were obtained for this study. The first pair of measurements compared VM/ESA Release 2 migration from the VM/ESA Release 1.0 (370 Feature) on 3380 (CKD) DASD. The second pair of measurements provided the same comparison except that 9336 (FBA) DASD was used. The hardware was configured so that the 3380 and 9336 systems were essentially identical to one another. Each 3380 volume had a corresponding 9336 volume with the same data on it, and each system had the same number of channel paths to the DASD volumes. This way, a comparison of 3380 versus 9336 performance can also be made. The PACEX4 workload was run on all measurements. All measurements had 64MB of real storage, which was more than ample to completely contain the VSE guest in storage. No paging occurred for any of the measurements.

The results of the measurements indicated that when migrating from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 2 on 3380 DASD, the internal throughput (ITR (H)) decreased 9.9% from 5.85 jobs per processor minute to 5.27. External throughput (ETR (H)) decreased 1.3% from 3.16 jobs per minute to 3.12. The equivalent measurements on 9336 DASD showed that internal throughput decreased 10.8% from 5.94 to 5.30, and external throughput decreased 0.5% from 1.98 to 1.97. Thus, the measurements indicated that migrating from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 2 on FBA DASD may experience a slightly lower (0.9%) relative performance difference in internal throughput than migrating from CKD systems. The primary contributing factor to the decrease in internal throughput for both pairs of measurements is the reduced processor speed when the 9221 is run in ESA mode.

External throughput was much lower for the 9336 measurements. This was expected because the hardware associated with the 9336 is slower than the corresponding 3380 hardware, and the PACEX4 workload is very I/O intensive. Because the PACEX4 workload is a set of 28 VSE batch jobs, the elapsed time for the 9336 measurements was much greater than the 3380 measurements. It was this elapsed time increase that caused the slightly less relative internal throughput for the 9336 measurements. VM/ESA entails a greater CP path length to enter into guest-enabled wait state than VM/370. Although the rate per second of entries into guest-enabled wait state (GUESTWT/SEC) was almost the same for both the VM/ESA Release 2 3380 and 9336 measurements, the increased run time for the 9336 measurement caused a much higher number of entries into guest-enabled wait per job (GUESTWT/CMD).

Fast path CCW translation support was extended to support FBA devices, resulting in a shorter CP path length in VM/ESA Release 2 than in the VM/ESA Release 1.0 (370 Feature). However, the reduced processor speed, along with the extra cost to enter guest-enabled wait state, resulted in a slightly higher CP time per job (CP/CMD (H)).

Note that there were more real I/Os per job (RIO/CMD (H,V)) in the 9336 measurements than in the 3380 measurements. This was consistent with VSE native measurements.

The measurement results are summarized in the following table. In the table, those metrics identified with an asterisk (*) are identical in meaning for both the VM/ESA Release 1.0 (370 Feature) and VM/ESA measurements. Exercise caution when comparing VM/ESA Release 1.0 (370 Feature) to VM/ESA using any of the remaining metrics. Although all these measures are closely analogous, their

exact meanings are not identical due to differences between these two operating systems and differences in how the data are collected and reduced. The derivation of each metric, for both VM/ESA Release 1.0 (370 Feature) and VM/ESA, is provided in the glossary.

DASD Type	3380 (CKD)		9336 (FBA)	
	VM/ESA 1.0 (370) HB74V33C	VM/ESA 2 HB54V335	VM/ESA 1.0 (370) HB74V33A	VM/ESA 2 HB54V330
Environment				
IML MODE	370	ESA	370	ESA
REAL STORAGE	64MB	64MB	64MB	64MB
EXP. STORAGE	0MB	0MB	0MB	0MB
VM MODE	370	370	370	370
VM SIZE	16MB	16MB	16MB	16MB
GUEST SETTING	V=V	V=V	V=V	V=V
VSE SUPERVISOR	370	370	370	370
Throughput (Min)				
Elapsed Time (H) *	532.1	538.6	847.4	854.9
ETR (H) *	3.16	3.12	1.98	1.97
ITR (H) *	5.85	5.27	5.94	5.30
ITR	5.86	5.29	5.93	5.31
ITRR (H) *	1.000	0.900	1.016	0.906
ITRR	1.000	0.902	1.012	0.906
Proc Usage (Sec)				
PBT/CMD (H) *	10.256	11.393	10.098	11.323
PBT/CMD	10.238	11.350	10.117	11.297
CP/CMD (H)	4.605	4.977	4.610	5.083
CP/CMD	4.313	4.232	4.264	4.274
EMUL/CMD (H)	5.651	6.416	5.488	6.240
EMUL/CMD	5.927	7.118	5.853	7.022
Processor Util.				
TOTAL (H) *	53.97	59.23	33.37	37.08
TOTAL	53.88	59.00	33.43	37.00
TOTAL EMUL (H)	29.74	33.35	18.13	20.44
TOTAL EMUL	31.19	37.00	19.34	23.00
TVR(H)	1.81	1.78	1.84	1.81
TVR	1.73	1.59	1.73	1.61
Storage				
NUCLEUS SIZE (V)	706KB	2460KB	704KB	2460KB
TRACE TABLE (V)	444KB	976KB	444KB	976KB
VSE WKSET (V)	836.5	1.9	876.25	1177
PGBLPGS	14708	14415	14708	14434
FREEPGS	1001	85	1001	72
FREE UTIL	na	0.57	na	0.49
SHRPGS	116	1007	159	986
Paging				
PAGE/CMD	0.000	0.000	0.000	0.000
XSTOR/CMD	na	0.000	na	0.000
FAST CLR/CMD	na	76.949	na	91.595
I/O				
VIO RATE	165.180	162.000	114.640	110.000
VIO/CMD	3138.762	3116.422	3469.295	3358.496
RIO RATE (H,V)	162.898	165.000	111.504	112.000
RIO/CMD (H,V)	3095.393	3174.134	3374.393	3419.560
DASD IO TOTAL (V)	82828	78824	94132	93521
DASD IO RATE (V)	162.41	164.22	112.06	111.33
DASD IO/CMD (V)	3086.09	3159.06	3391.28	3399.24
PRIVOPs				
PRIVOP/CMD (R)	5098.399	3127.795	5713.706	3384.687
DIAG/CMD (R)	174.915	248.311	214.478	286.374
SIE/CMD	na	7444.787	na	7755.073
SIE INTCPT/CMD	na	6030.277	na	6824.464
FREE TOTL/CMD	na	6906.146	na	7388.691
GUESTWT/CMD	na	1673	na	2588

Note: V=VMMAP(370)/VMPRF(ESA), H=Hardware Monitor, R=RTM, (H,V)=Hardware Monitor (370)/ VMPRF(ESA), Unmarked=VMMAP(370)/RTM(ESA), *=Identical meaning for VM/ESA Release 1.0 (370 Feature) and VM/ESA runs.

Table 31. VM/ESA Rel. 2 VSE Guest Migration on FBA and CKD DASD on 9221-170

9121-320 / VSE Guest

Workload: PACEX4

Hardware Configuration

Processor model: 9121-320⁵
 Processors used: 1
 Storage
 Real: 64MB
 Expanded: None
 Tape: 3480 (MONITOR)

DASD for 3370 Measurements:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			VM Sys.
					TDSK	VSAM	VSE Sys.	
3380-A	3880-03	1		1	1			1
3380-A	3880-03	1	1	1				
3380-D	3880-03	1	1		1			1
3380-A	3880-03	1	1		1			
3370	3880-01	1				3	1	
3370	3880-01	1				3	1	

DASD for 3380 Measurements:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			VM Sys.
					TDSK	VSAM	VSE Sys.	
3380	3880	1		1	1	3	1	1
3380	3880	1	1	1				
3380	3880	1	1		1			1
3380	3880	1	1		1			
3380	3880	1				3	1	

Software Configuration

VSE version: 1.2.0
 VSE supervisor mode: 370

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVED	Other Options
VSE	1	VSE V=V	16MB/370	100	OFF	
SMART	1	RTM/ESA	5MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	6MB/XA	100	OFF	

Measurement Discussion

Two V=V measurements were obtained to provide an indication of VM/ESA Release 2 VSE guest performance on FBA devices other than the integrated-adaptor/control-unit DASD such as the 9336. The first measurement was a PACEX4 measurement on a 9121-320 running VM/ESA Release 2 on 3380 (CKD)

⁵ See "Hardware Used" on page 25 for an explanation of how this processor model was defined.

DASD. The second measurement was the same as the first except 3370 (FBA) DASD were used. Each 3380 volume (that actually had I/O) had a corresponding 3370 volume with the same data on it, and each system had the same number of channel paths to the DASD.

The results of the measurements indicated that the external throughput rate (ETR (H)) for the 3370 measurement was 24% less than the external throughput rate for the 3380 measurement. This was expected because the 3370 is a slower device than the 3380. The internal throughput rate (ITR (H)) was actually 0.8% better for the 3370 measurement than the 3380 equivalent. The dominating factor in this improvement is that the VSE guest used less emulation time per job (EMUL/CMD (H)) in the 3370 measurement. As with the 9221-170 measurements, there were more I/Os for the FBA measurement than the CKD measurement.

DASD Type RELEASE RUN ID	3380 (CKD) VM/ESA 2 LB54V332	3370 (FBA) VM/ESA 2 LB54V331
Environment		
IML MODE	ESA	ESA
REAL STORAGE	64MB	64MB
EXP. STORAGE	0MB	0MB
VM MODE	370	370
VM SIZE	16MB	16MB
GUEST SETTING	V=V	V=V
VSE SUPERVISOR	ESA	ESA
Throughput (Min)		
Elapsed Time (H)	518.0	686.0
ETR (H)	3.24	2.45
ITR (H)	41.15	41.48
ITR	20.27	20.41
ITRR (H)	1.000	1.008
ITRR	1.000	1.007
Proc Usage (Sec)		
PBT/CMD (H)	2.916	2.893
PBT/CMD	2.960	2.940
CP/CMD (H)	1.007	1.074
CP/CMD	0.925	1.078
EMUL/CMD (H)	1.909	1.819
EMUL/CMD	2.035	1.862
Processor Util.		
TOTAL (H)	15.76	11.81
TOTAL	16.00	12.00
TOTAL EMUL (H)	10.32	7.42
TOTAL EMUL	11.00	7.60
TVR(H)	1.53	1.59
TVR	1.45	1.58
Storage		
NUCLEUS SIZE (V)	2420KB	2420KB
TRACE TABLE (V)	400KB	400KB
VSE WKSET (V)	1138	1184
PGBLPGS	14282	14284
FREEPGS	86	83
FREE UTIL	0.41	0.42
SHRPGS	918	918
Paging		
PAGE/CMD	0.000	0.000
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	92.500	73.500
I/O		
VIO RATE	175.000	139.000
VIO/CMD	3237.500	3405.500
RIO RATE (V)	166.000	137.000
RIO/CMD (V)	3071.000	3356.500
DASD IO TOTAL (V)	89341	90284
DASD IO RATE (V)	165.45	136.79
DASD IO/CMD (V)	3060.76	3351.45
PRIVOPs		
PRIVOP/CMD	3244.953	3419.521
DIAG/CMD	233.615	243.613
SIE/CMD	7252.000	7570.500
SIE INTCPT/CMD	6599.320	6964.860
FREE TOTL/CMD	7122.500	7276.500
GUESTWT/SEC	159	131
GUESTWT/CMD	2944	3208
Note: V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 32. VM/ESA Release 2 VSE Guest Comparison between FBA and CKD DASD on 9121-320

CP Configurability

Hardware Configuration

Processor model: 9021-340⁶
 Processors used: 1
 Storage
 Real: 256MB
 Expanded: 1GB

DASD:

<i>Type of DASD</i>	<i>Control Unit</i>	<i>Number of Paths</i>	<i>PAGE</i>	<i>SPOOL</i>	<i>- Number of Volumes -</i>			<i>System</i>
					<i>TDSK</i>	<i>User</i>	<i>Server</i>	
3390-2	3990-2	2						3
3390-3	3990-2	2	3	2	2			9
3390-2	3990-2	4	5	5	5			
3390-3	3990-2	4				8		
3390-3	3990-2	2				8		

Additional Information

Spool files: 30 021

Number of spool
pages in use: 131 184

Measurement Discussion

The generation of a VM/ESA Release 2 system can be simplified by externalizing the system definition data located in HCPSYS and the real I/O configuration data located in HCPRIO from the runnable code of the CP nucleus. This separation allows changes to the system definition data and real I/O configuration data without requiring updates to the HCPSYS and HCPRIO assemble files, reassembling them, building a nucleus, and then re-IPLing the system every time changes are made.

This study investigates the effects of the CP Configurability support on WARM START IPL and re-IPL.

For each IPL environment, certain variables were kept constant to improve measurement repeatability: real storage, expanded storage, processors online, trace table size, vectors enabled, spool files, hardware recognized by the system, all subsystem and DASD cache offline, and the directory. The HCPSYS and HCPRIO generated in the nucleus matched the system configuration file.

⁶ See "Hardware Used" on page 25 for an explanation of how this processor model was defined.

IPL Type	VM Release	CP Configurability Used	Elapsed Time (sec.)
WARM	VM/ESA 1.1	NA	90
	VM/ESA Rel. 2	NO	92
	VM/ESA Rel. 2	YES	85
re-IPL	VM/ESA 1.1	NA	81
	VM/ESA Rel. 2	NO	83
	VM/ESA Rel. 2	YES	74

Table 33. System IPL Comparison

IPL Type WARM START IPL was performed by issuing a load clear on the processor controller and selecting the default for both the type of start and TOD clock messages. Re-IPL was performed by entering the SHUTDOWN REIPL command on the operator's console.

Elapsed Time (sec) The average over three trials of the amount of time (in seconds) from the start of a warm start IPL or re-IPL to the time the operator user ID runs its profile.

As shown in the above table, the use of CP Configurability caused a decrease in elapsed time for the warm start IPL and re-IPL. This is partly due to the frame tables not being built as part of the CP nucleus. They are built in the dynamic paging area. This reduces the size of the CP nucleus, which reduces the time to read the nucleus at system IPL.

DIAGNOSE Code X'250'

VM/ESA Release 2 introduces the new DIAGNOSE code, X'250' as part of the extended support for VM Data Spaces. DIAGNOSE code X'250' performs block I/O synchronously or asynchronously between DASD and buffers in a virtual machine's primary address space or in a VM Data Space.

In addition to supporting VM Data Spaces, DIAGNOSE code X'250' provides significantly improved performance relative to the existing, IUCV-based block I/O support. A set of measurements was obtained to quantify these benefits. This section summarizes and discusses the results of those measurements.

Measurement Description

A test case was written to evaluate the block I/O function. A 4KB block size was used and the buffers were page-aligned. The blocks were written to, and subsequently read from, ascending contiguous block locations on 3390 DASD. This was done by way of a loop of 20 000 write requests followed by a loop of 20 000 read requests.

Different variations of this test case were produced to cover the following cases:

- IUCV-based block I/O (*BLOCKIO)
 - 1 block (using the single block interface)
 - 3 blocks
- DIAGNOSE code X'250'
 - 1 block, asynchronous
 - 1 block, synchronous
 - 3 blocks, asynchronous

Each of these test case variations was run from a single XA-mode virtual machine on a dedicated 9121-320 processor running VM/ESA Release 2. The QUERY TIME command was issued around the write loop and the read loop. This served as the primary source of performance information. Monitor data and hardware instrumentation data were also collected.

Results

The write loop and read loop results were obtained from the VTIME and TTIME information in the spooled console output of the QUERY TIME command. The results for each measured interval were divided by 20 000 to normalize by block I/O request. CP processor time was calculated as TTIME minus VTIME and converted to milliseconds. The following table summarizes these CP processor time results for *BLOCKIO and DIAGNOSE code X'250':

Case	*BLOCKIO	DIAG 250	Delta	%Delta
1: Read 1 block, no MDC	0.384	0.300	-0.084	-21.9
2: Read 3 blocks, no MDC	0.458	0.343	-0.115	-25.1
3: Write 1 block, no MDC	0.373	0.292	-0.081	-21.7
4: Write 3 blocks, no MDC	0.451	0.332	-0.119	-26.4
5: Write 1 block, MDC	0.415	0.329	-0.086	-20.7
6: Write 3 blocks, MDC	0.554	0.439	-0.115	-20.8

Table 34. DIAGNOSE Code X ζ 250 ζ Performance - Comparison to *BLOCKIO. CP processor time (in ms) per block I/O request, 9121-320, asynchronous interface

DIAGNOSE code X ζ 250 ζ achieves its performance benefit by greatly reducing the CP path length required to get to and from the CP block I/O routines. The path lengths of the block I/O routines themselves are only slightly affected by the DIAGNOSE code X ζ 250 ζ support.

For requests involving large numbers of blocks, expect the percentage reduction in CP processing requirements to be much smaller than the 21% to 26% range shown in the above table. This is because the processing required to communicate a block I/O request to CP is fairly constant, while the processing required to do that request is roughly proportional to the number of blocks involved. Because DIAGNOSE code X ζ 250 ζ only improves the communications portion of block I/O processing, the DIAGNOSE code X ζ 250 ζ improvement becomes a smaller fraction of total block I/O processing as the number of blocks in the request increases.

When minidisk caching is not a factor, the percentage improvement is significantly larger for the 3 block case than the 1-block case (compare cases 1 and 2, and cases 3 and 4). This is because, in the 1-block case, *BLOCKIO benefits from a separate interface that is optimized for 1 block, while DIAGNOSE code X ζ 250 ζ has no such special interface. *BLOCKIO loses this advantage when multiple blocks are involved.

This same factor is at work in a minidisk caching environment (cases 5 and 6). This is indicated by the fact that the increased absolute improvement going from case 5 to case 6 (write, with MDC: 0.086 ms to 0.115 ms) is quite similar to the increased improvement going from case 3 to case 4 (write, no MDC: 0.081 ms to 0.119 ms). The reason for the smaller *percentage* improvement associated with case 6 is that this similar absolute difference is divided by a bigger base value (0.554 ms). Writes take more processor time with MDC in effect because of the added processing required to insert the blocks into the cache.

The processor usage reported by QUERY TIME is fairly accurate. However, it does not include "system time" that CP cannot logically attribute to a given user. The monitor data showed that the system time for each DIAGNOSE code X ζ 250 ζ measurement was similar to the system time for the corresponding *BLOCKIO measurement. In normal operating environments, system time is fairly small. For example, for the CMS-intensive measurements shown in section "Measurement Variability" on page 173, system time was about 8% of total CP processor time.

DIAGNOSE code X ζ 250 ζ provides both an asynchronous and a synchronous interface. The synchronous interface is useful when the virtual machine does not choose to overlap the I/O with other processing. Synchronous invocation is

easier to implement and requires substantially less processor time, as illustrated by case 2 in the following table:

Case	ASYNCH	SYNCH	Delta	%Delta
1: Read 1 block, MDC	0.156	0.153	-0.003	-1.9
2: Write 1 block, MDC	0.329	0.226	-0.103	-31.3

Table 35. DIAGNOSE Code X'250' Performance - Synchronous Interface Benefits. CP processor time (in ms) per block I/O request, 9121-320

DIAGNOSE code X'250' allows the application to specify whether the minidisk cache, if present, should be used or bypassed. Our test case application requested that the minidisk cache be used. On a read request, when all requested blocks are found in the minidisk cache and the application has requested that the cache be used, DIAGNOSE code X'250' will do an immediate synchronous return even if asynchronous processing was specified. That is the situation with case 1, which therefore showed no significant additional improvement when synchronous processing was specified.

Because case 2 is a write operation, this "immediate return" minidisk cache optimization does not apply. Therefore, the asynchronous request was truly handled asynchronously. Case 2, then, shows the large reduction in CP processor usage that occurs when going from asynchronous to synchronous block I/O request processing. CP processor usage is reduced because CP does not have to reflect an external interrupt to the requesting virtual machine when the I/O completes.

Synchronous processing also reduces processor usage in the virtual machine. This is because there is no CMS processing required to wait for I/O completion and to handle the reflected external interrupt. For case 2, emulation processor time per block I/O went from 0.058 ms to 0.006 ms—a 90% decrease.

When DIAGNOSE code X'250' is used synchronously, it is functionally quite similar to DIAGNOSE code X'A4'. Measurement results have shown that DIAGNOSE code X'A4' offers somewhat better performance. DIAGNOSE code X'A4', then, is a better choice if it meets the application's functional requirements. DIAGNOSE code X'250' is more appropriate if the application needs to do block I/O directly into data spaces or if the application would benefit from its asynchronous capability.

*BLOCKIO is basically an asynchronous interface. The only exception to this is that, when all requested blocks in a read request are found in the minidisk cache, the application can optionally choose to have CP do an immediate, synchronous return. This case was not part of the evaluation matrix. However, past measurements indicate that, in this synchronous case, *BLOCKIO can be expected to benefit relative to asynchronous usage by about the same absolute amount as was measured here for DIAGNOSE code X'250' (case 2).

Extended CMS File System Interfaces

This support provides an interface to the minidisk file system for high-level languages through callable services library (CSL) routines. It has extended certain existing CSL routines so that they work on both SFS and minidisks. This allows a high-level language application to have single-path code that provides the function available to current minidisk applications and also allows for SFS exploitation when desired.

A path length comparison was made between a selection of these CSL routines and the FS macros that perform the same functions. For example, DMSERASE was compared to FSERASE. The scenario traced consisted of:

1. Open a nonexistent minidisk file for write
2. Write two records
3. Close
4. Open the file for read
5. Read a record
6. Close
7. Erase the file

The following table summarizes the path length comparison for each step of the scenario.

Function	FS Macro	CSL Routine
OPEN (write)	1527	2782
WRITE 1	1010	1285
WRITE 2	531	753
CLOSE	3421	3454
OPEN (read)	1587	3117
READ	440	703
CLOSE	1120	1079
ERASE	2734	3845
Total	12370	17018

Table 36. Path length comparison between FS macros and CSL functions.

There was an overall 38% increase in path length to perform the scenario using CSL routines as compared to using FS macros. The increased path length for CSL was due to the linkage overhead of locating the function routine. When opening a file, the CSL routines have to determine whether a namedef, dirid or file mode letter was specified. This cost was 822 instructions. Erase incurred the same overhead as seen with open.

Tuning Considerations

This chapter discusses VM performance tuning topics:

- The default CMS SFS file cache size increased from 12KB to 20KB in VM/ESA Release 2. The performance effects are examined in “CMS File Cache for SFS” on page 146.
- Although no changes occurred to I/O assist in VM/ESA Release 2, its performance characteristics are explained in “I/O Assist for Guests” on page 150.

CMS File Cache for SFS

Workload: FS7BMAXR

Hardware Configuration

Processor model: 9121-480
 Processors used: 2
 Storage
 Real: 192MB
 Expanded: 64MB (all retained for MDC)
 Tape: 3480 (MONITOR)

DASD:

<i>Type of DASD</i>	<i>Control Unit</i>	<i>Number of Paths</i>	<i>PAGE</i>	<i>SPOOL</i>	<i>- Number of Volumes -</i>			<i>System</i>
					<i>TDSK</i>	<i>User</i>	<i>Server</i>	
3390-2	3990-2	4	8	4	4			
3390-2	3990-3	4					10 R	
3390-2	3990-3	2	4 R	2 R	2 R			2 R

Note: *R* or *W* next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

<i>Control Unit</i>	<i>Number</i>	<i>Lines per Control Unit</i>	<i>Speed</i>
3088-02	1	NA	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

<i>Virtual Machine</i>	<i>Number</i>	<i>Type</i>	<i>Machine Size/Mode</i>	<i>RELSHARE</i>	<i>RESERVE</i>	<i>Other Options</i>
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	400	QUICKDSP ON
ROSERV1	1	SFS	32MB/XC	1500	OFF	QUICKDSP ON
RWSERV1	1	SFS	32MB/XC	1500	1000	QUICKDSP ON
RWSERV2	1	SFS	32MB/XC	1500	1000	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	100	OFF	
Users	1560	User	3MB/XC	100	OFF	

Measurement Discussion

On sequential I/O operations, CMS reads ahead and writes behind. The file cache size determines the maximum number of blocks that can be transferred at a time. For minidisk, the file cache size is 8KB. For SFS, the file cache size is specified during CMS generation and can range from 1KB to 96KB. The default was changed from 12KB to 20KB in VM/ESA Release 2. This section evaluates this change.

A VM/ESA Release 2 measurement was made on the 9121-480 with the file cache size set at 12KB, the VM/ESA Release 1.1 default, for comparison. Increasing the file cache size had a small effect on external response time with little effect on other system metrics such as internal throughput rate and paging rate.

The external response time (AVG LAST(T)) decreased somewhat. This is primarily due to the fact that many of the compiler work files are larger than the three block limit of the 12KB cache. As a result, increasing the file cache size to 20KB had the effect of reducing the number of server calls and server I/O requests that were required. This can be seen by the reduction in FP REQ/CMD(Q) of 9.3% and IO/CMD (Q) of 4.5%.

In principle, the file cache size represents a trade-off between SFS server efficiency and paging. In this comparison however, the increased file cache size did not drive up the paging rate. Apparently, the increased number of referenced buffer pages was balanced by the fact that they were held for a shorter time because the I/Os were done more efficiently.

FILE CACHE SIZE RELEASE RUN ID	12KB VM/ESA 2 L25M1565	20KB VM/ESA 2 L25M1564
Environment		
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
USERS	1560	1560
VTAMs	1	1
VSCSs	0	0
Response Time		
TRIV INT	0.124	0.128
NONTRIV INT	0.608	0.608
TOT INT	0.450	0.454
TOT INT ADJ	0.400	0.394
AVG FIRST (T)	0.312	0.286
AVG LAST (T)	0.512	0.486
Throughput		
AVG THINK (T)	25.47	25.58
ETR	50.20	49.08
ETR (T)	56.51	56.49
ETR RATIO	0.888	0.869
ITR (H)	62.01	61.92
ITR	27.57	26.94
EMUL ITR	41.45	40.20
ITRR (H)	1.000	0.999
ITRR	1.000	0.977
Proc. Usage		
PBT/CMD (H)	32.252	32.301
PBT/CMD	32.206	32.217
CP/CMD (H)	11.537	11.397
CP/CMD	10.794	10.621
EMUL/CMD (H)	20.716	20.903
EMUL/CMD	21.412	21.596
Processor Util.		
TOTAL (H)	182.26	182.47
TOTAL	182.00	182.00
UTIL/PROC (H)	91.13	91.24
UTIL/PROC	91.00	91.00
TOTAL EMUL (H)	117.07	118.08
TOTAL EMUL	121.00	122.00
MASTER TOTAL (H)	90.96	91.04
MASTER TOTAL	91.00	91.00
MASTER EMUL (H)	54.07	54.50
MASTER EMUL	56.00	56.00
TVR(H)	1.56	1.55
TVR	1.50	1.49
Storage		
NUCLEUS SIZE (V)	2304KB	2304KB
TRACE TABLE (V)	400KB	400KB
WKSET (V)	76	76
PGBLPGS	40620	40629
PGBLPGS/USER	26.0	26.0
FREEPGS	4308	4276
FREE UTIL	0.90	0.90
SHRPGS	1358	1377
Paging		
READS/SEC	501	492
WRITES/SEC	342	340
PAGE/CMD	14.917	14.728
PAGE IO RATE (V)	143.400	140.700
PAGE IO/CMD (V)	2.538	2.491
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	6.600	6.709

FILE CACHE SIZE RELEASE RUN ID	12KB VM/ESA 2 L25M1565	20KB VM/ESA 2 L25M1564
Environment		
REAL STORAGE	192MB	192MB
EXP. STORAGE	64MB	64MB
USERS	1560	1560
VTAMs	1	1
VSCSs	0	0
Queues		
DISPATCH LIST	40.49	41.09
ELIGIBLE LIST	0.00	0.00
I/O		
VIO RATE	531	530
VIO/CMD	9.396	9.382
RIO RATE (V)	345	344
RIO/CMD (V)	6.105	6.089
MDC READS	267	266
MDC WRITES	113	118
MDC MODS	74	78
MDC HIT RATIO	0.87	0.86
PRIVOPs		
PRIVOP/CMD	25.167	24.666
DIAG/CMD	23.803	24.027
DIAG 08/CMD	0.814	0.814
DIAG 14/CMD	0.018	0.018
DIAG 58/CMD	1.256	1.239
DIAG 98/CMD	1.221	1.257
DIAG A4/CMD	2.035	2.053
DIAG A8/CMD	1.716	1.699
DIAG 214/CMD	12.652	12.763
SIE/CMD	61.050	60.806
SIE INTCPT/CMD	43.346	43.172
FREE TOTL/CMD	74.693	73.286
VTAM Machines		
WKSET (V)	408	407
TOT CPU/CMD (V)	4.0405	4.0714
CP CPU/CMD (V)	1.7794	1.7899
VIRT CPU/CMD (V)	2.2611	2.2816
DIAG 98/CMD (V)	1.222	1.259
SFS Servers		
WKSET (V)	1578	1611
TOT CPU/CMD (V)	3.8144	3.6977
CP CPU/CMD (V)	1.8974	1.7899
VIRT CPU/CMD (V)	1.9170	1.9079
FP REQ/CMD(Q)	1.388	1.270
IO/CMD (Q)	1.885	1.804
IO TIME/CMD (Q)	0.034	0.032
SFS TIME/CMD (Q)	0.060	0.054
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Query Filepool Counters, Unmarked=RTM		

Table 37. SFS - CMS File Cache Size Comparison: 9121-480

I/O Assist for Guests

Although the I/O interpretation facilities are not changed in VM/ESA Release 2, this section discusses how it works, the performance benefit, and the associated performance data.

Description

The start-interpretive-execution assist (commonly referred to as SIE assist, I/O assist, and I/O passthru) is available for dedicated devices of V=R guests on most ESA/370* and ESA/390* processors. On processors with the PR/SM* feature, VM/ESA also allows I/O interpretation for dedicated devices of V=F guests. With this feature, some I/O instructions (for example, SIOF, SSCH, RSCH, TSCH, HSCH, and TPI) are handled by the processor and do not cause an interception of the SIE instruction. Under normal circumstances, even the interruption for these I/Os will go directly to the guest machine and will not cause an interruption of the SIE instruction. Processing of the I/O interruptions for an undispached guest depends on the processor support. On processors without PR/SM, CP enables for interruptions, processes the real interruption, and reflects it to the guest. On processors with PR/SM, the processor can usually queue and present these interruptions directly to the guest. CP dispatches the guest when notified that the interruption is pending.

To be eligible for this I/O assist feature, a device must:

- be I/O assist capable. Most devices are I/O assist capable. However, devices with a class of SWITCH or TERMINAL, 3203 printers, and tapes with dynamic pathing for a 370-mode guest are not supported.
- be dedicated to a preferred guest in read/write mode. A device attached as read-only never qualifies. A minidisk never qualifies.
- have a virtual subchannel number equal to the real subchannel number when dedicated to an ESA-mode guest.
- have a virtual device number less than or equal to 0FFF when dedicated to a 370-mode guest.

There are certain conditions for which the entire guest or an individual device is either temporarily or permanently (that is, until the next guest IPL) taken out of I/O assist.

The reasons why an entire guest must be taken out of I/O assist include:

- The command SET IOASSIST OFF is issued for the guest.
- The command SET CCWTRANS ON (or equivalently, SET NOTRANS OFF) is issued for the V=R guest.
- A trace trap of any kind is activated by the CP TRACE command.
- One or more virtual processors are stopped (Examples: the CP STOP command, PA1 with SET RUN OFF, guest loading a disabled wait PSW).
- A 370-mode guest has a BC-mode PSW with I/O interruptions enabled. This disqualifies the guest until the next guest IPL.
- The guest issues DIAGNOSE I/O. DIAGNOSE codes X'18', X'20', X'A4', and X'A8' disqualify the guest until the next guest IPL.

- The guest becomes undispached. On processors without PR/SM, CP enables for interruptions, processes the real interruption, and reflects it to the guest. On processors with PR/SM, the processor can usually queue and present these interruptions directly to the guest. CP dispatches the guest when notified that the interruption is pending.
- An individual device must be taken out of I/O assist and the processor does not have PR/SM support. On processors without PR/SM, when VM/ESA needs to take one device out of I/O assist, it takes the entire guest out (including all devices) and when that I/O operation completes (for example, after the last I/O operation in progress gives final status), it puts all the eligible devices back into I/O assist. On processors with PR/SM, only the devices that need to be taken out of I/O assist are affected.
- An ESA-mode guest is selectively enabled for some, but not all, of the virtual interruption subclasses (ISCs) to which its devices are assigned, in a way that is incompatible with the VM/ESA I/O assist support. The original I/O assist support for processors without PR/SM expected the guest to use only interruption subclasses 3, 5, and 7 in normal operation. The PR/SM support is more general, but selective enablement can still cause I/O interruptions to be taken and queued in CP.
- The guest is using address-limit-checking (SAL instruction). This would ordinarily be used only by a VM/XA or VM/ESA guest running its own V=R second level guest.
- The guest has specified an invalid measurement-block origin for an SCHM instruction. CP must simulate the architected error conditions at the time of the I/Os.
- The guest appears to be in MVS's disabled console communications function (DCCF). This is a special case of selective enablement that requires special treatment to prevent a potential console lockout.

The reasons for which an individual device must be taken out of I/O assist include:

- The entire guest was taken out of I/O assist.
- CP-initiated I/O: This can arise for several reasons, including commands (for example, HALT, VARY, VARY PATH, QUERY rdev STATUS), 3990 state change events, and errors that trigger dynamic path management actions to restore path groups at the control unit. The device is temporarily taken out of I/O assist for this operation.
- The guest attempts to issue restricted commands (for example, 3990 subsystem control functions to manage cache and duplexing): These must be simulated to ensure the guest is authorized to perform the requested operation. The device is temporarily taken out of I/O assist for this operation.
- Unit check status from the device: CP drains sense data immediately to prevent suspension of CP's or other guests' I/Os to other devices on the control unit. Also, some conditions (for example 3990 state change notifications) require special CP handling. CP must then keep the device out of I/O assist long enough to simulate the guest's sense.
- Asynchronous message from the control unit: This message may report status pertaining to the control unit or entire subsystem. CP must receive it to determine how all the devices are affected. The device is temporarily taken out of I/O assist for this case.

- The secondary volume of a 3990 duplexed pair is not assisted. (The guest should not be directing any I/Os to it anyway, but CP needs to react if it does.)

Performance data

The performance data available to VM/ESA for a device in I/O assist depends on whether the guest has issued a valid SCHM (Set Channel Monitor) instruction to start the measurement-block-update facility. If the guest does not issue an SCHM, then all performance data will be collected in CP's measurement-blocks. If the guest issues an SCHM, it is simulated by CP, and a real SCHM is issued to set up the measurement-block-update facility, replacing CP's measurement-block with the guest's measurement-block. Thus, any I/O successfully completed in I/O assist will be known only to the guest. For devices taken out of I/O assist, CP's measurement-block is reinstated and updated by the measurement-block-update facility. These updates are reflected to the guest's measurement-block by CP.

Table 38 shows the source of the update for CP's measurement-blocks and the guest's measurement-blocks by the type of I/O operation.

Type of I/O	Update source by measurement-block owner	
	CP's	Guest's
CP I/O	channel subsystem	none
nonassisted I/O and without a guest issued SCHM	channel subsystem	none
assisted I/O and without a guest issued SCHM	channel subsystem	none
nonassisted I/O and with a guest issued SCHM	channel subsystem	by CP at virtual interrupt reflection
assisted I/O and with a guest issued SCHM	none	channel subsystem

Table 38. Measurement-Block Updates for I/O

The information from CP's measurement-blocks is contained in monitor domain 6, record 3. The information from the guest's measurement-blocks is not in any VM/ESA monitor data. VM/ESA collects information in CP control blocks relative to transitions between assist and nonassist on a device basis. This information is collected in monitor domain 6, record 3.

Some cache control units collect data similar to the data collected by the measurement-block-update facility. This data is retrieved by VM/ESA and reported in monitor domain 6, record 4, even for devices that are in I/O assist. VMPRF uses this data to produce the PRF016 CACHE_DASD_BY_ACTIVITY and PRF095 DASD_BY_ACTIVITY_EF reports. I/Os to retrieve this information are counted in the CP measurement-block by the channel subsystem but are not included in the count retrieved from the control unit.

The relevant flags for the current I/O assist status of any device is kept in the virtual device block. These flags can be viewed on a running system by using an EXEC.

Examples

A sample of various reports is included for 1 selected MVS and 2 selected VSE guest measurements published in this report.

For the MVS guest measurement 225MVSD4 (detailed in Table 15 on page 90), the following reports are included for 4 selected devices.

- MVS/ESA RMF Direct Access Device Activity (Figure 12 on page 154): This report contains information collected in the guest's measurement-block.
- RTM VM/ESA Device Screen (Figure 13 on page 154): This report contains information collected in CP control blocks.
- VMPRF DASD_BY_ACTIVITY (Figure 14 on page 154): VMPRF derives the values for this report from monitor domain 6, record 3 data. These monitor records contains information collected in CP's measurement-blocks.
- VMPRF DASD_IO_ASSIST (Figure 15 on page 155): This report contains information collected in CP control blocks. VMPRF APAR VM52375 replaced several values in this report with "*****" because of anomalies with the frequency of monitor sampling versus the frequency of CP updates to the counters. CP updates timers and counters when the state of a device changes, while monitor takes a snapshot of counters on a fixed interval. VMPRF derives the values for this report from monitor domain 6, record 4 data.
- VMPRF DASD_BY_ACTIVITY_EF (Figure 16 on page 155): CP obtains the data in this record by issuing the Perform Subsystem Function command to execute Read Subsystem Data on a 3990-3 control unit.

Because MVS issues an SCHM, the assisted I/Os are not counted in CP's measurement-block.

The sample RTM data (Figure 13 on page 154) shows all 4 devices in I/O PASSTHRU.

Device 0D23 is not attached through a cache control unit and has no I/Os from any source. Thus, all sample reports show the expected count of 0.

Device 0810 is attached through a cache control unit and has only CP-initiated I/O to retrieve statistics from the cache control unit. These are the 41 I/Os shown in the sample VMPRF DASD activity data (Figure 14 on page 154) that are only counted in CP's measurement-block. The sample RMF data (Figure 12 on page 154) and the sample VMPRF extended function data (Figure 16 on page 155) show the expected count of 0. The device was out of assist at the start of the measurement and remained out of assist until the last monitor interval of the measurement. The 1 transition shown in the sample VMPRF I/O assist data (Figure 15 on page 155) occurred during this last monitor interval of the measurement.

Device 0D9E is not attached through a cache control unit and has both assisted and nonassisted I/Os from MVS. The I/O activity rate of 2.860 shown in the sample RMF data (Figure 12 on page 154) represents all of the MVS-initiated I/Os. The 27 I/Os shown in the sample VMPRF DASD activity data (Figure 14 on

page 154) are the nonassisted I/Os that are counted in CP's measurement-block and then reflected to the MVS measurement-block. The 27 transitions shown in the sample VMPRF I/O assist data (Figure 15 on page 155) implies 1 transition for each nonassisted I/O.

Device 09A2 is attached through a cache control unit. It has both assisted and nonassisted I/Os from MVS, as well as CP-initiated I/Os to retrieve statistics from the cache control unit. The I/O rate of 8.786 shown in the sample RMF data (Figure 12) and the I/O activity rate of 8.7 shown in the sample VMPRF extended function data (Figure 16 on page 155) represent all of the MVS-initiated I/Os. The I/O count of 186 shown in the sample VMPRF DASD activity data (Figure 14) includes the CP-initiated I/Os and the nonassisted MVS-initiated I/Os. The count of 80 transitions shown in the sample VMPRF I/O assist data (Figure 15 on page 155) is smaller than number of nonassisted I/Os and thus multiple I/Os were completed during a single transition.

DIRECT ACCESS DEVICE ACTIVITY																		PAGE 1	
MVS/ESA SP4.2.2				SYSTEM ID UNKN RPT VERSION 4.2.2				DATE 10/15/92 TIME 15.40.20				INTERVAL 19.23.738 CYCLE 1.000 SECONDS							
TOTAL SAMPLES = 1,163 IODF = ?? NO CREATION INFORMATION AVAILABLE																			
STORAGE GROUP	DEV NUM	DEVICE TYPE	VOLUME SERIAL	LCU	ACTIVITY RATE	AVG RESP TIME	AVG IOSQ	AVG DPB DLY	AVG CUB DLY	AVG DB DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	AVG NUMBER ALLOC	% ANY ALLOC	% MT PEND
	D23	3380K	CBW2PK		0.000	0	0		0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
	810	3380K	PAGEF1		0.000	0	0		0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0	1.0	100.0	0.0
	D9E	3380K	USERPK		2.860	27	7		0.0	0.7	0.8	12.7	6.9	1.96	5.60	0.0	85.9	100.0	0.0
CHGD	9A2	3380K	DOA001		8.786	64	20		0.0	0.4	0.6	14.4	29.1	25.57	38.20	0.0	0.6	100.0	0.0

Figure 12. MVS/ESA RMF Direct Access Device Activity report.

```

<>VM/ESA CPU9021 SERIAL 310376 1017M DATE 10/15/92 START 14:40:22 END 14:59:34<>
<--- DEVICE ---> <----- DEVICE RDEV DATA -----> <-- MEASUREMENT FACILITY -->
0D23 3380 CBW2PK     DEVICE IN I/O PASSTHRU SINCE DATA LAST RESET
0810 3380 PAGEF1    DEVICE IN I/O PASSTHRU SINCE DATA LAST RESET
0D9E 3380 USERPK    DEVICE IN I/O PASSTHRU SINCE DATA LAST RESET
09A2 3380 DOA001    DEVICE IN I/O PASSTHRU SINCE DATA LAST RESET
    
```

Figure 13. RTM VM/ESA Device Screen.

PRF012 Run 10/21/1992 14:38:21																		DASD_BY_ACTIVITY		Page 1
From 10/15/1992 14:40:11																		DASD Activity Ordered by Activity		Sample
To 10/15/1992 15:00:13																				CPU 9021 SN 10376
For 1202 Secs 00:20:02																		MVS guest measurement 225MVSD4		VM/ESA 20.01 SLU 0000
<-----Device----->				<-----SSCH+RSCH----->				<-----Time----->				<-SSCHs in>				<--Queue-->				
Num-ber	Volume Serial	Control Unit	Owner	Mini-disk Links	On-line Secs	Count	Rate	Plus Avoided	Plus Rate	Pct Busy	Pct Pend	Disc	Conn	Serv	Resp	Mean	Max Err			
0D23	CBW2PK 3380-K	3880-03	MVS	0	1202	0	0	0	0	0	0	0	0	0	0	0	0			
0810	PAGEF1 3380-K	3990-3	MVS	0	1202	41	0.0	41	0.0	0.0	0.0	0.0	0.8	0.9	0.9	0	0			
0D9E	USERPK 3380-K	3880-03	MVS	0	1202	27	0.0	27	0.0	0.0	0.2	0.0	0.3	0.5	0.5	0	0			
09A2	DOA001 3380-K	3990-3	MVS	0	1202	186	0.2	186	0.2	0.3	3.1	1.5	11.7	16.2	16.2	0	0			

Figure 14. VMPRF DASD_BY_ACTIVITY.

```

PRF098 Run 10/21/1992 14:38:21          DASD_IO_ASSIST                      Page 2
                                           SIE Assist Activity by Device

From 10/15/1992 14:40:11                Sample
To   10/15/1992 15:00:13                CPU 9021          SN 10376
For  1202 Secs 00:20:02                  MVS guest measurement 225MVSD4      VM/ESA 20.01 SLU 0000
    
```

<-----Device----->				<Eligible>		<--Transition-->			<---Percentage of Time--->		
Num-ber	Volume Serial Type	Control Unit	Owner	XA	370	On-line	Count	Rate	Inside IO	Leaving IO	Outside IO
0D23	CBW2PK 3380-K	3880-03	MVS	YES	YES	1202	0	*****	*****	*****	*****
0810	PAGEF1 3380-K	3990-3	MVS	YES	YES	1202	1	*****	*****	*****	*****
0D9E	USERPK 3380-K	3880-03	MVS	YES	YES	1202	27	*****	*****	*****	*****
09A2	DOA001 3380-K	3990-3	MVS	YES	YES	1202	80	*****	*****	*****	*****

Figure 15. VMPRF DASD_IO_ASSIST.

```

PRF095 Run 10/21/1992 14:38:22          DASD_BY_ACTIVITY_EF                      Page 3
                                           DASD Extended Function Activity Ordered by Activity

From 10/15/1992 14:40:11                Sample
To   10/15/1992 15:00:13                CPU 9021          SN 10376
For  1202 Secs 00:20:02                  MVS guest measurement 225MVSD4      VM/ESA 20.01 SLU 0000
    
```

<-----Device----->				<-----Rate----->				<-----Percent----->															
Num-ber	Volume Serial Type	Stg Ctlr ID	Control Unit	C A	D F	D U	Total I/O	Read NonSq	Read Seq	Write Seq	<-----Hits----->						DFW Tot	CFW Tot	Cache I/O	Norm Stge	Seq Stge	De- Stge	DFW By-pass
											W L	ST	FW	Read	Tot	Read							
0810	PAGEF1 3380-K	000C	3990-3	A	A	A	0.0	0.0	0	0	100	67	67	0	0	0	0	0	100	0	0	0	0
09A2	DOA001 3380-K	003A	3990-3	A	A	A	8.7	0.2	0	0.1	61	99	99	98	98	0	100	0	3	0	0	0	0

Figure 16. VMPRF DASD_BY_ACTIVITY_EF.

For the VSE measurement LB58REX2 (detailed in Table 18 on page 98), the following reports are included for 2 selected devices.

- RTM VM/ESA Device Screen (Figure 17 on page 157): This report contains information collected in CP control blocks.
- VMPRF DASD_BY_ACTIVITY (Figure 18 on page 157): VMPRF derives the values for this report from monitor domain 6, record 3 data. These monitor records contains information collected in CP's measurement-blocks.
- VMPRF DASD_IO_ASSIST (Figure 19 on page 157): This report contains information collected in CP control blocks. VMPRF APAR VM52375 replaced several values in this report with "*****" because of anomalies with the frequency of monitor sampling versus the frequency of CP updates to the counters. CP updates timers and counters when the state of a device changes, while monitor takes a snapshot of counters on a fixed interval.

No devices dedicated to the VSE guest were attached through a cache control unit and thus a sample of the VMPRF DASD extended function data is not included.

Because an SCHM instruction was not issued by the VSE guest, all of the I/O activity was collected in CP's measurement-blocks. The sample VMPRF DASD activity data thus reflects all of the I/Os.

The sample RTM data (Figure 17 on page 157) shows both devices in I/O PASSTHRU.

For device 03C0, the I/O count of 14992 shown in the sample VMPRF DASD activity data (Figure 18 on page 157) represents all of the VSE-initiated I/Os. The count of 972 transitions shown in the sample VMPRF I/O assist data (Figure 19 on page 157) represents the number of times this device was temporarily taken out of I/O assist. There is no counter that represents the number of I/Os that were completed while the device was temporarily out of assist.

For device 07D0, the I/O count of 2917 shown in the sample VMPRF DASD Activity data (Figure 18 on page 157) represents all of the VSE-initiated I/Os. The count of 0 transitions shown in the sample VMPRF I/O assist data (Figure 19 on page 157) indicates the device maintained the same I/O assist status throughout the measurement period. In this case, device 07D0 was in I/O assist at the start of the measurement and remained in I/O assist for the duration of the measurement. However, a device that was out of I/O assist at the start of the measurement and remained out of I/O assist for the duration of the measurement would look the same.

```
<>VM/ESA CPU9121 SERIAL 121207 190M DATE 09/14/92 START 17:25:55 END 17:35:05<>
<--- DEVICE ---> <----- DEVICE RDEV DATA -----> <-- MEASUREMENT FACILITY -->
DEV TYPE VOLSER IOREQST SEC %Q %ER R %LK LNK PA %UT ACC FPT DCT CN %CN
03C0 3380 DOSRES DEVICE IN I/O PASSTHRU SINCE DATA LAST RESET
07D0 3380 SYSWK1 DEVICE IN I/O PASSTHRU SINCE DATA LAST RESET
```

Figure 17. RTM VM/ESA Device Screen.

```
PRF012 Run 09/28/1992 14:59:28 DASD_BY_ACTIVITY Page 1
DASD Activity Ordered by Activity
From 09/14/1992 17:25:58 Sample
To 09/14/1992 17:34:58 CPU 9121 SN 21207
For 540 Secs 00:09:00 VSE guest measurement LB58REX2 VM/ESA 20.01 SLU 0000
```

```
<-----Device-----> <-----SSCH+RSCH-----> <-----Time-----> <--SSCHs in-->
<--Queue-->
Num- Volume Control Mini- On- Plus
ber Serial Type Unit Owner disk line Count Rate Avoided Rate Busy Pend Disc Conn Serv Resp Mean Max Err
03C0 DOSRES 3380-A 3880-03 VSEVR 0 540 14992 27.8 14992 27.8 61.2 0.2 18.7 3.2 22.0 22.0 0 0 0
07D0 SYSWK1 3380-A 3880-03 VSEVR 0 540 2917 5.4 2917 5.4 11.6 0.1 19.4 1.9 21.4 21.4 0 0 0
```

Figure 18. VMPRF DASD_BY_ACTIVITY.

```
PRF098 Run 09/28/1992 14:59:28 DASD_IO_ASSIST Page 2
SIE Assist Activity by Device
From 09/14/1992 17:25:58 Sample
To 09/14/1992 17:34:58 CPU 9121 SN 21207
For 540 Secs 00:09:00 VSE guest measurement LB58REX2 VM/ESA 20.01 SLU 0000
```

```
<-----Device-----> <Eligible> <--Transition--> <---Percentage of Time--->
On- Inside Leaving Outside
line IO IO IO
Num- Volume Control XA 370 Secs Count Rate Assist Assist Assist
ber Serial Type Unit Owner YES YES 540 972 ***** ***** ***** *****
03C0 DOSRES 3380-A 3880-03 VSEVR YES YES 540 0 ***** ***** ***** *****
07D0 SYSWK1 3380-A 3880-03 VSEVR YES YES 540 0 ***** ***** ***** *****
```

Figure 19. VMPRF DASD_IO_ASSIST.

For the VSE measurement VR1170DA (detailed in Table 51 on page 207), the following reports are included for 2 selected devices.

- EXPLORE/VSE** XA DEVICE STATISTICS (Figure 20 on page 158): This report contains information collected in VSE control blocks and in the guest's measurement-block.
- RTM VM/ESA Device Screen (Figure 21 on page 158): This report contains information collected in CP control blocks.
- VMPRF DASD_BY_ACTIVITY (Figure 22 on page 159): VMPRF derives the values for this report from monitor domain 6, record 3 data. These monitor records contains information collected in CP's measurement-blocks.
- VMPRF DASD_IO_ASSIST (Figure 23 on page 159): This report contains information collected in CP control blocks. VMPRF APAR VM52375 replaced several values in this report with "*****" because of anomalies with the frequency of monitor sampling versus the frequency of CP updates to the counters. CP updates timers and counters when the state of a device changes, while monitor takes a snapshot of counters on a fixed interval.

No devices dedicated to the VSE guest were attached through a cache control unit and thus a sample of the VMPRF DASD extended function data is not included.

An SCHM is issued by EXPLORE/VSE at start of the monitor, and thus the assisted I/Os are not counted in CP's measurement-blocks.

The sample RTM data (Figure 21) shows both devices in I/O PASSTHRU.

For device 0647, the I/O count of 8227 shown in the sample EXPLORE/VSE XA DEVICE STATISTICS (Figure 20) represents all of the VSE-initiated I/Os. The count of 32 transitions shown in the sample VMPRF I/O assist data (Figure 23 on page 159) represents the number of times this device was temporarily taken out of I/O assist. The I/O count of 32 shown in the sample VMPRF DASD Activity data (Figure 22 on page 159) represents all of the non-assisted VSE-initiated I/Os and is equal to the number of transitions.

For device 0347, the I/O count of 3655 shown in the sample EXPLORE/VSE XA DEVICE STATISTICS (Figure 20) represents all of the VSE-initiated I/Os. The count of 0 transitions shown in the sample VMPRF I/O assist data (Figure 23 on page 159) indicates the device maintained the same I/O assist status throughout the measurement period. The count of 0 I/Os shown in the sample VMPRF DASD Activity data (Figure 22 on page 159) represents all of the non-assisted VSE-initiated I/Os and implies that this device remained in I/O assist for the duration of the measurement.

GOAL SYSTEMS INTERNATIONAL, INC.	E X P L O R E / R P T		COPYRIGHT 1991 RELEASE: 5.75 07/19/91 14.44	
FROM 10/06/92 13:45:00	VSE GUEST MEASUREMENT VR1170DA		REPORT 3 PAGE 1	
TO 10/06/92 14:00:00	XA DEVICE STATISTICS		REPORT DATE 10/06/92 14.14.58	
SHIFT NONE				
EACH DAY				
PERIOD NONE				

DATE	TIME	VOLSER ID	DISK CUU	SSCH COUNT	SSCH RATE	SAMPLE COUNT	CONN TIME	PEND TIME	DISC TIME	SERV TIME
10/06/92	00.00	CICS02	0647	8227	9.14	8227	0.00333	0.00026	0.01869	0.02227
		VSAM08	0347	3655	4.06	3655	0.00922	0.00051	0.00512	0.01485

Figure 20. EXPLORE/VSE XA DEVICE STATISTICS

```

<>VM/ESA CPU9221 SERIAL 056789 222M DATE 10/06/92 START 13:45:00 END 13:59:47<>
<--- DEVICE ---> <----- DEVICE RDEV DATA -----> <-- MEASUREMENT FACILITY -->
DEV TYPE VOLSER IOREQST SEC %Q %ER R %LK LNK PA %UT ACC FPT DCT CN %CN
0647 3380 CICS02 DEVICE IN I/O PASSTHRU SINCE DATA LAST RESET
0347 3380 VSAM08 DEVICE IN I/O PASSTHRU SINCE DATA LAST RESET
    
```

Figure 21. RTM VM/ESA Device Screen.

```

PRF012 Run 10/06/1992 14:11:15      DASD_BY_ACTIVITY                      Page 1
                                     DASD Activity Ordered by Activity

From 10/06/1992 13:45:13           Sample
To 10/06/1992 14:00:13             CPU 9221      SN 56789
For 900 Secs 00:15:00              VSE guest measurement VR1170DA      VM/ESA 20.01 SLU 000
    
```

<-----Device----->			<-----SSCH+RSCH----->					<-----Time----->					<-SSCHs in <--Queue-->				
Num-ber	Volume Serial Type	Control Unit Owner	Mini-disk Links	On-line Secs	Count	Rate	Plus Avoided	Plus Avoid Rate	Pct Busy	Pend	Disc	Conn	Serv	Resp	Mean	Max	Err
0647	CICS02 3380-D	3880-03 USERVR	0	900	32	0.0	32	0.0	0.0	0.2	0.3	0.3	0.8	0.8	0	0	0
0347	VSAM08 3380-A	3880-03 USERVR	0	900	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 22. VMPRF DASD_BY_ACTIVITY.

```

PRF098 Run 01/26/1993 11:20:51     DASD_IO_ASSIST                      Page 2
                                     SIE Assist Activity by Device

From 10/06/1992 13:45:13           Sample
To 10/06/1992 14:00:13             CPU 9221      SN 56789
For 900 Secs 00:15:00              VSE guest measurement VR1170DA      VM/ESA 20.01 SLU 0000
    
```

<-----Device----->			<Eligible>		<--Transition-->		<---Percentage of Time--->			
Num-ber	Volume Serial Type	Control Unit Owner	XA	370	On-line Secs	Count	Rate	Inside IO Assist	Leaving IO Assist	Outside IO Assist
0647	CICS02 3380-D	3880-03 USERVR	YES	YES	900	32	*****	*****	*****	*****
0347	VSAM08 3380-A	3880-03 USERVR	YES	YES	900	0	*****	*****	*****	*****

Figure 23. VMPRF DASD_IO_ASSIST.

Additional Evaluations

This chapter explores some additional performance-related topics.

- “Processor Capacity” on page 162 explains measurement results from running a CMS-intensive minidisk workload on VM/ESA Release 2 on a 9021-900 with different numbers of processors.
- “Minidisk to SFS” on page 168 compares the performance of SFS and the CMS minidisk file system.
- “Measurement Variability” on page 173 presents data showing how much variability in measurement results can be expected on a 9121-480 processor using a CMS-intensive minidisk workload.
- “RACF” on page 185 showed RACF/VM 1.9.2 with improved performance relative to RACF/VM 1.9.0. Measurements indicate that the multiple server support introduced by RACF/VM 1.9.2 can provide constraint relief and other benefits without increasing overall processing requirements.
- “LPAR Performance” on page 179 illustrates how running VM/ESA in an LPAR affects the CP monitor data.
- “VSE Guest Using Shared DASD” on page 199 demonstrates the performance characteristics of various VSE guest environments using dedicated and shared DASD.

Processor Capacity

The processor capacity measurements were made to determine the performance of VM/ESA Release 2 when running on different size processors within the 9021 520-based family. The performance of VM/ESA Release 2 scaled as expected on these processors, based on measurements made on the 9021-900 configured with 1, 2, 3, and 6 processors online, and based on similar measurements made on the 9021 340-based processors using VM/ESA Release 1.1. The following graph represents the internal throughput rate (ITR) as a function of the number of processors online for the 9021-900.

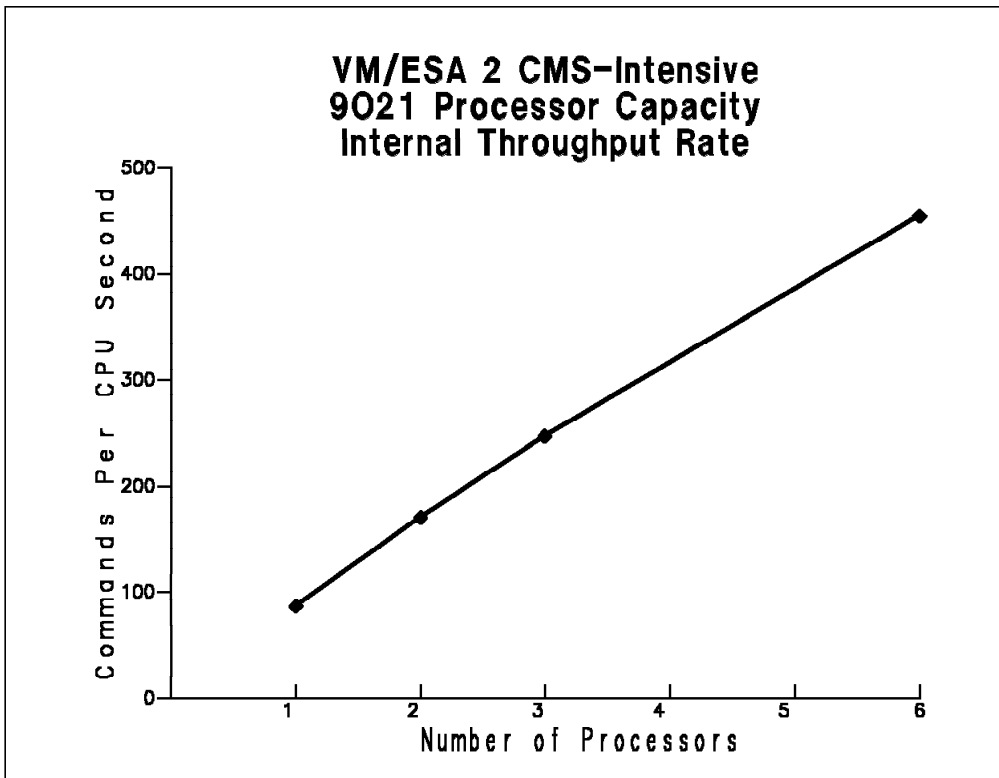


Figure 24. Internal Throughput Rate for the 9021-900 Processor.

Workload: FS7B0R

Hardware Configuration

Processor Model	Processors Used	Real Storage	Expanded Storage	Retained for MDC
9021-900	1	256MB	512MB	50MB
9021-900	2	384MB	1536MB	150MB
9021-900	3	512MB	2048MB	200MB
9021-900	6	1024MB	4096MB	400MB

Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3380-D	3880-03	4	18	30	30			
3380-E	3880-03	4	2		4			
3380-A	3880-03	4		2	6			
3380-D	3880-23	4				16 R		4 R
3380-K	3990-02	4	10	10		60		
3390-2	3990-02	4				40 R		

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Note: For each of the measurements, all of the DASD volumes were online. Because of the reduced number of users run when fewer processors were online, not all of the DASD volumes were required or used.

Communications:

Processor Model	Control Unit	Number	Lines per Control Unit	Speed
9021-900	3745-410	1	44	56Kb
9021-900	3745-410	2	44	56Kb
9021-900	3745-410	3	44	56Kb
9021-900	3745-410	4	44	56Kb
	3088-02	1	NA	4.5MB

Software Configuration

Driver: TPNS (1 Per processor online)
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

<i>Virtual Machine</i>	<i>Number</i>	<i>Type</i>	<i>Machine Size/Mode</i>	<i>RELSHARE</i>	<i>RESERVED</i>	<i>Other Options</i>
For the 9021-900 (1-Way):						
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	350	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
Users	2160	User	3MB/XC	100	OFF	
For the 9021-900 (2-Way):						
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	350	QUICKDSP ON
VSCS1	1	VSCS	64MB/XA	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
Users	4300	User	3MB/XC	100	OFF	
For the 9021-900 (3-Way):						
VTAMXA	1	VTAM	64MB/XA	10000	350	QUICKDSP ON
VSCSn	3	VSCS	64MB/XA	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
Users	6240	User	3MB/XC	100	OFF	
For the 9021-900 (6-Way):						
VTAMXA	1	VTAM	64MB/XA	10000	350	QUICKDSP ON
VSCSn	6	VSCS	64MB/XA	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
Users	11420	User	3MB/XC	100	OFF	

Measurement Discussion

The processor capacity measurements were made using a 9021-900 processor. The processor utilization of the existing 6-way measurement was chosen as the target for these measurements. For each measurement, the 9021-900 (6-way) was configured for the selected real and expanded storage sizes (or the next greatest size available on the 9021-900), and the SYSSTORE macro in HCPSYS was used to generate the exact real storage size. Any excess expanded storage was attached to an inactive user. The VARY OFFLINE PROCESSORS command was used to disable excess processors. The 6-way measurement reserved a fixed 400MB of expanded storage for minidisk caching (see "9021-900 / Minidisk" on page 33 for more detail). For the other measurements, a fixed amount of expanded storage was reserved for minidisk caching with the amount reserved having the same proportion to the total expanded storage as for the 6-way measurement.

While each measurement had a single VTAM, the number and type of VSCSs active as VTAM Service Machines (VSMs) varied, as did the number of TPNS machines. The 1-way measurement had an internal VSCS machine to handle the user traffic driven by one TPNS machine. The 2-way measurement used both an internal and an external VSCS machine, with user traffic driven by two TPNS machines. The 3-way and 6-way measurements did not use the internal VSCS. These measurements had external VSCS and TPNS machines equal to the number of processors.

Table 39 on page 166 summarizes the four processor capacity measurements. The results show that VM/ESA Release 2 scaled as expected across a selection of the 9021 520-based family. The ITRs ranged from 86.9 on the 1-way to 455.0 for the 6-way with ITRRs of 1, 1.96, 2.85, and 5.24. These results were similar to

studies on the 9021 340-based family using VM/ESA Release 1.1, which are published in the *VM/ESA Release 1.1 Performance Report*.

The ITRRs show a relative drop in ITR/processor caused by

- normal costs of inter-processor communications
- additional processing requirements generated by using external VSCS machines.
- additional expanded storage paging (XSTOR/CMD)

Paging to expanded storage increased because the amount of real storage did not increase in proportion to the number of processors (see RSTOR RATIO), resulting in less real storage per user for the measurements with more processors online. The real storage available was limited by the amount orderable for any given model of the 9021. The goal in selecting storage sizes for the processor capacity measurements was for the sum of real and expanded storage to scale with the number of processors.

The drop in ITR/processor can be seen by examining the CP/CMD (H) and EMUL/CMD (H) components of PBT/CMD (H). EMUL/CMD (H) varied by less than 4% across the measurements, while the CP processor time (CP/CMD (H)) increased by 4.7% to 9.4% to 20.9% for the 1-way to 2-way, 2-way to 3-way, and 3-way to 6-way measurements respectively. Approximately 29% of the 14.5% increase in CP/CMD (H) between the 1-way and the 3-way was attributed to the increased cost of the VTAM and VSCS virtual machines. Likewise, 16% of the small EMUL/CMD (H) increase (1.2%) was from additional resource consumption in VTAM and VSCS. The total processor and emulation processor usage in the VTAM machines (TOT CPU/CMD (V) and VIRT CPU/CMD (V)) increased as the number of external VSCS machines and processors increased.

Note the drop in DIAGNOSE X'98' instructions issued per command as the number of external VSCS machines increased. This reflected an increase in VTAM's ability to chain I/O buffers. The increases in PRIVOP/CMD are related to the increased IUCV usage associated with the external VSCS machines.

RELEASE RUN ID	VM/ESA 2 215R2160	VM/ESA 2 225R4301	VM/ESA 2 235R6243	VM/ESA 2 265RB426
Environment				
REAL STORAGE	256MB	384MB	512MB	1024MB
RSTOR RATIO	1.00	1.50	2.00	4.00
EXP. STORAGE	512MB	1536MB	2048MB	4096MB
XSTOR RATIO	1.00	3.00	4.00	8.00
USERS	2160	4300	6240	11420
VTAMs	1	1	1	1
VSCSS	0	1	3	6
PROCESSORS	1	2	3	6
Response Time				
TRIV INT	0.021	0.026	0.052	0.080
NONTRIV INT	0.337	0.244	0.211	0.266
TOT INT	0.167	0.144	0.156	0.198
TOT INT ADJ	0.226	0.169	0.154	0.192
AVG FIRST (T)	0.266	0.282	0.312	0.319
AVG LAST (T)	0.465	0.422	0.402	0.406
Throughput				
AVG THINK (T)	26.09	25.99	26.03	25.98
ETR	103.88	179.58	219.27	394.82
ETR (T)	76.64	153.20	222.34	406.79
ETR RATIO	1.356	1.172	0.986	0.971
ITR (H)	86.89	170.17	247.31	454.96
ITR	117.90	99.88	81.52	73.73
EMUL ITR	165.62	141.67	118.70	113.52
ITRR (H)	1.000	1.958	2.846	5.236
ITRR	1.000	0.847	0.691	0.625
Proc. Usage				
PBT/CMD (H)	11.509	11.753	12.130	13.188
PBT/CMD	11.483	11.749	12.099	13.176
CP/CMD (H)	3.630	3.800	4.157	5.028
CP/CMD	3.262	3.460	3.778	4.622
EMUL/CMD (H)	7.878	7.953	7.974	8.160
EMUL/CMD	8.221	8.290	8.321	8.555
Processor Util.				
TOTAL (H)	88.20	180.05	269.70	536.47
TOTAL	88.00	180.00	269.00	536.00
UTIL/PROC (H)	88.20	90.03	89.90	89.41
UTIL/PROC	88.00	90.00	89.67	89.33
TOTAL EMUL (H)	60.37	121.84	177.29	331.93
TOTAL EMUL	63.00	127.00	185.00	348.00
MASTER TOTAL (H)	88.20	90.26	90.96	94.21
MASTER TOTAL	88.00	90.00	91.00	94.00
MASTER EMUL (H)	60.37	55.68	48.35	32.02
MASTER EMUL	63.00	58.00	51.00	34.00
TVR(H)	1.46	1.48	1.52	1.62
TVR	1.40	1.42	1.45	1.54
Storage				
NUCLEUS SIZE (V)	2376KB	2376KB	2376KB	2376KB
TRACE TABLE (V)	200KB	400KB	600KB	1200KB
WKSET (V)	67	64	65	61
PGBLPGS	52659	78300	101 k	209 k
PGBLPGS/USER	24.4	18.2	16.2	18.3
FREEPGS	5733	11352	16404	30652
FREE UTIL	0.94	0.90	0.91	0.90
SHRPGS	1497	1818	1933	2609
Paging				
READS/SEC	287	323	543	666
WRITES/SEC	148	150	322	446
PAGE/CMD	5.676	3.087	3.891	2.734
PAGE IO RATE (V)	51.500	48.600	114.600	142.600
PAGE IO/CMD (V)	0.672	0.317	0.515	0.351
XSTOR IN/SEC	290	776	1267	3021
XSTOR OUT/SEC	454	993	1697	3670
XSTOR/CMD	9.708	11.547	13.331	16.448
FAST CLR/CMD	6.290	6.403	6.405	6.409

RELEASE RUN ID	VM/ESA 2 215R2160	VM/ESA 2 225R4301	VM/ESA 2 235R6243	VM/ESA 2 265RB426
Environment				
REAL STORAGE	256MB	384MB	512MB	1024MB
RSTOR RATIO	1.00	1.50	2.00	4.00
EXP. STORAGE	512MB	1536MB	2048MB	4096MB
XSTOR RATIO	1.00	3.00	4.00	8.00
USERS	2160	4300	6240	11420
VTAMs	1	1	1	1
VSCSs	0	1	3	6
PROCESSORS	1	2	3	6
Queues				
DISPATCH LIST	42.31	69.84	80.30	174.63
ELIGIBLE LIST	0.00	0.00	0.00	0.00
I/O				
VIO RATE	665	1314	1894	3506
VIO/CMD	8.677	8.577	8.519	8.619
RIO RATE (V)	238	394	593	1016
RIO/CMD (V)	3.106	2.572	2.667	2.498
MDC READS	437	866	1258	2313
MDC WRITES	207	420	597	1097
MDC MODS	171	347	495	905
MDC HIT RATIO	0.92	0.92	0.92	0.92
PRIVOPs				
PRIVOP/CMD	14.093	15.046	15.998	15.886
DIAG/CMD	25.607	24.614	24.175	24.092
DIAG 08/CMD	0.757	0.770	0.751	0.767
DIAG 14/CMD	0.013	0.020	0.022	0.022
DIAG 58/CMD	1.200	1.208	1.214	1.219
DIAG 98/CMD	0.339	0.248	0.202	0.258
DIAG A4/CMD	3.928	3.949	3.904	3.938
DIAG A8/CMD	1.879	1.841	1.898	1.900
DIAG 214/CMD	13.218	13.257	13.147	13.260
SIE/CMD	45.188	45.594	44.977	46.707
SIE INTCPT/CMD	30.276	29.636	29.235	29.425
FREE TOTL/CMD	83.682	84.857	85.456	86.039
VTAM Machines				
WKSET (V)	458	2880	7184	12025
TOT CPU/CMD (V)	1.4789	1.6319	1.8016	2.2249
CP CPU/CMD (V)	0.6524	0.7434	0.8446	1.1384
VIRT CPU/CMD (V)	0.8264	0.8885	0.9570	1.0865
DIAG 98/CMD (V)	0.347	0.251	0.206	0.260
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

Table 39. VM/ESA Release 2 on the 9021-900 Using Different Numbers of Processors

Minidisk to SFS

Workload: FS7B0R / FS7BMAXR

Hardware Configuration

Processor model: 9121-480
 Processors used: 2
 Storage
 Real: 192MB
 Expanded: 64MB (all retained for MDC)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3390-2	3990-2	4	8	4	4			
3390-2	3990-3	4					10 R	
3390-2	3990-3	2	4 R	2 R	2 R			2 R

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3088-02	1	NA	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVE	Other Options
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	400	QUICKDSP ON
ROSERV1	1	SFS	32MB/XC	1500	OFF	QUICKDSP ON
RWSERVn	2	SFS	32MB/XC	1500	1000	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	100	OFF	
Users	1870/1560	User	3MB/XC	100	OFF	

Measurement Discussion

The measurements in this section compare the performance of the CMS minidisk file system to the Shared File System (SFS) to demonstrate the effects of migrating files from minidisks to SFS. For these measurements, all files were moved from minidisks to SFS, except those on the S and Y minidisks (not supported by SFS). The SFS environment consists of all read/write data in FILECONTROL directories and all read-only data in DIRCONTROL directories in

VM Data Spaces. The read-only DIRCONTROL directories are in a separate read-only file pool (ROSERV1).

Measurements were obtained on VM/ESA Release 1.1 and VM/ESA Release 2. Below are the major run characteristics for this comparison:

- L24R1876** VM/ESA Release 1.1 minidisk with the number of users selected to obtain an approximate processor utilization of 90%.
- L24M1566** VM/ESA Release 1.1 maximum SFS, with the number of users selected to obtain an approximate processor utilization of 90%.
- L25R187F** VM/ESA Release 2 minidisk, with the number of users set to the same number used for the L24R1876 measurement.
- L25M1564** VM/ESA Release 2 maximum SFS, with the number of users set to the same number used for the L24M1566 measurement.

When migrating all supported data from minidisks to SFS while keeping the processor utilization at approximately 90%, VM/ESA Release 2 measurements showed a decrease in internal throughput rate (ITR (H)) similar to the decrease shown in the VM/ESA Release 1.1 measurements.

The difference in response time between minidisks and SFS, both internal (TOT INT ADJ) and external (AVG LAST (T)), decreased significantly in VM/ESA Release 2. For VM/ESA Release 1.1, the internal response time for SFS was 0.091 seconds greater than minidisks while for VM/ESA Release 2, this difference was only 0.037 seconds. The external response time for VM/ESA Release 1.1 showed SFS was 0.050 seconds greater than minidisks and for VM/ESA Release 2, SFS was 0.012 seconds less than minidisks. The reason for the response time improvement for SFS in VM/ESA Release 2 was due to the performance enhancements for SFS, primarily the SFS checkpoint enhancements and increasing the file cache size from 12KB to 20KB. See "Performance Improvements" on page 6 for a description of these enhancements.

Paging (PAGE/CMD) has shown a significant increase for SFS in VM/ESA Release 2. In VM/ESA Release 1.1, there was an 11% reduction in paging for SFS compared to minidisks. In VM/ESA Release 2, SFS has a 1% increase in paging over minidisks. This is due to the increased real storage requirements for VM/ESA Release 2. With minidisks, this was offset by the improvements in virtual storage requirements and page references for the read-only minidisk directories. SFS does not benefit from this improvement, and thus experiences higher paging.

Virtual file I/Os per command can be approximated by adding:

- DIAG A4/CMD
- DIAG A8/CMD (this includes some additional I/O that is not file system related, but it should be about the same in both cases)
- IO/CMD (Q) (for SFS run)

There was a 1.5% reduction in virtual file I/Os for SFS compared to minidisks in VM/ESA Release 1.1. For VM/ESA Release 2, the virtual file I/O reduction for SFS was 3.5%. This further improvement was due to a 10% reduction in SFS IO/CMD as result of increasing the default file cache size from 12KB to 20KB.

Note that these observations are for an equal processor utilization environment and that the results would be different for an equal user environment. It is

expected that the results for an equal user environment would be similar to those documented in the *VM/ESA Release 1.1 Performance Report*.

FILE SYSTEM RELEASE RUN ID	MINIDISK VM/ESA 1.1 L24R1876	SFS VM/ESA 1.1 L24M1566	MINIDISK VM/ESA 2 L25R187F	SFS VM/ESA 2 L25M1564
Environment				
REAL STORAGE	192MB	192MB	192MB	192MB
EXP. STORAGE	64MB	64MB	64MB	64MB
USERS	1870	1560	1870	1560
VTAMs	1	1	1	1
VSCSs	0	0	0	0
Response Time				
TRIV INT	0.123	0.120	0.127	0.128
NONTRIV INT	0.523	0.677	0.539	0.608
TOT INT	0.384	0.498	0.395	0.454
TOT INT ADJ	0.351	0.442	0.357	0.394
AVG FIRST (T)	0.334	0.305	0.330	0.286
AVG LAST (T)	0.492	0.542	0.498	0.486
Throughput				
AVG THINK (T)	25.54	25.62	25.43	25.58
ETR	61.26	49.34	60.60	49.08
ETR (T)	66.98	55.60	66.99	56.49
ETR RATIO	0.915	0.887	0.905	0.869
ITR (H)	73.51	61.61	73.86	61.92
ITR	33.65	27.35	33.44	26.94
EMUL ITR	50.63	41.87	49.18	40.20
ITRR (H)	1.000	0.838	1.005	0.842
ITRR	1.000	0.813	0.994	0.801
Proc. Usage				
PBT/CMD (H)	27.209	32.464	27.078	32.301
PBT/CMD	27.172	32.375	27.021	32.217
CP/CMD (H)	9.794	11.963	9.364	11.397
CP/CMD	9.107	11.151	8.659	10.621
EMUL/CMD (H)	17.415	20.501	17.714	20.903
EMUL/CMD	18.065	21.224	18.362	21.596
Processor Util.				
TOTAL (H)	182.24	180.49	181.38	182.47
TOTAL	182.00	180.00	181.00	182.00
UTIL/PROC (H)	91.12	90.25	90.69	91.24
UTIL/PROC	91.00	90.00	90.50	91.00
TOTAL EMUL (H)	116.64	113.98	118.66	118.08
TOTAL EMUL	121.00	118.00	123.00	122.00
MASTER TOTAL (H)	90.97	90.19	90.48	91.04
MASTER TOTAL	91.00	90.00	90.00	91.00
MASTER EMUL (H)	52.74	52.73	53.86	54.50
MASTER EMUL	55.00	55.00	56.00	56.00
TVR(H)	1.56	1.58	1.53	1.55
TVR	1.50	1.53	1.47	1.49
Storage				
NUCLEUS SIZE (V)	2792KB	2792KB	2304KB	2304KB
TRACE TABLE (V)	400KB	400KB	400KB	400KB
WKSET (V)	73	64	77	76
PGBLPGS	39883	40824	39665	40629
PGBLPGS/USER	21.3	26.2	21.2	26.0
FREEPGS	4741	4171	4876	4276
FREE UTIL	0.97	0.91	0.95	0.90
SHRPGS	1109	1311	1177	1377
Paging				
READS/SEC	562	412	586	492
WRITES/SEC	370	278	389	340
PAGE/CMD	13.915	12.410	14.555	14.728
PAGE IO RATE (V)	184.200	132.700	177.300	140.700
PAGE IO/CMD (V)	2.750	2.387	2.647	2.491
XSTOR IN/SEC	0	0	0	0
XSTOR OUT/SEC	0	0	0	0
XSTOR/CMD	0.000	0.000	0.000	0.000
FAST CLR/CMD	5.748	5.863	6.404	6.709

FILE SYSTEM RELEASE RUN ID	MINIDISK VM/ESA 1.1 L24R1876	SFS VM/ESA 1.1 L24M1566	MINIDISK VM/ESA 2 L25R187F	SFS VM/ESA 2 L25M1564
Environment				
REAL STORAGE	192MB	192MB	192MB	192MB
EXP. STORAGE	64MB	64MB	64MB	64MB
USERS	1870	1560	1870	1560
VTAMs	1	1	1	1
VSCSs	0	0	0	0
Queues				
DISPATCH LIST	39.12	42.26	45.11	41.09
ELIGIBLE LIST	0.00	0.00	0.00	0.00
I/O				
VIO RATE	624	534	616	530
VIO/CMD	9.316	9.605	9.196	9.382
RIO RATE (V)	391	345	376	344
RIO/CMD (V)	5.838	6.205	5.613	6.089
MDC READS	363	259	367	266
MDC WRITES	182	114	181	118
MDC MODS	150	75	151	78
MDC HIT RATIO	0.91	0.86	0.92	0.86
PRIVOPs				
PRIVOP/CMD	17.610	29.533	17.507	24.666
DIAG/CMD	24.604	21.925	26.135	24.027
DIAG 08/CMD	0.746	0.791	0.776	0.814
DIAG 14/CMD	0.015	0.018	0.015	0.018
DIAG 58/CMD	1.254	1.259	1.224	1.239
DIAG 98/CMD	0.851	1.223	0.866	1.257
DIAG A4/CMD	3.912	2.068	3.852	2.053
DIAG A8/CMD	1.971	1.727	1.926	1.699
DIAG 214/CMD	11.601	10.828	13.197	12.763
SIE/CMD	52.374	64.912	51.877	60.806
SIE INTCPT/CMD	36.662	48.035	35.276	43.172
FREE TOTL/CMD	83.115	95.632	64.626	73.286
VTAM Machines				
WKSET (V)	421	421	415	407
TOT CPU/CMD (V)	3.7739	4.1568	3.7736	4.0714
CP CPU/CMD (V)	1.7086	1.8586	1.6836	1.7899
VIRT CPU/CMD (V)	2.0653	2.2982	2.0900	2.2816
DIAG 98/CMD (V)	0.852	1.223	0.872	1.259
SFS Servers				
WKSET (V)	na	1573	na	1611
TOT CPU/CMD (V)	na	4.1568	na	3.6977
CP CPU/CMD (V)	na	2.0384	na	1.7899
VIRT CPU/CMD (V)	na	2.1184	na	1.9079
FP REQ/CMD(Q)	na	1.392	na	1.270
IO/CMD (Q)	na	2.000	na	1.804
IO TIME/CMD (Q)	na	0.033	na	0.032
SFS TIME/CMD (Q)	na	0.083	na	0.054
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Query Filepool Counters, Unmarked=RTM				

Table 40. Minidisk to SFS Comparison / Equal Processor Utilization on 9121-480

Measurement Variability

Workload: FS7B0R

Hardware Configuration

Processor model: 9121-480
 Processors used: 2
 Storage
 Real: 192MB
 Expanded: 64MB (All reserved for MDC)
 Tape: 3480 (MONITOR)

DASD:

<i>Type of DASD</i>	<i>Control Unit</i>	<i>Number of Paths</i>	<i>PAGE</i>	<i>SPOOL</i>	<i>- Number of Volumes -</i>			<i>System</i>
					<i>TDSK</i>	<i>User</i>	<i>Server</i>	
3390-2	3990-2	4	8	4	4			
3390-2	3990-3	4				10 R		
3390-2	3990-3	2	4	2	2			2 R

Note: *R* or *W* next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

<i>Control Unit</i>	<i>Number</i>	<i>Lines per Control Unit</i>	<i>Speed</i>
3088-02	1	NA	3.0MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

<i>Virtual Machine</i>	<i>Number</i>	<i>Type</i>	<i>Machine Size/Mode</i>	<i>RELSHARE</i>	<i>RESERVED</i>	<i>Other Options</i>
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	400	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	100	OFF	
Users	1870	User	2MB/XC	100	OFF	

Measurement Discussion

The same measurement was repeated four times to get a better understanding of FS7B0R workload run variability. Table 41 on page 175 shows the results for each individual measurement. Table 42 on page 177 shows the average, range, standard deviation, and percent standard deviation for these four runs. Range is the absolute value of the difference between the lowest and highest value observed. Standard deviation is calculated using the standard formula that assumes independent, normally distributed variables. Percent standard deviation is 100 times the standard deviation divided by the average.

If, for repeated measurements, one assumes that the values observed for any given metric are normally distributed, then that value should fall within one standard deviation of the average 68% of the time and within two standard deviations of the average 95% of the time.

Observe that some metrics show far more variability than others. For example, the percent standard deviation for AVG LAST (T) (external response time) is 5.0% while for ITR (H) (internal throughput rate) it is only 0.2%.

These results were obtained with VM/ESA Release 1.1. From informal observations, it is believed that VM/ESA Release 2 measurements show very similar run variabilities.

These results, of course, only apply to the measured environment: 9121-480, 192MB/64MB. However, the relative variabilities of the various metrics are probably similar to these results in other environments as well. For example, AVG LAST (T) is almost certainly much more variable than ITR (H) regardless of what processor is used.

RELEASE RUN ID	VM/ESA 1.1 L24R1871	VM/ESA 1.1 L24R1873	VM/ESA 1.1 L24R1874	VM/ESA 1.1 L24R1875
Response Time				
TRIV INT	0.122	0.120	0.125	0.126
NONTRIV INT	0.556	0.554	0.533	0.548
TOT INT	0.402	0.399	0.392	0.402
TOT INT ADJ	0.375	0.376	0.359	0.368
AVG FIRST (T)	0.363	0.370	0.332	0.332
AVG LAST (T)	0.543	0.555	0.502	0.507
Throughput				
AVG THINK (T)	25.40	25.43	25.53	25.51
ETR	62.51	63.32	61.35	61.37
ETR (T)	67.07	67.11	66.95	67.05
ETR RATIO	0.932	0.944	0.916	0.915
ITR (H)	73.61	73.81	73.71	73.44
ITR	34.33	34.85	33.79	33.63
EMUL ITR	51.62	52.57	51.08	50.77
ITRR (H)	1.000	1.003	1.001	0.998
ITRR	1.000	1.015	0.984	0.979
Proc. Usage				
PBT/CMD (H)	27.171	27.097	27.132	27.234
PBT/CMD	27.134	27.119	27.185	27.295
CP/CMD (H)	9.766	9.791	9.835	9.846
CP/CMD	9.094	9.238	9.261	9.247
EMUL/CMD (H)	17.406	17.306	17.297	17.388
EMUL/CMD	18.040	17.881	17.924	18.047
Processor Util.				
TOTAL (H)	182.25	181.85	181.65	182.59
TOTAL	182.00	182.00	182.00	183.00
UTIL/PROC (H)	91.13	90.93	90.82	91.30
UTIL/PROC	91.00	91.00	91.00	91.50
TOTAL EMUL (H)	116.75	116.14	115.80	116.58
TOTAL EMUL	121.00	120.00	120.00	121.00
MASTER TOTAL (H)	90.95	90.79	90.62	91.17
MASTER TOTAL	91.00	91.00	91.00	91.00
MASTER EMUL (H)	52.81	52.57	52.35	52.77
MASTER EMUL	55.00	55.00	54.00	55.00
TVR(H)	1.56	1.57	1.57	1.57
TVR	1.50	1.52	1.52	1.51
Storage				
NUCLEUS SIZE (V)	2792KB	2792KB	2792KB	2792KB
TRACE TABLE (V)	400KB	400KB	400KB	400KB
WKSET (V)	72	73	72	73
PGBLPGS	39842	39841	39842	39848
PGBLPGS/USER	21.3	21.3	21.3	21.3
FREEPGS	4789	4780	4782	4781
FREE UTIL	0.96	0.96	0.96	0.96
SHRPGS	1148	1096	1105	1061
Paging				
READS/SEC	565	534	539	542
WRITES/SEC	370	370	370	370
PAGE/CMD	13.940	13.470	13.577	13.603
PAGE IO RATE (V)	183.600	181.100	184.300	184.100
PAGE IO/CMD (V)	2.737	2.698	2.753	2.746
XSTOR IN/SEC	0	0	0	0
XSTOR OUT/SEC	0	0	0	0
XSTOR/CMD	0.000	0.000	0.000	0.000
FAST CLR/CMD	5.740	5.752	5.691	5.742
Queues				
DISPATCH LIST	44.81	45.39	44.67	40.54
ELIGIBLE LIST	0.02	0.00	0.00	0.00

RELEASE RUN ID	VM/ESA 1.1 L24R1871	VM/ESA 1.1 L24R1873	VM/ESA 1.1 L24R1874	VM/ESA 1.1 L24R1875
I/O				
VIO RATE	617	618	621	618
VIO/CMD	9.199	9.209	9.276	9.218
RIO RATE (V)	389	393	400	402
RIO/CMD (V)	5.800	5.856	5.975	5.996
MDC READS	367	359	363	361
MDC WRITES	181	177	182	180
MDC MODS	149	147	151	149
MDC HIT RATIO	0.92	0.92	0.92	0.92
PRIVOPs				
PRIVOP/CMD	17.711	17.619	17.612	17.642
DIAG/CMD	24.501	24.352	24.501	24.476
DIAG 08/CMD	0.745	0.745	0.747	0.746
DIAG 14/CMD	0.015	0.015	0.015	0.015
DIAG 58/CMD	1.252	1.326	1.329	1.253
DIAG 98/CMD	0.820	0.805	0.822	0.850
DIAG A4/CMD	3.921	3.844	3.898	3.878
DIAG A8/CMD	1.879	1.922	1.897	1.909
DIAG 214/CMD	11.614	11.503	11.501	11.589
SIE/CMD	52.106	52.063	52.263	52.457
SIE INTCPT/CMD	36.474	36.444	36.584	36.720
FREE TOTL/CMD	82.923	82.504	82.808	82.944
VTAM Machines				
WKSET (V)	426	419	404	424
TOT CPU/CMD (V)	3.8100	3.7748	3.7756	3.7785
CP CPU/CMD (V)	1.7145	1.7053	1.7094	1.7070
VIRT CPU/CMD (V)	2.0955	2.0695	2.0662	2.0716
DIAG 98/CMD (V)	0.821	0.813	0.831	0.858
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

Table 41. Measurement Variability of FS7B0R on the 9121-480 - Results

STATISTIC	AVERAGE	RANGE	STD DEV.	% STD DEV.
Response Time				
TRIV INT	0.123	0.006	0.003	2.2
NONTRIV INT	0.548	0.023	0.010	1.9
TOT INT	0.399	0.010	0.005	1.2
TOT INT ADJ	0.370	0.017	0.008	2.1
AVG FIRST (T)	0.349	0.038	0.020	5.8
AVG LAST (T)	0.527	0.053	0.026	5.0
Throughput				
AVG THINK (T)	25.47	0.13	0.06	0.2
ETR	62.14	1.97	0.96	1.5
ETR (T)	67.05	0.16	0.07	0.1
ETR RATIO	0.927	0.029	0.014	1.5
ITR (H)	73.64	0.37	0.16	0.2
ITR	34.15	1.22	0.55	1.6
EMUL ITR	51.51	1.80	0.79	1.5
ITRR (H)	1.001	0.005	0.002	0.2
ITRR	0.995	0.036	0.016	1.6
Proc. Usage				
PBT/CMD (H)	27.159	0.137	0.059	0.2
PBT/CMD	27.183	0.176	0.080	0.3
CP/CMD (H)	9.810	0.080	0.037	0.4
CP/CMD	9.210	0.167	0.078	0.8
EMUL/CMD (H)	17.349	0.109	0.056	0.3
EMUL/CMD	17.973	0.166	0.083	0.5
Processor Util.				
TOTAL (H)	182.09	0.94	0.42	0.2
TOTAL	182.25	1.00	0.50	0.3
UTIL/PROC (H)	91.05	0.48	0.21	0.2
UTIL/PROC	91.13	0.50	0.25	0.3
TOTAL EMUL (H)	116.32	0.95	0.43	0.4
TOTAL EMUL	120.50	1.00	0.58	0.5
MASTER TOTAL (H)	90.88	0.55	0.23	0.3
MASTER TOTAL	91.00	0.00	0.00	0.0
MASTER EMUL (H)	52.63	0.46	0.21	0.4
MASTER EMUL	54.75	1.00	0.50	0.9
TVR(H)	1.57	0.01	0.01	0.3
TVR	1.51	0.02	0.01	0.6
Storage				
NUCLEUS SIZE (V)	2792KB	0KB	0KB	0.0
TRACE TABLE (V)	400KB	0KB	0KB	0.0
WKSET (V)	73	1	1	0.8
PGBLPGS	39843	7	3	0.0
PGBLPGS/USER	21.3	0.0	0.0	0.0
FREEPGS	4783	9	4	0.1
FREE UTIL	0.96	0.00	0.00	0.0
SHRPGS	1103	87	36	3.2
Paging				
READS/SEC	545	31	14	2.5
WRITES/SEC	370	0	0	0.0
PAGE/CMD	13.648	0.470	0.203	1.5
PAGE IO RATE (V)	183.275	3.200	1.480	0.8
PAGE IO/CMD (V)	2.734	0.055	0.025	0.9
XSTOR IN/SEC	0	0	0	0.0
XSTOR OUT/SEC	0	0	0	0.0
XSTOR/CMD	0.000	0.000	0.000	0.0
FAST CLR/CMD	5.731	0.061	0.027	0.5

STATISTIC	AVERAGE	RANGE	STD DEV.	% STD DEV.
Queues				
DISPATCH LIST	43.85	4.85	2.23	5.1
ELIGIBLE LIST	0.01	0.02	0.01	200.0
I/O				
VIO RATE	619	4	2	0.3
VIO/CMD	9.226	0.077	0.035	0.4
RIO RATE (V)	396	13	6	1.5
RIO/CMD (V)	5.907	0.196	0.094	1.6
MDC READS	363	8	3	0.9
MDC WRITES	180	5	2	1.2
MDC MODS	149	4	2	1.1
MDC HIT RATIO	0.92	0.00	0.00	0.0
PRIVOPs				
PRIVOP/CMD	17.646	0.099	0.045	0.3
DIAG/CMD	24.458	0.149	0.071	0.3
DIAG 08/CMD	0.746	0.002	0.001	0.1
DIAG 14/CMD	0.015	0.000	0.000	0.0
DIAG 58/CMD	1.290	0.077	0.043	3.4
DIAG 98/CMD	0.824	0.045	0.019	2.3
DIAG A4/CMD	3.885	0.077	0.033	0.8
DIAG A8/CMD	1.902	0.043	0.018	1.0
DIAG 214/CMD	11.552	0.113	0.058	0.5
SIE/CMD	52.222	0.394	0.179	0.3
SIE INTCPT/CMD	36.556	0.276	0.125	0.3
FREE TOTL/CMD	82.795	0.440	0.203	0.2
VTAM Machines				
WKSET (V)	418	22	10	2.4
TOT CPU/CMD (V)	3.7847	0.0352	0.0169	0.4
CP CPU/CMD (V)	1.7091	0.0092	0.0040	0.2
VIRT CPU/CMD (V)	2.0757	0.0293	0.0134	0.6
DIAG 98/CMD (V)	0.831	0.045	0.020	2.4
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

Table 42. Measurement Variability of FS7B0R on the 9121-480 - Statistics

LPAR Performance

This section explores the performance characteristics of running VM/ESA Release 2 in logical partitions (LPARs) on a 9121-480 processor. This is not to imply that there is something unique about VM/ESA Release 2 running in an LPAR (there is not). Rather, the measurement data are presented to help understand how to interpret key metrics in an LPAR environment running VM/ESA guests.

The 9121-480 processor was configured with two LPARs (both shared for one measurement, both dedicated for another) with each running identical VM/ESA systems using the same workload with the same number of users. This simplified the examination of the data to understand how system resources were used. Although the data cannot be used as a basis for capacity planning in other situations, the discussion shows how the data can be interpreted.

Workload: FS7B0R

Hardware Configuration

Processor model: 9121-480
 Processors used: 2
 Storage
 Real: 192MB
 Expanded: 64MB (all used for MDC)

Storage per LPAR
 Central Initial Storage: 90MB
 Expanded Initial Storage: 32MB (all used for MDC)

LPAR Settings
 Weight: 500 for shared, DED for dedicated
 Processor Running Time: Dynamically Determined
 Wait Completion: No for shared, Yes for dedicated

Tape: 3480 (MONITOR)

DASD:

<i>Type of DASD</i>	<i>Control Unit</i>	<i>Number of Paths</i>	<i>PAGE</i>	<i>SPOOL</i>	<i>- Number of Volumes -</i>			<i>System</i>
					<i>TDSK</i>	<i>User</i>	<i>Server</i>	
LPAR 1 DASD:								
3390-2	3990-2	4	8	2	4			
3390-2	3990-3	2	4	2	2			2 R
3390-2	3990-3	4				5 R		
LPAR 2 DASD:								
3390-2	3990-2	2		1				2 R
3390-2	3990-2	4	5	2	5			
3390-2	3990-3	2	3		2			
3390-2	3990-3	4				5 R		

Note: *R* or *W* next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

<i>Control Unit</i>	<i>Number</i>	<i>Lines per Control Unit</i>	<i>Speed</i>
3088-08	2	NA	4.5MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines (per LPAR):

<i>Virtual Machine</i>	<i>Number</i>	<i>Type</i>	<i>Machine Size/Mode</i>	<i>RELSHARE</i>	<i>RESERVED</i>	<i>Other Options</i>
VTAMXA	1	VTAM/VSCS	64MB/XC	10000	400	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
Users	890/910	User	3MB/XC	100	OFF	

Measurement Discussion

Table 45 on page 183 presents data for one of the two LPAR VM/ESA systems in each of the dedicated LPAR and shared LPAR measurements. Table data for the other LPAR system in each measurement is not provided here because of its similarity to the first LPAR in each case: the systems and workload in each LPAR were the same although the I/O configuration differed slightly.

The data do not reveal significant insights about VM/ESA Release 2 performance that have not been explored in other parts of this report. In order to understand the broader context (VM/ESA Release 2 and LPAR working in concert to complete work), metrics that give insight to system resource use are examined more closely rather than confining observations to what the VM systems were doing inside the LPARs.

Shared LPAR: The following table summarizes certain CP monitor data gathered for the shared LPAR measurement.

Data Source	LPAR 1	LPAR 2
VMPRF PROCESSORS_BY_TIME report		
Total	88.8	90.3
User	82.2	83.5
Syst	6.6	6.7
Emul	59.4	60.5
VMPRF LPAR report		
Pct Processor Util		
Logical	89.2	90.6
Physical	44.6	45.3
CP monitor domain 0, record 17		
LPAR management time	0.3	

Table 43. Shared LPAR CP monitor data summary (percent, rounded to tenths)

The table shows VM in LPAR 1 reported 88.8% processor utilization and VM in LPAR 2 reported 90.3% processor utilization, for an average of 89.6%. A different view is taken in VMPRF's LPAR report: 89.2% utilization was reported for the *logical* processor associated with LPAR 1 and 90.6% for LPAR 2 for an average of 89.9%. The 0.3% difference between this and the VM-reported average can be attributed to the amount of LPAR processing time that can be directly charged to support of the work being done in the LPARs.

Alternatively, a total of 44.6% of all available physical processor capacity (2 processors) was attributed to support of LPAR 1 and 45.3% to LPAR 2. Thus a total of 89.9% of available physical processor capacity was reported by VMPRF in direct support of LPAR work. Note that this is in agreement with the view taken of the logical processors.

Some work (such as scheduling) being done by LPAR is not directly charged to support of work being done in the LPARs. This LPAR management time can be found in CP monitor domain 0, record 17. In this case there was an average of about 0.3% processor utilization attributed to LPAR management.

Adding directly-attributable LPAR processor utilization (89.9%) to LPAR management time (0.3%) gives overall physical processor utilization. The total, 90.2%, is also what was reported by hardware instrumentation data for the measurement.

Dedicated LPAR: The same set of metrics are provided in the following table for the dedicated LPAR measurement.

Data Source	LPAR 1	LPAR 2
VMPRF PROCESSORS_BY_TIME report		
Total	91.5	91.0
User	84.7	84.1
Syst	6.8	6.8
Emul	61.3	61.0
VMPRF LPAR report		
Pct Processor Util		
Logical	99.9	99.7
Physical	49.9	49.8
CP monitor domain 0, record 17		
LPAR management time	0.2	

Table 44. Dedicated LPAR CP monitor data summary (percent, rounded to tenths)

As in the shared LPAR case, VMPRF reports "Total" processor utilization from VM's viewpoint in the dedicated LPARs: 91.5% and 91.0% for LPAR 1 and LPAR 2 respectively, for an average of 91.2%.

Unlike the shared LPAR case, however, the VMPRF LPAR report data shows nearly 100% utilization of both logical and physical processor resource. The dedicated LPAR configuration for the system used the default Wait Completion = Yes. From LPAR's viewpoint, the processor is 100% committed to the guest running in it (less any time LPAR needs it). This is regardless of what the guest may be doing with the processor—some of the time was spent in a wait state.

The LPAR management time in this case was reported as about 0.2%. Working backward, it is known that hardware instrumentation reported 91.4% processor utilization, so subtracting LPAR management time yields 91.2% physical processor utilization directly in support of the LPARs' work.

The processor utilization reported from VM's view and the utilization calculated by subtracting LPAR management time from overall utilization were the same in this case: 91.2%. This indicates that, as expected, there was negligible processor resource used by LPAR in direct support of either partition. This was because of the dedication of physical processors to the logical partitions, in contrast to the shared LPAR measurement discussed previously.

Low Utilization: Some measurements were made (data not shown) to assess the behavior of the system in a low-processor-utilization environment. The configurations were the same as those for the dedicated- and shared-LPAR measurements discussed previously except only 180 users were active, giving a processor utilization in the range of 9 to 10 percent. In each case, the LPAR management utilization was found to be nearly identical to those found in the measurement data shown previously. The LPAR time charged to support of specific LPARs was roughly 25% of that in the previous measurements. (The numbers were sufficiently small to allow for significant margin of error due to roundoff in the calculations.) Therefore, in this configuration, little if any "low utilization effect" was observed for the LPAR processor requirements.

Comparison to Basic Mode Measurement: Simple calculations show that, even allowing for LPAR overhead, there is a difference in the number of users that can be supported per unit of processor utilization compared to the "native" measurement of a single VM/ESA Release 2 system running on the 9121-480 in basic mode (as shown in "9121-480 / Minidisk" on page 42). Although not examined in detail, because the difference is small, there are 2 factors that could have an effect:

1. Efficiencies exist when system management is consolidated in one rather than two control program images. The cost of only one CP monitor is incurred, one scheduler is responsible for the whole system, and so on.
2. There was less storage per user available for the LPAR measurements than for the native measurement. In this case 6MB of storage was set aside for use by LPAR (though not all was required) and each of the VM/ESA control programs occupied 2472KB for the nucleus and 200KB for the trace table. In addition, the LPAR configuration required 2 sets of shared pageable storage, 2 sets of minidisk caches, 2 VTAM virtual machines, and so on. The single-image non-LPAR system would only have one set of storage requirements for these items and none set aside for LPAR use.

Conclusions: The measurement data show that the processor resource consumed by LPAR for these environments was very small—significantly less than 1% in each case. The LPAR management time in all cases was 0.2 to 0.3 percent, and the amount of LPAR utilization charged directly to support of the VM/ESA guests was in the range of 0 to 0.3 percent. Thus, for this configuration, the total cost of running VM/ESA Release 2 in LPARs was 0.6% or less of available processor resource. Other costs were incurred by splitting a single control program image (in the basic mode case) into 2 images (in the LPAR cases) on the 9121-480.

RELEASE RUN ID LPAR TYPE	VM/ESA 2 LPARSHR1 SHARED	VM/ESA 2 LPARDED1 DEDICATED
Environment		
REAL STORAGE	90MB	90MB
EXP. STORAGE	32MB	32MB
USERS	890	910
VTAMs	1	1
VSCSs	0	0
PROCESSORS	1 (logical)	1 (dedicated)
Response Time		
TRIV INT	0.098	0.102
NONTRIV INT	0.479	0.648
TOT INT	0.372	0.485
TOT INT ADJ	0.319	0.425
AVG FIRST (T)	0.224	0.266
AVG LAST (T)	0.393	0.526
Throughput		
AVG THINK (T)	25.51	25.34
ETR	27.31	28.55
ETR (T)	31.89	32.61
ETR RATIO	0.856	0.876
ITR (H)	NA	NA
ITR	30.81	31.24
EMUL ITR	46.04	46.64
ITRR (H)	NA	NA
ITRR	1.000	1.014
Proc. Usage		
PBT/CMD (H)	NA	NA
PBT/CMD	27.906	27.906
CP/CMD (H)	NA	NA
CP/CMD	9.093	9.200
EMUL/CMD (H)	NA	NA
EMUL/CMD	18.813	18.706
Processor Util.		
TOTAL (H)	90.15	91.39
TOTAL	89.00	91.00
TOTAL EMUL (H)	NA	NA
TOTAL EMUL	60.00	61.00
TVR(H)	NA	NA
TVR	1.48	1.49
Storage		
NUCLEUS SIZE (V)	2472KB	2472KB
TRACE TABLE (V)	200KB	200KB
WKSET (V)	80	80
PGBLPGS	18172	18069
PGBLPGS/USER	20.4	19.8
FREEPGS	2465	2534
FREE UTIL	0.93	0.93
SHRPGS	963	1051
Paging		
READS/SEC	330	329
WRITES/SEC	206	215
PAGE/CMD	16.806	16.682
PAGE IO RATE (V)	90.100	95.200
PAGE IO/CMD (V)	2.825	2.919
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	6.114	6.317
Queues		
DISPATCH LIST	18.83	24.17
ELIGIBLE LIST	0.00	0.00

RELEASE RUN ID LPAR TYPE	VM/ESA 2 LPARSHR1 SHARED	VM/ESA 2 LPARDED1 DEDICATED
Environment		
REAL STORAGE	90MB	90MB
EXP. STORAGE	32MB	32MB
USERS	890	910
VTAMs	1	1
VSCSs	0	0
PROCESSORS	1 (logical)	1 (dedicated)
I/O		
VIO RATE	303	312
VIO/CMD	9.501	9.568
RIO RATE (V)	197	203
RIO/CMD (V)	6.177	6.225
MDC READS	174	181
MDC WRITES	84	87
MDC MODS	70	72
MDC HIT RATIO	0.92	0.92
PRIVOPs		
PRIVOP/CMD	17.765	17.897
DIAG/CMD	27.357	27.627
DIAG 08/CMD	0.753	0.767
DIAG 14/CMD	0.000	0.000
DIAG 58/CMD	1.254	1.227
DIAG 98/CMD	1.129	1.104
DIAG A4/CMD	3.825	3.895
DIAG A8/CMD	1.944	1.993
DIAG 214/CMD	12.573	12.788
SIE/CMD	54.589	54.801
SIE INTCPT/CMD	38.758	38.361
FREE TOTL/CMD	67.507	67.742
VTAM Machines		
WKSET (V)	1284	1275
TOT CPU/CMD (V)	4.5771	4.5647
CP CPU/CMD (V)	1.9822	1.9739
VIRT CPU/CMD (V)	2.5949	2.5908
DIAG 98/CMD (V)	1.149	1.122
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 45. VM/ESA Release 2 in Shared and Dedicated LPARs on 9121-480. Data for only one of the two LPAR systems in each case is presented; data for the partner LPAR in each case were similar. Unlike most other measurements in this report, these were done using the General Availability level of VM/ESA Release 2.

RACF

The following sections discuss the performance of the two RACF releases that run on VM/ESA Release 2: RACF 1.9.0 for VM and RACF 1.9.2 for VM.

- “9021-900 / RACF Release Comparison” on page 186 compares RACF 1.9.0 for VM and RACF 1.9.2 for VM, when running with security label checking not activated.
- “9021-900 / RACF Security Levels” on page 190 discusses the functional enhancement for security label checking. The activation of security labels is required if a system is configured at the DoD B1 level of trust.
- “9021-900 / RACF Multiple Servers” on page 194 discusses the RACF 1.9.2 for VM enhancement allowing more than one RACF service machine.

The system configurations used for the measurements discussed in these sections were not configured at either the C2 or B1 level of security. Information on security labels and DoD configured systems can be found in the *C2/B1 Trusted Facility Manual for VM/ESA with RACF*. The security label active measurements can be thought of as a simulation of a B1 environment.

All runs use the Teleprocessing Network Simulator (TPNS) to drive the workload. VTAM is used in this environment. VTAM was not identified as a component of either C2 or B1 system configuration. Other products, which were not present for these measurements, are required if the system was being configured as either a C2 or B1 system.

9021-900 / RACF Release Comparison

Workload: FS7B0R

Hardware Configuration

Processor model: 9021-900
 Processors used: 6
 Storage
 Real: 1024MB
 Expanded: 4096MB (400MB retained for MDC)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3380-D	3880-03	4	18	30	30			
3380-E	3880-03	4	2		4			
3380-A	3880-03	4		2	6			
3380-D	3880-23	4				16 R		4 R
3380-K	3990-02	4	10	10		60		
3390-2	3990-02	4				40 R		

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3745-410	4	44	56Kb
3088-02	1	NA	4.5MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVED	Other Options
VTAMXA	1	VTAM	64MB/XA	10000	350	QUICKDSP ON
VSCSn	6	VSCS	64MB/XA	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
RACFVM	1	RACF	16MB/370	10000	OFF	QUICKDSP ON
Users	11420	User	3MB/XC	100	OFF	

Note: The resident number of index blocks used for the runs with RACF active was 100. This is the default supplied in the data set name table (ICHRDSNT) shipped with RACF for VM. For more information on the Dataset Name Table see *System Programming Library: RACF*. The SMF CLOSE value in the SMF control file was set to 200. For more information on the SMF CLOSE value, see *Resource Access Control Facility (RACF) Program Directory for VM Installations*.

Measurement Discussion

This section summarizes the results of the measurements on the 9021-900 with VM/ESA Release 2 running the minidisk workload comparing RACF 1.9.0 for VM and 1.9.2 for VM. Both runs used the same control and audit parameters. RACF 1.9.2 for VM showed improvement in internal throughput and response time.

RACF was configured to provide default-level user identification, user authentication, and access control to minidisks and VM/ESA CP commands. Both runs used the default control and audit parameters. The default events for which RACF was called for authorization were APPCPWVL, DIAGNOSE code XçA0ç, DIAGNOSE code XçD4ç, DIAGNOSE code XçE4ç, LINK, LOGOFF, LOGON, MDISK, RSTDSEG, STORE class C, TAG, TRANSFER class D, TRANSFER class G, and TRSOURCE. User access to the system (LOGON) and unsuccessful accesses to minidisks were audited by RACF.

RACF APAR VM52545 is required to run RACF 1.9.0 for VM on the VM/ESA Release 2 platform.

The 9021-900 experienced a 0.7% increase in internal throughput rate (ITR (H)) and a 0.2 second decrease in external response time (AVG LAST (T)). CP/CMD (H) decreased 1.8% while EMUL/CMD (H) was equivalent.

Note: APAR VM54161 was not able to be included in either of these runs. Including it would have improved the response time of both runs and may have reduced the external response time difference. For more information on this APAR, refer to "IUCV Storage Management" on page 8.

VM RELEASE RACF RELEASE RUN ID	VM/ESA Release 2 RACF 1.9.0 for VM 265BB425	VM/ESA Release 2 RACF 1.9.2 for VM 265BB427
Environment		
REAL STORAGE	1024MB	1024MB
EXP. STORAGE	4096MB	4096MB
USERS	11420	11420
VTAMs	1	1
VSCSs	6	6
PROCESSORS	6	6
Response Time		
TRIV INT	0.075	0.092
NONTRIV INT	0.389	0.383
TOT INT	0.222	0.242
TOT INT ADJ	0.318	0.303
AVG FIRST (T)	0.770	0.602
AVG LAST (T)	0.932	0.735
Throughput		
AVG THINK (T)	26.17	26.10
ETR	565.43	499.48
ETR (T)	395.07	398.43
ETR RATIO	1.431	1.254
ITR (H)	431.50	434.41
ITR	103.08	90.90
EMUL ITR	166.46	145.62
ITRR (H)	1.000	1.007
ITRR	1.000	0.882
Proc. Usage		
PBT/CMD (H)	13.905	13.812
PBT/CMD	13.896	13.804
CP/CMD (H)	5.709	5.609
CP/CMD	5.290	5.195
EMUL/CMD (H)	8.196	8.203
EMUL/CMD	8.606	8.609

VM RELEASE RACF RELEASE RUN ID	VM/ESA Release 2 RACF 1.9.0 for VM 265BB425	VM/ESA Release 2 RACF 1.9.2 for VM 265BB427
Environment		
REAL STORAGE	1024MB	1024MB
EXP. STORAGE	4096MB	4096MB
USERS	11420	11420
VTAMs	1	1
VSCSs	6	6
PROCESSORS	6	6
Processor Util.		
TOTAL (H)	549.35	550.30
TOTAL	549.00	550.00
UTIL/PROC (H)	91.56	91.72
UTIL/PROC	91.50	91.67
TOTAL EMUL (H)	323.80	326.83
TOTAL EMUL	340.00	343.00
MASTER TOTAL (H)	95.76	96.08
MASTER TOTAL	96.00	96.00
MASTER EMUL (H)	30.84	31.06
MASTER EMUL	33.00	33.00
TVR(H)	1.70	1.68
TVR	1.61	1.60
Storage		
NUCLEUS SIZE (V)	2380KB	2384KB
TRACE TABLE (V)	1200KB	1200KB
WKSET (V)	57	58
PGBLPGS	209 k	209 k
PGBLPGS/USER	18.3	18.3
FREEPGS	30581	30465
FREE UTIL	0.89	0.89
SHRPGS	2584	2641
Paging		
READS/SEC	654	663
WRITES/SEC	442	444
PAGE/CMD	2.774	2.778
PAGE IO RATE (V)	141.600	143.300
PAGE IO/CMD (V)	0.358	0.360
XSTOR IN/SEC	3069	3093
XSTOR OUT/SEC	3730	3754
XSTOR/CMD	17.209	17.185
FAST CLR/CMD	6.616	6.596
Queues		
DISPATCH LIST	300.46	263.88
ELIGIBLE LIST	0.00	0.00
I/O		
VIO RATE	3495	3495
VIO/CMD	8.846	8.772
RIO RATE (V)	1006	1014
RIO/CMD (V)	2.546	2.545
MDC READS	2270	2279
MDC WRITES	1078	1093
MDC MODS	890	903
MDC HIT RATIO	0.92	0.92
PRIVOPs		
PRIVOP/CMD	15.285	15.294
DIAG/CMD	24.529	24.434
DIAG 08/CMD	0.772	0.773
DIAG 14/CMD	0.023	0.023
DIAG 58/CMD	1.397	1.280
DIAG 98/CMD	0.263	0.259
DIAG A4/CMD	3.976	3.973
DIAG A8/CMD	1.868	1.925
DIAG 214/CMD	13.359	13.348
SIE/CMD	48.092	47.688
SIE INTCPT/CMD	29.336	29.566
FREE TOTL/CMD	68.341	67.767

VM RELEASE RACF RELEASE RUN ID	VM/ESA Release 2 RACF 1.9.0 for VM 265BB425	VM/ESA Release 2 RACF 1.9.2 for VM 265BB427
Environment		
REAL STORAGE	1024MB	1024MB
EXP. STORAGE	4096MB	4096MB
USERS	11420	11420
VTAMs	1	1
VSCSs	6	6
PROCESSORS	6	6
VTAM Machines		
WKSET (V)	12228	11828
TOT CPU/CMD (V)	2.2838	2.2702
CP CPU/CMD (V)	1.2424	1.2208
VIRT CPU/CMD (V)	1.0414	1.0494
DIAG 98/CMD (V)	0.265	0.259
RACF Machines		
WKSET (V)	330	366
TOT CPU/CMD (V)	0.0576	0.0543
CP CPU/CMD (V)	0.0099	0.0098
VIRT CPU/CMD (V)	0.0477	0.0445
*RPI/CMD (V)	0.011	0.011
*RPI/sec (V)	4.5	4.4
*RPI/CPU sec (V)	198	203
CPU/*RPI (V)	5.0505	4.9261
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 46. VM/ESA Release 2 RACF Release Comparison on a 9021-900.

9021-900 / RACF Security Levels

Workload: FS7B0R

Hardware Configuration

Processor model: 9021-900
 Processors used: 6
 Storage
 Real: 1024MB
 Expanded: 4096MB (400MB retained for MDC)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3380-D	3880-03	4	18	30	30			
3380-E	3880-03	4	2		4			
3380-A	3880-03	4		2	6			
3380-D	3880-23	4				16 R		4 R
3380-K	3990-02	4	10	10		60		
3390-2	3990-02	4				40 R		

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3745-410	4	44	56Kb

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVED	Other Options
VTAMXA	1	VTAM	64MB/XA	10000	350	QUICKDSP ON
VSCSn	4	VSCS	64MB/XA	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
RACFVM	1	RACF	16MB/370	10000	OFF	QUICKDSP ON
Users	8000	User	3MB/XC	100	OFF	

Note: The resident number of index blocks that was used for the runs with RACF active was 100. This is the default supplied in the data set name table (ICHRDSNT) shipped with RACF for VM. For more information on the Dataset Name Table see *System Programming Library: RACF*. The SMF CLOSE value in the SMF control file was set to 200. For more information on the SMF CLOSE value, see *Resource Access Control Facility (RACF) Program Directory for VM Installations*.

Measurement Discussion

This section summarizes the results of the measurements on the 9021-900 with VM/ESA Release 2 for 8000 users running the minidisk workload. The three measurements discussed are VM/ESA Release 2 with no external security manager, VM/ESA Release 2 with RACF 1.9.2 for VM without security label (SECLABEL) processing activated, and VM/ESA Release 2 with RACF 1.9.2 for VM with security label (SECLABEL) processing activated.

RACF was configured to provide user identification and authentication, access control to minidisks and some VM/ESA CP commands. User access to the system (LOGON) and unsuccessful accesses to minidisks were audited by RACF.

- Run 265B8005 used the default control and audit parameters. The default events for which RACF was called for authorization were APPCPWVL, DIAGNOSE code X'00', DIAGNOSE code X'D4', DIAGNOSE code X'E4', LINK, LOGOFF, LOGON, MDISK, RSTDSEG, STORE class C, TAG, TRANSFER class D, TRANSFER class G, and TRSOURCE. Security labels (SECLABELS) were defined to both users and minidisks with SECLABEL checking enabled. SECLABEL processing was activated using the RACF set RACF options command (SETROPTS). SETROPTS settings relative to SECLABEL processing in effect were:
 - RACF SECLABEL class was active and RACLISTed
 - MLACTIVE FAIL
 - MLS FAIL
 - RACF VMMAC class was active
 - profiles for users and resources were assigned a SECLABEL

The activation of SECLABELS would be required if a system was being configured at the DoD B1 level of trust. Information on the RACF SETROPTS command can be found in *Resource Access Control Facility (RACF) Command Language Reference* and *Resource Access Control Facility (RACF) Security Administrator's Guide*.

- Run 265B8003 did not enable SECLABELS checking; all other control and audit parameters were the same as 265B8005.
- Run 265B8000 did not use control or audit parameters and RACF was not used.

The following three runs experienced similar external response times (AVG LAST (T)). The differences are believed to be run variation. The internal throughput rate (ITR(H)) decreased about 1% as security increased. This is from the combined effects of a total increased CP/CMD (H) of about 1% and EMUL/CMD(H) of about 0.7%. These were influenced by the total 9.6% increase in READS/SEC, with most of this paging increase occurring when security label checking was enabled.

VM RELEASE RACF RELEASE SECURITY LEVEL RUN ID	VM/ESA Release 2 NA NA 265B8000	VM/ESA Release 2 RACF 1.9.2 for VM no SECLABELS 265B8003	VM/ESA Release 2 RACF 1.9.2 for VM SECLABELS 265B8005
Environment			
REAL STORAGE	1024MB	1024MB	1024MB
EXP. STORAGE	4096MB	4096MB	4096MB
USERS	8000	8000	8000
VTAMs	1	1	1
VSCSs	6	6	6
PROCESSORS	6	6	6
Response Time			
TRIV INT	0.022	0.022	0.022
NONTRIV INT	0.140	0.137	0.137
TOT INT	0.103	0.101	0.101
TOT INT ADJ	0.099	0.097	0.096
AVG FIRST (T)	0.233	0.242	0.229
AVG LAST (T)	0.303	0.312	0.298
Throughput			
AVG THINK (T)	25.90	26.00	26.06
ETR	271.11	271.05	270.07
ETR (T)	282.48	283.39	282.75
ETR RATIO	0.960	0.956	0.955
ITR (H)	470.09	467.89	466.22
ITR	75.31	75.03	74.57
EMUL ITR	113.63	113.14	112.51
ITRR (H)	1.000	0.995	0.992
ITRR	1.000	0.996	0.990
Proc. Usage			
PBT/CMD (H)	12.764	12.824	12.869
PBT/CMD	12.744	12.774	12.838
CP/CMD (H)	4.668	4.698	4.716
CP/CMD	4.283	4.305	4.350
EMUL/CMD (H)	8.095	8.126	8.154
EMUL/CMD	8.461	8.469	8.488
Processor Util.			
TOTAL (H)	360.55	363.40	363.89
TOTAL	360.00	362.00	363.00
UTIL/PROC (H)	60.09	60.57	60.65
UTIL/PROC	60.00	60.33	60.50
TOTAL EMUL (H)	228.68	230.28	230.55
TOTAL EMUL	239.00	240.00	240.00
MASTER TOTAL (H)	64.90	65.30	65.48
MASTER TOTAL	65.00	65.00	65.00
MASTER EMUL (H)	22.47	22.43	22.36
MASTER EMUL	24.00	24.00	24.00
TVR(H)	1.58	1.58	1.58
TVR	1.51	1.51	1.51
Storage			
NUCLEUS SIZE (V)	2376KB	2384KB	2384KB
TRACE TABLE (V)	1200KB	1200KB	1200KB
WKSET (V)	64	64	64
PGBLPGS	223 k	223 k	223 k
PGBLPGS/USER	27.9	27.9	27.9
FREEPGS	21444	21531	21602
FREE UTIL	0.88	0.88	0.88
SHRPGS	2228	2234	2191
Paging			
READS/SEC	281	286	308
WRITES/SEC	111	108	113
PAGE/CMD	1.388	1.390	1.489
PAGE IO RATE (V)	13.700	13.200	14.300
PAGE IO/CMD (V)	0.048	0.047	0.051
XSTOR IN/SEC	1505	1493	1480
XSTOR OUT/SEC	1717	1702	1695
XSTOR/CMD	11.406	11.274	11.229
FAST CLR/CMD	6.489	6.641	6.649

VM RELEASE RACF RELEASE SECURITY LEVEL RUN ID	VM/ESA Release 2 NA NA 265B8000	VM/ESA Release 2 RACF 1.9.2 for VM no SECLABELS 265B8003	VM/ESA Release 2 RACF 1.9.2 for VM SECLABELS 265B8005
Environment			
REAL STORAGE	1024MB	1024MB	1024MB
EXP. STORAGE	4096MB	4096MB	4096MB
USERS	8000	8000	8000
VTAMs	1	1	1
VSCSs	6	6	6
PROCESSORS	6	6	6
Queues			
DISPATCH LIST	65.00	66.31	67.21
ELIGIBLE LIST	0.00	0.00	0.00
I/O			
VIO/CMD	8.585	8.610	8.654
RIO RATE (V)	626	627	630
RIO/CMD (V)	2.216	2.213	2.228
MDC READS	1613	1611	1623
MDC WRITES	770	767	766
MDC MODS	637	632	632
MDC HIT RATIO	0.92	0.92	0.92
PRIVOPs			
PRIVOP/CMD	17.759	17.839	17.814
DIAG/CMD	24.234	24.330	24.429
DIAG 08/CMD	0.750	0.755	0.757
DIAG 14/CMD	0.021	0.021	0.021
DIAG 58/CMD	1.232	1.242	1.224
DIAG 98/CMD	0.223	0.226	0.226
DIAG A4/CMD	3.954	3.938	3.965
DIAG A8/CMD	1.866	1.870	1.910
DIAG 214/CMD	13.265	13.261	13.315
SIE/CMD	46.021	45.874	45.976
SIE INTCPT/CMD	32.215	32.112	32.183
FREE TOTL/CMD	70.802	70.575	70.733
VTAM Machines			
WKSET (V)	10753	10591	10561
TOT CPU/CMD (V)	2.2059	2.2094	2.2143
CP CPU/CMD (V)	1.1401	1.1449	1.1474
VIRT CPU/CMD (V)	1.0658	1.0645	1.0669
DIAG 98/CMD (V)	0.224	0.227	0.226
RACF Machines			
WKSET (V)	NA	473	481
TOT CPU/CMD (V)	NA	0.0509	0.0550
CP CPU/CMD (V)	NA	0.0098	0.0079
VIRT CPU/CMD (V)	NA	0.0411	0.0471
*RPI/CMD (V)	NA	0.011	0.011
*RPI/sec (V)	NA	3.2	3.2
*RPI/CPU sec (V)	NA	222	206
CPU/*RPI (V)	NA	4.5045	4.8543
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM			

Table 47. VM/ESA Release 2 RACF SECLABELS Comparison on a 9021-900.

9021-900 / RACF Multiple Servers

Workload: FS7B0R

Hardware Configuration

Processor model: 9021-900
 Processors used: 6
 Storage
 Real: 1024MB
 Expanded: 4096MB (400MB retained for MDC)
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	PAGE	SPOOL	- Number of Volumes -			System
					TDSK	User	Server	
3380-D	3880-03	4	18	30	30			
3380-E	3880-03	4	2		4			
3380-A	3880-03	4		2	6			
3380-D	3880-23	4				16 R		4 R
3380-K	3990-02	4	10	10		60		
3390-2	3990-02	4				40 R		

Note: R or W next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively.

Communications:

Control Unit	Number	Lines per Control Unit	Speed
3745-410	4	44	56Kb
3088-02	1	NA	4.5MB

Software Configuration

Driver: TPNS
 Think time distribution: Bactrian
 CMS block size: 4KB

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVED	Other Options
VTAMXA	1	VTAM	64MB/XA	10000	350	QUICKDSP ON
VSCSn	6	VSCS	64MB/XA	10000	OFF	QUICKDSP ON
WRITER	1	CP MONITOR	2MB/XA	10000	OFF	QUICKDSP ON
SMART	1	RTM	16MB/370	10000	OFF	QUICKDSP ON
RACFVMn	1 or 4	RACF	16MB/370	10000	OFF	QUICKDSP ON
Users	11420	User	3MB/XC	100	OFF	

Note: The resident number of index blocks that was used for the runs with RACF active was 100. This is the default supplied in the data set name table (ICHRDSNT) shipped with RACF for VM. For more information on the Dataset Name Table see *System Programming Library: RACF*. The SMF CLOSE value in the SMF control file was set to 200. For more information on the SMF CLOSE value, see *Resource Access Control Facility (RACF) Program Directory for VM Installations*.

Measurement Discussion

This section summarizes the results of the measurements on the 9021-900 with VM/ESA Release 2 running the minidisk workload with SECLABELS activated comparing RACF 1.9.2 for VM with one and four RACF service machines.

RACF for VM was configured to provide default-level user identification, user authentication, and access control to minidisks and VM/ESA CP commands. Both runs used the default control and audit parameters and security labels (SECLABELS) defined to both users and minidisks with SECLABEL checking enabled. The default events for which RACF for VM was called for authorization were APPCPWVL, DIAGNOSE code XçA0ç, DIAGNOSE code XçD4ç, DIAGNOSE code XçE4ç, LINK, LOGOFF, LOGON, MDISK, RSTDSEG, STORE class C, TAG, TRANSFER class D, TRANSFER class G, and TRSOURCE. User access to the system (LOGON) and unsuccessful accesses to minidisks were audited by RACF for VM. The activation of SECLABELS would be required if a system was being configured at the DoD B1 level of trust.

The 9021-900 experienced similar internal throughput rate (ITR (H)) and external response time (AVG LAST (T)). The differences are believed to be run variation.

Note: APAR VM54161 was not able to be included in either of these runs. Including it would have improved the response time of both runs. For more information on this APAR, refer to "IUCV Storage Management" on page 8.

In addition to the performance improvement, multiple RACF service machines may provide:

- Virtual storage

On VM, most security-relevant processing is being performed within the RACF virtual machine. Adding an additional RACF service machine on systems that have large numbers of concurrently logged on users increases the effective virtual storage that is used to maintain security-relevant information.

- Availability

If one of the RACF service machines should become disabled, the remaining service machine(s) may continue to provide security services.

- Improved throughput capacity for security-relevant requests

Multiple RACF service machines are able to process more security-relevant requests than a single RACF service machine. Note that processing more security relevant work does not necessarily translate into improved end user response time or an improved calculated system throughput. Processing a larger security-relevant workload does mean that end users should wait less time while attempting resource access that involves RACF services.

The following considerations can help to determine if more than a single RACF service machine may be needed:

- High frequency of RACF commands

Installations having applications that utilize RACF commands that are frequently invoked by large numbers of users should consider an additional RACF service machine.

- RACROUTE service

Installations having products or applications that employ the full function RACROUTE interface (which was introduced in RACF 1.9.0 for VM) should consider dedicating a RACF service machine to process RACROUTE requests. (Dedicating a service machine is described in the *System Programming Library: RACF* .)

- Discrete profiles for minidisks with large access lists

Installations should consider the use of an additional RACF service machine when the majority of the minidisks have been defined with discrete profiles having large access lists.

- Auditing

For installations that typically audit a large number of security-relevant events, additional RACF service machines may help improve the distribution of the overhead in recording the audit records.

VM RELEASE RACF RELEASE RACF SERVER(S) RUN ID	VM/ESA Release 2 RACF 1.9.2 for VM 1 265BB420	VM/ESA Release 2 RACF 1.9.2 for VM 4 265BB422
Environment		
REAL STORAGE	1024MB	1024MB
EXP. STORAGE	4096MB	4096MB
USERS	11420	11420
VTAMs	1	1
VSCSs	6	6
PROCESSORS	6	6
Response Time		
TRIV INT	0.072	0.068
NONTRIV INT	0.402	0.385
TOT INT	0.219	0.213
TOT INT ADJ	0.334	0.316
AVG FIRST (T)	0.855	0.832
AVG LAST (T)	1.037	1.010
Throughput		
AVG THINK (T)	26.15	26.17
ETR	601.30	584.86
ETR (T)	393.95	394.67
ETR RATIO	1.526	1.482
ITR (H)	429.03	430.70
ITR	109.28	106.56
EMUL ITR	176.86	172.19
ITRR (H)	1.000	1.004
ITRR	1.000	0.975
Proc. Usage		
PBT/CMD (H)	13.985	13.931
PBT/CMD	13.987	13.936
CP/CMD (H)	5.766	5.727
CP/CMD	5.356	5.321
EMUL/CMD (H)	8.219	8.204
EMUL/CMD	8.631	8.615
Processor Util.		
TOTAL (H)	550.94	549.81
TOTAL	551.00	550.00
UTIL/PROC (H)	91.82	91.64
UTIL/PROC	91.83	91.67
TOTAL EMUL (H)	323.80	323.77
TOTAL EMUL	340.00	340.00
MASTER TOTAL (H)	95.95	95.82
MASTER TOTAL	96.00	96.00
MASTER EMUL (H)	30.74	30.40
MASTER EMUL	33.00	32.00
TVR(H)	1.70	1.70
TVR	1.62	1.62

VM RELEASE RACF RELEASE RACF SERVER(S) RUN ID	VM/ESA Release 2 RACF 1.9.2 for VM 1 265BB420	VM/ESA Release 2 RACF 1.9.2 for VM 4 265BB422
Environment		
REAL STORAGE	1024MB	1024MB
EXP. STORAGE	4096MB	4096MB
USERS	11420	11420
VTAMs	1	1
VSCSs	6	6
PROCESSORS	6	6
Storage		
NUCLEUS SIZE (V)	2384KB	2384KB
TRACE TABLE (V)	1200KB	1200KB
WKSET (V)	57	58
PGBLPGS	209 k	209 k
PGBLPGS/USER	18.3	18.3
FREEPGS	30508	30534
FREE UTIL	0.89	0.89
SHRPGS	2672	2629
Paging		
READS/SEC	642	650
WRITES/SEC	446	451
PAGE/CMD	2.762	2.790
XSTOR IN/SEC	3043	3075
XSTOR OUT/SEC	3708	3742
XSTOR/CMD	17.137	17.273
FAST CLR/CMD	6.653	6.646
Queues		
DISPATCH LIST	323.56	292.00
ELIGIBLE LIST	0.02	0.00
I/O		
VIO RATE	3480	3477
VIO/CMD	8.834	8.810
MDC READS	2257	2251
MDC WRITES	1076	1075
MDC MODS	888	887
MDC HIT RATIO	0.92	0.92
PRIVOPs		
PRIVOP/CMD	15.173	15.290
DIAG/CMD	24.570	24.511
DIAG 08/CMD	0.774	0.773
DIAG 14/CMD	0.023	0.023
DIAG 58/CMD	1.363	1.300
DIAG 98/CMD	0.256	0.261
DIAG A4/CMD	3.973	3.958
DIAG A8/CMD	1.901	1.954
DIAG 214/CMD	13.405	13.368
SIE/CMD	48.230	48.142
SIE INTCPT/CMD	29.420	29.366
FREE TOTL/CMD	68.537	68.412
VTAM Machines		
WKSET (V)	12281	12001
TOT CPU/CMD (V)	2.2664	2.2894
CP CPU/CMD (V)	1.2389	1.2502
VIRT CPU/CMD (V)	1.0275	1.0392
DIAG 98/CMD (V)	0.259	0.262

VM RELEASE RACF RELEASE RACF SERVER(S) RUN ID	VM/ESA Release 2 RACF 1.9.2 for VM 1 265BB420	VM/ESA Release 2 RACF 1.9.2 for VM 4 265BB422
Environment		
REAL STORAGE	1024MB	1024MB
EXP. STORAGE	4096MB	4096MB
USERS	11420	11420
VTAMs	1	1
VSCSs	6	6
PROCESSORS	6	6
RACF Machines		
WKSET Total (V)	368	1074
WKSET RACFVM (V)	368	271
WKSET RACFVM1 (V)	NA	267
WKSET RACFVM2 (V)	NA	257
WKSET RACFVM3 (V)	NA	279
TOT CPU/CMD (V)	0.0591	0.0632
CP CPU/CMD (V)	0.0098	0.0193
VIRT CPU/CMD (V)	0.0493	0.0439
*RPI/CMD (V)	0.011	0.011
*RPI/sec Total (V)	4.5	4.4
*RPI/sec RACFVM (V)	4.455	1.013
*RPI/sec RACFVM1 (V)	NA	1.235
*RPI/sec RACFVM2 (V)	NA	1.086
*RPI/sec RACFVM3 (V)	NA	1.113
*RPI/CPU sec (V)	193	176
CPU/*RPI (V)	5.1813	5.6818
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 48. VM/ESA Release 2 RACF 1.9.2 for VM Multiple Server Comparison on a 9021-900.

VSE Guest Using Shared DASD

This section examines VSE/ESA 1.2.0 guest performance running under VM/ESA Release 2 to illustrate the performance characteristics of shared versus dedicated DASD in one- and two-guest environments. Two different workloads were used for the measurements: PACE and VSECICS. The PACE workload is batch-only and is characterized by heavy I/O. Two variants, PACEX6 and PACEX8, differ in the volume of work done. The VSECICS workload simulates terminal users in a CICS* environment and is characterized by light I/O. (See Appendix B, "Workloads" on page 223 for detailed descriptions of the workloads.) The VSE guests included V=R and V=V machines with dedicated DASD, DASD configured as full pack minidisks, and multiple guest machines with shared DASD volumes. See "I/O Assist for Guests" on page 150 for a discussion of issues relevant to these environments.

Workloads: PACEX8, PACEX6, and VSECICS

Hardware Configuration

Processor models: 9221-200 and 9221-170
 Processors used: 2 and 1, respectively
 Storage
 Real: 256MB
 Expanded: None
 Tape: 3480 (MONITOR)

DASD:

Type of DASD	Control Unit	Number of Paths	- Number of Volumes -					
			PAGE	SPOOL	TDSK	VSAM	VSE Sys.	VM Sys.
3380	3880	2					2	
3380	3880	2				10		
3380	3880	2						4

Note: For the measurements with 2 VSE guests, twice the number of VSAM and VSE system volumes were used, all of the same type and number of channel paths.

Software Configuration

VSE version: 1.2.0

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	RELSHARE	RESERVED	Other Options
USERVR	1	VSE V=R	24MB/ESA	100	OFF	IOASSIST ON CCWTRANS OFF
SMART	1	RTM	16MB/370	100	OFF	
WRITER	1	CP MONITOR	2MB/XA	100	OFF	

For the two-guest measurements, the following was used also:
 USERVF3 1 VSE V=F 24MB/ESA 100 OFF IOASSIST ON

Note: EXPLORE/VSE active during measurements of the VSECICS workload.

Single Guest, PACEX8 Workload

Figure 25 and Figure 26 summarize the PACEX8 workload running as a single VSE/ESA* guest of VM/ESA Release 2 on a 9221-200 processor.




Figure 25. Internal Throughput for Single VSE Guest on 9221-200. PACEX8 workload on VSE under VM/ESA Release 2.




Figure 26. External Throughput for Single VSE Guest on 9221-200. PACEX8 workload on VSE under VM/ESA Release 2.

The V=R guest using dedicated DASD had the best internal throughput and shortest batch run time. The V=R guest was also run with the VSAM work file DASD defined as full pack minidisks (FPMD). With DASD defined as FPMD, the batch run time increased by 20% and the internal throughput decreased 20%. The V=V guest machine with dedicated DASD, compared to the V=R guest with dedicated DASD, had a 28% increase in batch run time with a 15% decrease in internal throughput. Additional cost for defining the V=V guest machine DASD

as full pack minidisks is negligible. Utilization per processor decreased from 50% with a V=R guest and dedicated DASD, to 46% with a V=V guest. See Table 49 on page 201 for detailed data for these measurements.

Note the significant differences in the I/O rates reported for the V=R dedicated case versus the others. The VIO and RIO rates reported do not count the assisted DASD I/Os to the dedicated devices. Otherwise they would be more in line with the V=R FPMD and V=V cases.

Virtual Machine Type	V = R		V = V	
	DASD Usage		DASD Usage	
RELEASE RUN ID	dedicated VM/ESA 2 VR1200D8	FPMD VM/ESA 2 VR1200F8	dedicated VM/ESA 2 VV1200D8	FPMD VM/ESA 2 VV1200F8
Environment				
IML MODE	ESA	ESA	ESA	ESA
REAL STORAGE	256MB	256MB	256MB	256MB
EXP. STORAGE	0MB	0MB	0MB	0MB
VM MODE	ESA	ESA	ESA	ESA
VM SIZE	24MB	24MB	24MB	24MB
GUEST SETTING	V = R	V = R	V = V	V = V
VSE SUPERVISOR	ESA	ESA	ESA	ESA
PROCESSORS	2	2	2	2
Throughput (Min)				
Elapsed Time (H)	603.5	726.0	773.6	790.0
ETR (H)	5.57	4.63	4.34	4.25
ITR (H)	7.45	5.06	4.70	4.64
ITR	5.51	4.41	4.67	4.57
ITRR (H)	1.000	0.679	0.631	0.623
ITRR	1.000	0.800	0.847	0.830
Proc. Usage (Sec)				
PBT/CMD (H)	8.051	11.862	12.755	12.926
PBT/CMD	10.885	13.613	12.848	13.119
CP/CMD (H)	1.904	5.393	5.858	6.061
CP/CMD	0.323	4.149	4.697	4.796
EMUL/CMD (H)	6.147	6.469	6.897	6.865
EMUL/CMD	10.562	9.465	8.151	8.323
Processor Util.				
TOTAL (H)	74.70	91.49	92.33	91.63
TOTAL	101.00	105.00	93.00	93.00
UTIL/PROC (H)	37.35	45.74	46.16	45.81
UTIL/PROC	50.50	52.50	46.50	46.50
TOTAL EMUL (H)	57.03	49.89	49.92	48.67
TOTAL EMUL	98.00	73.00	59.00	59.00
MASTER TOTAL (H)	37.35	45.74	46.16	45.81
MASTER TOTAL	1.30	5.30	21.00	14.00
MASTER EMUL (H)	76.35	54.54	54.07	53.11
MASTER EMUL	0.24	0.24	11.00	5.90
TVR(H)	1.31	1.83	1.85	1.88
TVR	1.03	1.44	1.58	1.58
Storage				
NUCLEUS SIZE (V)	2448KB	2448KB	2452KB	2452KB
TRACE TABLE (V)	128KB	128KB	128KB	128KB
PGBLPGS	14938	14937	31598	31575
FREEPGS	89	91	86	86
FREE UTIL	0.50	0.49	0.52	0.53
SHRPGS	1062	1062	934	934
Paging				
PAGE/CMD	0.000	0.000	0.000	0.000
XSTOR/CMD	0.000	0.000	0.000	0.000
FAST CLR/CMD	0.000	0.000	82.889	98.746

Virtual Machine Type	V = R		V = V	
DASD Usage	dedicated	FPMD	dedicated	FPMD
RELEASE RUN ID	VM/ESA 2 VR1200D8	VM/ESA 2 VR1200F8	VM/ESA 2 VV1200D8	VM/ESA 2 VV1200F8
Environment				
IML MODE	ESA	ESA	ESA	ESA
REAL STORAGE	256MB	256MB	256MB	256MB
EXP. STORAGE	0MB	0MB	0MB	0MB
VM MODE	ESA	ESA	ESA	ESA
VM SIZE	24MB	24MB	24MB	24MB
GUEST SETTING	V = R	V = R	V = V	V = V
VSE SUPERVISOR	ESA	ESA	ESA	ESA
PROCESSORS	2	2	2	2
I/O				
VIO RATE	4.000 *	219.000	227.000	225.000
VIO/CMD	43.110 *	2839.373	3135.980	3173.972
RIO RATE (V)	1.000 *	217.000	222.000	221.000
RIO/CMD (V)	10.777 *	2813.442	3066.906	3117.546
DASD IO TOTAL (V)	172520	173905	173069	172200
DASD IO RATE (V)	287.53	241.53	221.88	220.77
DASD IO/CMD (V)	3098.87	3131.54	3065.29	3114.29
PRIVOPs				
PRIVOP/CMD (R)	52.010	2849.143	3135.455	3180.462
DIAG/CMD (R)	222.702	228.017	229.023	232.693
SIE/CMD	538.871	9399.749	10803.245	10904.358
SIE INTCPT/CMD	501.150	9399.749	10047.018	10250.096
FREE TOTL/CMD	732.865	6819.680	7059.409	7420.042
Note: V=VMPRF, H=Hardware Monitor, Unmarked=RTM. Per command (/CMD), for this workload, means per batch job. I/O-related data marked with an asterisk(*) is as reported, but does not include assisted I/O for dedicated devices.				

Table 49. Single VSE Guest, PACEX8 Workload on 9221-200

The same measurement configurations were used on a 9221-170 and are summarized in Figure 27 and Figure 28.




Figure 27. Internal Throughput for Single VSE Guest on 9221-170. PACEX8 workload on VSE under VM/ESA Release 2.




Figure 28. External Throughput for Single VSE Guest on 9221-170. PACEX8 workload on VSE under VM/ESA Release 2.

As was the case for the 9221-200, the V=R guest with dedicated DASD on a 9221-170 had the highest internal throughput and shortest batch run time. The V=R machine with FPMD DASD yielded a 39% loss in internal throughput compared to the same guest with dedicated DASD. I/O assist was not available for the full pack minidisks. Also fast path CCW translation was not used—note the disproportionate number of CP free storage requests (FREE TOTL/CMD) indicating CCW fast path benefit was not in effect. The V=V guest with or without

dedicated DASD had a 30% decrease in internal throughput compared to the V=R guest machine with dedicated DASD.

The V=V configuration had a 25% increase in total run time compared to the V=R guest with dedicated DASD. Processor busy increased from 76% for the V=R guest with dedicated DASD, to 86% for the V=V guest. There was a negligible difference between the dedicated and FPMD cases for the V=V guest machine. Data for the 9221-170 runs are provided in Table 50.

Virtual Machine Type	V = R		V = V	
	dedicated	FPMD	dedicated	FPMD
RELEASE RUN ID	VM/ESA 2 VR1170D1	VM/ESA 2 VR1170P3	VM/ESA 2 VV1170P2	VM/ESA 2 VV1170P1
Environment				
IML MODE	ESA	ESA	ESA	ESA
REAL STORAGE	256MB	256MB	256MB	256MB
EXP. STORAGE	0MB	0MB	0MB	0MB
VM MODE	ESA	ESA	ESA	ESA
VM SIZE	24MB	24MB	24MB	24MB
VSE SUPERVISOR	ESA	ESA	ESA	ESA
PROCESSORS	1	1	1	1
Throughput (Min)				
Elapsed Time (H)	657.2	968.6	814.8	826.8
ETR (H)	5.11	3.47	4.12	4.06
ITR (H)	6.73	4.12	4.76	4.73
ITR	6.64	4.08	4.69	4.62
ITRR (H)	1.000	0.612	0.707	0.702
ITRR	1.000	0.615	0.706	0.696
Proc. Usage (Sec)				
PBT/CMD (H)	8.914	14.567	12.618	12.692
PBT/CMD	9.037	14.702	12.804	12.992
CP/CMD (H)	2.769	8.244	5.907	5.972
CP/CMD	1.056	7.091	4.656	4.724
EMUL/CMD (H)	6.145	6.323	6.711	6.720
EMUL/CMD	7.981	7.610	8.148	8.268
Processor Util.				
TOTAL (H)	75.96	84.22	86.72	85.97
TOTAL	77.00	85.00	88.00	88.00
TOTAL EMUL (H)	52.36	36.56	46.12	45.52
TOTAL EMUL	68.00	44.00	56.00	56.00
TVR(H)	1.45	2.30	1.88	1.89
TVR	1.13	1.93	1.57	1.57
Storage				
NUCLEUS SIZE (V)	2448KB	2448KB	2452KB	2452KB
TRACE TABLE (V)	64KB	64KB	64KB	64KB
PGBLPGS	14977	14975	31627	31621
FREEPGS	77	79	80	80
FREE UTIL	0.53	0.54	0.55	0.55
SHRPGS	928	930	930	912
Paging				
PAGE/CMD	0.000	0.000	0.000	0.000
XSTOR/CMD	0.000	0.000	0.000	0.000
FAST CLR/CMD	0.000	0.000	87.299	88.583
I/O				
VIO RATE	4.000 *	164.000	217.000	216.000
VIO/CMD	46.945 *	2836.543	3157.314	3188.974
RIO RATE (V)	1.000 *	162.000	207.000	207.000
RIO/CMD (V)	11.736 *	2801.951	3011.815	3056.100
DASD IO TOTAL (V)	173171	173376	173560	173764
DASD IO RATE (V)	262.38	180.60	206.62	206.86
DASD IO/CMD (V)	3079.39	3123.66	3006.27	3054.06

Virtual Machine Type	V = R		V = V	
DASD Usage	dedicated	FPMD	dedicated	FPMD
RELEASE RUN ID	VM/ESA 2 VR1170D1	VM/ESA 2 VR1170P3	VM/ESA 2 VV1170P2	VM/ESA 2 VV1170P1
Environment				
IML MODE	ESA	ESA	ESA	ESA
REAL STORAGE	256MB	256MB	256MB	256MB
EXP. STORAGE	0MB	0MB	0MB	0MB
VM MODE	ESA	ESA	ESA	ESA
VM SIZE	24MB	24MB	24MB	24MB
VSE SUPERVISOR	ESA	ESA	ESA	ESA
PROCESSORS	1	1	1	1
PRIVOPs				
PRIVOP/CMD	57.568	2844.086	3162.487	3200.915
DIAG/CMD	222.954	236.884	230.714	232.265
SIE/CMD	1854.346	9357.133	11159.722	11264.755
SIE INTCPT/CMD	1316.585	7766.420	8927.777	8899.156
FREE TOTL/CMD	1596.146	27206.600	7100.318	7440.939
Note: V=VMPRF, H=Hardware Monitor, Unmarked=RTM. Per command (/CMD), for this workload, means per batch job. I/O-related data marked with an asterisk(*) is as reported, but does not include assisted I/O for dedicated devices.				

Table 50. Single VSE Guest, PACEX8 Workload on 9221-170

Single GUEST, VSECICS WORKLOAD




Figure 29. Internal Throughput for VSE CICS Workload on 9221-170

The VSECICS workload measurements of a single guest of VM/ESA Release 2 on the 9221-170 processor are summarized in Figure 29. The measurements were used to examine the performance effects for a CICS workload when running a guest machine with dedicated DASD and with full pack minidisks. The V=R guest using full pack minidisks achieved about 85% of the internal throughput rate of the V=R dedicated DASD configuration environment. Very little difference in throughput rates occurred when running the V=V machine with dedicated DASD or full pack minidisks. There was less of a percentage ITR difference between the V=R dedicated DASD case and the other cases when compared to the PACE measurements discussed previously. This is attributed to the lighter I/O burden for the VSECICS workload. Table 51 provides the detailed data for these measurements.

Virtual Machine Type	V = R		V = V	
	dedicated	FPMD	dedicated	FPMD
RELEASE RUN ID	VM/ESA 2 VR1170DA	VM/ESA 2 VR1170FA	VM/ESA 2 VV1170DA	VM/ESA 2 VV1170FA
Environment				
RUN MODE	ESA	ESA	ESA	ESA
REAL STORAGE	256MB	256MB	256MB	256MB
EXP. STORAGE	0MB	0MB	0MB	0MB
USERS	328	276	276	276
VM MODE	ESA	ESA	ESA	ESA
VM SIZE	24MB	24MB	24MB	24MB
VSE MODE	ESA	ESA	ESA	ESA
PROCESSORS	1	1	1	1
Throughput/Response Time				
AVG RESP (C)	0.699	0.697	0.765	0.790
ETR (C)	21.43	18.05	17.95	17.94
ITR (H)	23.61	20.00	19.68	19.82
ITRR (H)	1.000	0.847	0.834	0.840

Virtual Machine Type	V = R		V = V	
DASD Usage	dedicated	FPMD	dedicated	FPMD
RELEASE RUN ID	VM/ESA 2 VR1170DA	VM/ESA 2 VR1170FA	VM/ESA 2 VV1170DA	VM/ESA 2 VV1170FA
Environment				
RUN MODE	ESA	ESA	ESA	ESA
REAL STORAGE	256MB	256MB	256MB	256MB
EXP. STORAGE	0MB	0MB	0MB	0MB
USERS	328	276	276	276
VM MODE	ESA	ESA	ESA	ESA
VM SIZE	24MB	24MB	24MB	24MB
VSE MODE	ESA	ESA	ESA	ESA
PROCESSORS	1	1	1	1
Proc. Usage				
PBT/CMD (H)	42.363	49.989	50.805	50.454
PBT/CMD	42.468	49.858	50.687	50.725
CP/CMD (H)	4.260	10.964	8.836	8.577
CP/CMD	1.867	8.864	6.127	6.132
EMUL/CMD (H)	38.104	39.025	41.969	41.878
EMUL/CMD	40.602	40.995	44.560	44.593
Processor Util.				
TOTAL (H)	90.78	90.24	91.21	90.51
TOTAL	91.00	90.00	91.00	91.00
TOTAL EMUL (H)	81.65	70.45	75.35	75.13
TOTAL EMUL	87.00	74.00	80.00	80.00
TVR(H)	1.11	1.28	1.21	1.20
TVR	1.05	1.22	1.14	1.14
Storage				
NUCLEUS SIZE (V)	2448KB	2448KB	2452KB	2452KB
TRACE TABLE (V)	64KB	64KB	64KB	64KB
PGBLPGS	14952	14944	31611	31609
FREEPGS	78	83	79	79
FREE UTIL	0.54	0.55	0.56	0.57
SHRPGS	930	930	913	930
Paging				
PAGE/CMD	0.000	0.000	0.000	0.000
XSTOR/CMD	0.000	0.000	0.000	0.000
FAST CLR/CMD	0.000	0.000	3.119	2.954
I/O				
VIO RATE	67.12 (EX)	56	56	55
VIO/CMD	3.13 (EX)	3.102	3.119	3.066
RIO RATE (V)	1 *	56	57	56
RIO/CMD (V)	0.047 *	3.102	3.175	3.122
PRIVOPs				
PRIVOP/CMD	0.053	3.173	3.177	3.131
DIAG/CMD	3.467	3.320	3.290	3.171
SIE/CMD	7.280	16.342	20.386	20.011
SIE INTCPT/CMD	4.659	11.603	12.028	11.807
FREE TOTL/CMD	1.913	31.799	8.578	8.027
<p>Note: C=CICSPARS, V=VMPRF(ESA), H=Hardware Monitor, EX=EXPLORE/VSE Unmarked=RTM. Per command (/CMD), for this workload, means per CICS transaction. I/O-related data marked with an asterisk(*) is as reported, but does not include assisted I/O for dedicated devices.</p>				

Table 51. Single VSE Guest, VSECICS Workload on 9221-170

Two Guests, PACEX8 Workload

In this section, measurements that used 2 VSE guests of VM/ESA Release 2 running on a 9221-200 processor are described.




Figure 30. Internal Throughput for Two VSE Guests on 9221-200. PACEX8 workload on VSE under VM/ESA Release 2.




Figure 31. External Throughput for Two VSE Guests on 9221-200. PACEX8 workload on VSE under VM/ESA Release 2.

Figure 30 and Figure 31 summarize the ITRs and ETRs, respectively, for 2 VSE guests running the PACEX8 workload. In the multiple preferred guests (MPG) environments, one guest was run V=R and the other V=F. Comparisons were made between 2 guests, each with its own dedicated DASD, and 2 guests sharing DASD (using the VSE lock file) for VSAM work files. With one V=R guest and one V=F guest sharing the VSAM work file DASD, the external throughput rate decreased 47% and the internal throughput rate decreased 39%.

However, the number of DASD used for the shared environment was reduced by 58%. Performance of multiple V=V guest sharing DASD was comparable to the multiple preferred guest environment. Table 52 on page 211 provides the detailed data for these measurements. The multiple V=V guest with dedicated-DASD environment was not measured on the 9221-200 because there was an elongation of job completion times causing problems in interpreting the measurement results.

Virtual Machine Types	V=R, V=F		V=V, V=V
DASD Usage	dedicated	shared	shared
RELEASE RUN ID	VM/ESA 2 VR2200D8	VM/ESA 2 VR2200S8	VM/ESA 2 VV2200S8
Environment			
IML MODE	ESA	ESA	ESA
REAL STORAGE	256MB	256MB	256MB
EXP. STORAGE	0MB	0MB	0MB
VM MODE	ESA	ESA	ESA
VM SIZE	24MB	24MB	24MB
VSE SUPERVISOR	ESA	ESA	ESA
PROCESSORS	2	2	2
Throughput (Min)			
Elapsed Time (H)	748.7	1091.8	1143.9
ETR (H)	8.98	6.15	5.87
ITR (H)	na	4.24	3.75
ITR	5.91	3.60	3.72
ITRR	1.000	0.610	0.630
Proc. Usage (Sec)			
PBT/CMD (H)	na	na	16.006
PBT/CMD	10.161	16.670	16.138
CP/CMD (H)	na	na	7.930
CP/CMD	0.602	4.972	6.026
EMUL/CMD (H)	na	na	8.076
EMUL/CMD	9.559	11.698	10.112
Processor Util.			
TOTAL (H)	na	na	156.71
TOTAL	152.00	171.00	158.00
UTIL/PROC (H)	na	na	78.35
UTIL/PROC	76.00	85.50	79.00
TOTAL EMUL (H)	na	na	79.07
TOTAL EMUL	143.00	120.00	99.00
MASTER TOTAL (H)	na	na	78.35
MASTER TOTAL	53.00	71.00	78.00
MASTER EMUL (H)	na	na	50.46
MASTER EMUL	45.00	43.00	48.00
TVR(H)	na	na	1.98
TVR	1.06	1.43	1.60
Storage			
NUCLEUS SIZE (V)	2448KB	2452KB	2452KB
TRACE TABLE (V)	128KB	128KB	128KB
PGBLPGS	14931	14929	31578
FREPGS	90	91	94
FREE UTIL	0.52	0.53	0.53
SHRPGS	934	934	934
Paging			
PAGE/CMD	0.000	0.000	0.000
XSTOR/CMD	0.000	0.000	0.000
FAST CLR/CMD	0.000	0.000	91.923

Virtual Machine Types	V=R, V=F		V=V, V=V
DASD Usage	dedicated	shared	shared
RELEASE RUN ID	VM/ESA 2 VR2200D8	VM/ESA 2 VR2200S8	VM/ESA 2 VV2200S8
Environment			
IML MODE	ESA	ESA	ESA
REAL STORAGE	256MB	256MB	256MB
EXP. STORAGE	0MB	0MB	0MB
VM MODE	ESA	ESA	ESA
VM SIZE	24MB	24MB	24MB
VSE SUPERVISOR	ESA	ESA	ESA
PROCESSORS	2	2	2
I/O			
VIO RATE	6.000	314.000	329.000
VIO/CMD	40.108	3060.966	3360.303
RIO RATE (V)	1.000	326.000	326.000
RIO/CMD (V)	6.685	3177.946	3329.662
DASD IO TOTAL (V)	345549	370617	370617
DASD IO RATE (V)	443.01	325.10	325.10
DASD IO/CMD (V)	2961.40	3169.20	3320.50
PRIVOPs			
PRIVOP/CMD	40.907	3071.634	3363.132
DIAG/CMD	179.364	219.569	220.363
SIE/CMD	1042.813	10625.647	12327.919
SIE INTCPT/CMD	813.394	9244.312	10108.894
FREE TOTL/CMD	1036.128	9104.912	9131.036
Note: V=VMPRF, H=Hardware Monitor, Unmarked=RTM. Per command (/CMD), for this workload, means per batch job.			

Table 52. Two VSE Guests, PACEX8 Workload on 9221-200

Two Guests, PACEX6 Workload

The PACEX6 workload was used to measure the environments under study. The results are summarized in Figure 32 and Figure 33.

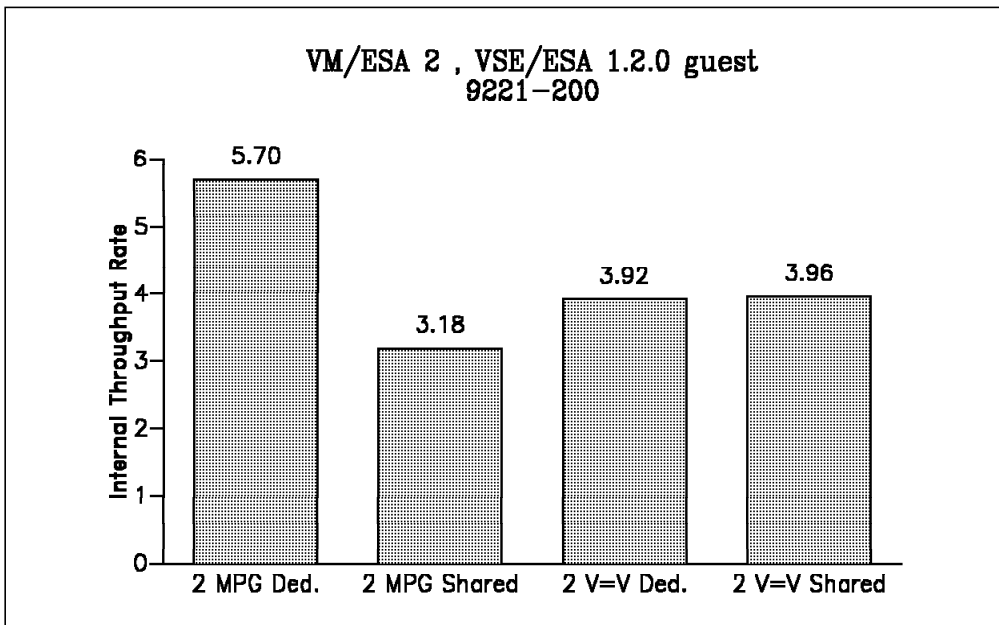


Figure 32. Internal Throughput for Two VSE Guests on 9221-200. PACEX6 workload on VSE under VM/ESA Release 2.




Figure 33. External Throughput for Two VSE Guests on 9221-200. PACEX6 workload on VSE under VM/ESA Release 2.

Multiple V=V guest machines running the PACEX6 workload with dedicated DASD yielded a 19% increase in batch run time with a 31% reduction in internal throughput compared to the multiple preferred guest environment with dedicated DASD. I/O assist was not available for the MPG shared measurement. Also, fast path CCW translation was not used—note the disproportionate number of CP free

storage requests (FREE TOTL/CMD) indicating CCW fast path benefit was not in effect. Data for these measurements are provided in Table 53 on page 215.

Virtual Machine Types	V=R, V=F		V=V, V=V	
DASD Usage	dedicated	shared	dedicated	shared
RELEASE RUN ID	VM/ESA 2 VR2200D6	VM/ESA 2 VR2200S6	VM/ESA 2 VV2200D6	VM/ESA 2 VV2200S6
Environment				
IML MODE	ESA	ESA	ESA	ESA
REAL STORAGE	256MB	256MB	256MB	256MB
EXP. STORAGE	0MB	0MB	0MB	0MB
VM MODE	ESA	ESA	ESA	ESA
VM SIZE	24MB	24MB	24MB	24MB
VSE SUPERVISOR	ESA	ESA	ESA	ESA
PROCESSORS	2	2	2	2
Throughput (Min)				
Elapsed Time (H)	626.9	1024.0	744.1	1010.8
ETR (H)	8.04	4.92	6.77	4.99
ITR (H)	na	na	3.95	3.96
ITR	5.70	3.18	3.92	na
ITRR	1.000	0.557	0.687	na
Proc. Usage (Sec)				
PBT/CMD (H)	na	na	15.186	15.146
PBT/CMD	10.522	18.895	15.325	na
CP/CMD (H)	na	na	7.257	7.411
CP/CMD	0.672	4.632	5.403	na
EMUL/CMD (H)	na	na	7.929	7.735
EMUL/CMD	9.851	14.263	9.921	na
Processor Util.				
TOTAL (H)	na	na	171.44	125.87
TOTAL	141.00	155.00	173.00	na
UTIL/PROC (H)	na	na	85.72	62.94
UTIL/PROC	70.50	77.50	86.50	na
TOTAL EMUL (H)	na	na	89.51	64.28
TOTAL EMUL	132.00	117.00	112.00	na
MASTER TOTAL (H)	na	na	85.72	62.94
MASTER TOTAL	41.00	55.00	85.00	na
MASTER EMUL (H)	na	na	52.21	51.07
MASTER EMUL	33.00	33.00	56.00	na
TVR(H)	na	na	1.92	1.96
TVR	1.07	1.32	1.54	na
Storage				
NUCLEUS SIZE (V)	2448KB	2448KB	2452KB	2452KB
TRACE TABLE (V)	128KB	128KB	128KB	128KB
PGBLPGS	14936	14930	31588	na
FREEPGS	86	90	91	na
FREE UTIL	0.54	0.53	0.53	na
SHRPGS	932	934	933	na
Paging				
PAGE/CMD	0.000	0.000	0.000	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	0.000	0.000	88.582	na
I/O				
VIO RATE	4.000	250.000	352.000	na
VIO/CMD	29.851	3047.582	3118.081	na
RIO RATE (V)	1.000	247.000	358.000	287.000
RIO/CMD (V)	7.463	3011.011	3171.230	3453.404
DASD IO TOTAL (V)	259103	276772	257162	274719
DASD IO RATE (V)	392.58	271.35	357.17	286.17
DASD IO/CMD (V)	2929.72	3307.79	3163.87	3443.36

Virtual Machine Types	V=R, V=F		V=V, V=V	
DASD Usage	dedicated	shared	dedicated	shared
RELEASE RUN ID	VM/ESA 2 VR2200D6	VM/ESA 2 VR2200S6	VM/ESA 2 VV2200D6	VM/ESA 2 VV2200S6
Environment				
IML MODE	ESA	ESA	ESA	ESA
REAL STORAGE	256MB	256MB	256MB	256MB
EXP. STORAGE	0MB	0MB	0MB	0MB
VM MODE	ESA	ESA	ESA	ESA
VM SIZE	24MB	24MB	24MB	24MB
VSE SUPERVISOR	ESA	ESA	ESA	ESA
PROCESSORS	2	2	2	2
PRIVOPs				
PRIVOP/CMD	35.703	3054.287	3123.813	na
DIAG/CMD	156.329	223.691	217.339	na
SIE/CMD	1067.170	10361.778	11728.236	na
SIE INTCPT/CMD	875.079	9325.600	9265.307	na
FREE TOTL/CMD	1126.872	9130.555	6900.526	na
Note: V=VMPRF, H=Hardware Monitor, Unmarked=RTM. Per command (/CMD), for this workload, means per batch job.				

Table 53. Two VSE Guests, PACEX6 Workload on 9221-200

Appendixes

Appendix A. SFS Counter Data

The SFS counts and timings in this appendix are provided to supplement the information provided for the SFS measurements. These were acquired by issuing the QUERY FILEPOOL STATUS command once at the beginning of the measurement interval and once at the end.

The QUERY FILEPOOL STATUS information was obtained for each SFS file pool server and the CRR recovery server. The counts and timings for each server were added together.

The first section in the following table consists of the counters normalized by the number of commands (as determined by TPNS). The beginning values were subtracted from the ending values and divided by the number of commands in the interval. Counts and timings which have a value of zero for all measurements are not shown. A zero entry indicates that at least one occurrence was counted but the result of normalizing per command is so small that it rounds to zero. A description of the SFS counts and timings can be found in *SFS and CRR Planning, Administration, and Operation*.

The second section consists of derived relationships which were calculated from a combination of two or more individual counts or timings. See the glossary for definitions of these derived values.

RELEASE RUN ID PROCESSOR REAL STORAGE EXP. STORAGE FILE CACHE SIZE WORKLOAD USERS	VM/ESA 1.1 L24M1566 9121-480 192MB 64MB 12KB FS7BMAXR 1560	VM/ESA 2 L25M1564 9121-480 192MB 64MB 20KB FS7BMAXR 1560	VM/ESA 2 L25M1565 9121-480 192MB 64MB 12KB FS7BMAXR 1560
NORMALIZED BY COMMAND			
Close File Requests	0.4030	0.4027	0.3980
Commit Requests	0.0344	0.0344	0.0341
Connect Requests	0.0074	0.0077	0.0078
Delete File Requests	0.0928	0.0936	0.0924
Lock Requests	0.0252	0.0249	0.0250
Open File New Requests	0.0019	0.0018	0.0018
Open File Read Requests	0.2395	0.2382	0.2347
Open File Replace Requests	0.1395	0.1410	0.1393
Open File Write Requests	0.0221	0.0221	0.0220
Query File Pool Requests	0.0000	0.0000	0.0000
Query User Space Requests	0.0198	0.0201	0.0198
Read File Requests	0.2279	0.1671	0.2325
Refresh Directory Requests	0.0217	0.0225	0.0228
Rename Requests	0.0047	0.0047	0.0048
Unlock Requests	0.0250	0.0244	0.0252
Write File Requests	0.1269	0.0645	0.1279
Total File Pool Requests	1.3919	1.2696	1.3881
File Pool Request Service Time	82.7521	54.4733	59.7872
Local File Pool Requests	1.3919	1.2696	1.3881
Begin LUWs	0.5287	0.5143	0.5063
Agent Holding Time (ms)	195.0043	150.1943	152.8759
SAC Calls	6.3716	6.1283	6.1006
Catalog Lock Conflicts	0.0031	0.0024	0.0025
Total Lock Conflicts	0.0031	0.0024	0.0025
Lock Wait Time (ms)	0.4030	0.1466	0.1941
File Blocks Read	0.9485	0.9478	0.9586
File Blocks Written	0.5721	0.5701	0.5728
Catalog Blocks Read	0.4614	0.4633	0.4490
Catalog Blocks Written	0.2409	0.2489	0.2334
Control Minidisk Blocks Read	0.0000	0.0000	0.0002
Control Minidisk Blocks Written	0.0532	0.0510	0.0509
Log Blocks Written	0.5294	0.5395	0.5251
Total DASD Block Transfers	2.8055	2.8206	2.7899
BIO Requests to Read File Block	0.5028	0.4379	0.5041
BIO Requests to Write File Blocks	0.2787	0.2168	0.2792
BIO Requests to Read Catalog Blocks	0.4614	0.4633	0.4490
BIO Requests to Write Catalog Blocks	0.2409	0.1908	0.1769
BIO Requests to Read Cntrl Mdisk Blks	0.0000	0.0000	0.0002
BIO Requests to Write Cntrl Mdisk Blks	0.0022	0.0025	0.0025
BIO Requests to Write Log Blocks	0.5294	0.4684	0.4550
Total BIO Requests	2.0154	1.7796	1.8668
Total BIO Request Time (ms)	32.6312	32.0990	34.0270
I/O Requests to Read File Blocks	0.4674	0.4308	0.4939
I/O Requests to Write File Blocks	0.2960	0.2349	0.2955
I/O Requests to Read Catalog Blocks	0.4614	0.4633	0.4490
I/O Requests to Write Catalog Blocks	0.2409	0.2012	0.1869
I/O Requests to Read Cntrl Mdisk Blks	0.0000	0.0000	0.0002
I/O Requests to Write Cntrl Mdisk Blks	0.0048	0.0046	0.0046
I/O Requests to Write Log Blocks	0.5294	0.4688	0.4554
Total I/O Requests	1.9999	1.8036	1.8854
Get Logname Requests	0.0031	0.0032	0.0033
Get LUWID Requests	0.0031	0.0032	0.0033
Total CRR Requests	0.0062	0.0064	0.0065
CRR Request Service Time (ms)	0.0900	0.1042	0.0945
Log I/O Requests	0.0062	0.0066	0.0067

RELEASE	VM/ESA 1.1	VM/ESA 2	VM/ESA 2
RUN ID	L24M1566	L25M1564	L25M1565
PROCESSOR	9121-480	9121-480	9121-480
REAL STORAGE	192MB	192MB	192MB
EXP. STORAGE	64MB	64MB	64MB
FILE CACHE SIZE	12KB	20KB	12KB
WORKLOAD	FS7BMAXR	FS7BMAXR	FS7BMAXR
USERS	1560	1560	1560
DERIVED RESULTS			
Agents Held	10.8	8.5	8.6
Agents In-call	4.6	3.1	3.4
Avg LUW Time (ms)	368.8	292.0	301.9
Avg File Pool Request Time (ms)	59.5	42.9	43.1
Avg Lock Wait Time (ms)	130.0	61.1	77.6
SAC Calls / FP Request	4.58	4.83	4.39
Deadlocks (delta)	0	0	0
Rollbacks Due to Deadlock (delta)	0	0	0
Rollback Requests (delta)	0	0	0
LUW Rollbacks (delta)	0	0	0
Checkpoints Taken (delta)	37	36	36
Checkpoint Duration (sec)	6.3	3.8	4.1
Seconds Between Checkpoints	49.0	50.3	50.3
Checkpoint Utilization	12.9	7.6	8.1
BIO Request Time (ms)	16.19	18.04	18.23
Blocking Factor (Blocks/BIO)	1.39	1.58	1.49
Chaining Factor (Blocks/IO)	1.40	1.56	1.48

Appendix B. Workloads

CMS-Intensive (FS7B)

Workload Description

FS7B simulates a CMS user environment, with variations simulating a minidisk environment, an SFS environment, or some combination of the two. Table 54 shows the search-order characteristics of the two environments used for measurements discussed in this document.

Filemode	ACCESS	Number of Files	FS7B0R	FS7BMAXR
A	R/W	100	minidisk	SFS
B	R/W	0	minidisk	SFS
C	R/O	500	minidisk	SFS (DS)
D	R/W	500	minidisk	SFS
E	R/O	500	minidisk	SFS (DS)
F	R/O	500	minidisk	SFS (DS)
G	R/O	500	minidisk	SFS (DS)
S	R/O	<i>m</i>	minidisk	minidisk
Y	R/O	<i>n</i>	minidisk	minidisk

Note: *m* and *n* are the number of files normally found on the the S- and Y-disks respectively. (DS) signifies the use of VM Data Spaces.

Table 54. FS7B Workload Characteristics

The measurement environments have the following characteristics in common:

- The HELP disk has the FSTs saved in a shared segment.
- The CMSINST, VMLIB, and VMMLIB shared segments are used.
- A Bactrian-distribution think time averaging 30 seconds is used. (See “Glossary of Performance Terms” on page 245 for an explanation of Bactrian distribution.)
- The workload is continuous in that scripts, repeated as often as required, are always running during the measurement period.
- Teleprocessing Network Simulator (TPNS) simulates users for the workloads. TPNS runs in a separate processor and simulates LU2 terminals. User traffic travels between the processors through channel-attached 3745 communications controllers, 3088 multisystem channel communication units, or both.

FS7B Variations

Two FS7B workload variants were used for measurements, one for minidisk-based CMS users, and the other for SFS-based CMS users.

FS7B0R Workload

All filemodes are accessed as minidisk; SFS is not used. All of the files on the C-disk have their FSTs saved in a shared segment.

FS7BMAXR Workload

All file modes, except S and Y (which SFS does not support), are accessed as SFS directories. The CMSFILES shared segment is used. All read-only SFS directories are defined with PUBLIC READ authority and are mapped to VM data spaces. The read/write SFS directory accessed as file mode D is defined with

PUBLIC READ and PUBLIC WRITE authority. The read/write SFS directories accessed as file modes A and B are private directories.

FS7B Licensed Programs

The following licensed programs were used in the FS7B measurements described in this document:

VS COBOL II Compiler and Library R3.2
Document Composition Facility V1.4
VS FORTRAN Compiler/Library/Debug V2R5
Assembler H V2.1
OS PL/I V2R3 Compiler & Library
C & PL/I Common Library V1R2
VTAM V3R4
NCP V5R4

Measurement Methodology

A calibration is made to determine how many simulated users are required to attain the desired processor utilization for the baseline measurement. That number of users is used for all subsequent measurements on the same processor and for the same environment.

The measurement proceeds as follows:

- All of the users are logged on by TPNS.
- A script is started for each user after a random delay of up to 5 minutes. (The random delay prevents all users from starting at once.)
- A stabilization period (the length depending on the processor used) is allowed to elapse so that start-up anomalies and user synchronization are eliminated.
- At the end of stabilization, measurement tools are started simultaneously to gather data for the measurement interval.
- At the end of the measurement interval, the performance data is reduced and analyzed.

FS7B Script Description

FS7B consists of 3 initialization scripts and 17 workload scripts. The LOGESA script is run at logon to set up the required search order and CMS configuration. Then users run the WAIT script, during which they are inactive and waiting to start the CMSSTRT script. The CMSSTRT script is run to stagger the start of user activity over a five minute interval. After the selected interval, each user starts running a general workload script. The scripts are summarized in Table 55 on page 226.

Script Name	% Used	Script Description
LOGESA	*	Logon and Initialization
WAIT	*	Wait state
CMSSTRT	*	Stagger start of user activity
ASM617A	5	Assemble (HASM) and Run
ASM627A	5	Assemble and Run
XED117	5	Edit a VS BASIC Program
XED127	10	Edit a VS BASIC Program
XED137	10	Edit a COBOL Program
XED147	10	Edit a COBOL Program
COB217	5	COBOL Compile
COB417	5	Run a COBOL Program
FOR217	5	VS FORTRAN Compile
FOR417	5	FORTRAN Run
PRD517	5	Productivity Aids Session
DCF517	5	Edit and Script a File
PLI317	5	PL/I Optimizer Session
PLI717	5	PL/I Optimizer Session
WND517A	8	Run Windows with IPL CMS
WND517AL	2	Run Windows with LOGON/LOGOFF
HLP517	5	Use HELP

Note: Scripts with an asterisk (*) in the “% Used” column are run only once each for each user during initialization.

Table 55. FS7B Workload Script Summary

The following are descriptions of each script used in the FS7B workload.

LOGESA: Initialization Script

```

SET AUTOREAD ON
IF A-disk is a minidisk
THEN
    Check the format of the A-disk
    Check the format of the B-disk
    Erase extraneous files from A-disk
    ACCESS 191 as A-disk
    Run PROFILE EXEC to access correct search order,
    SET ACNT OFF, SPOOL PRT CL D, and TERM LINEND OFF
ELSE
    Erase extraneous files from A-directory
    Run PROFILE EXEC to set correct search order, SET ACNT OFF,
    SPOOL PRT CL D, and TERM LINEND OFF
END
SET REMOTE ON

```

WAIT: Ten-Second Pause

Leave the user inactive in a 10-second wait loop.

CMSSTRT: Random-Length Pause

Delay, for up to 5 minutes, the start for each user to prevent all users from starting scripts at the same time.

ASM617A: Assemble (HASM) and Run

QUERY reader and printer
SPOOL PRT CLASS D
XEDIT an assembler file and QQUIT
GLOBAL appropriate MACLIBs
LISTFILE the assembler file
Assemble the file using HASM (NOLIST option)
Erase the text deck
Repeat all the above except for XEDIT
Reset GLOBAL MACLIBs
Load the text file (NOMAP option)
Generate a module (ALL and NOMAP options)
Run the module
Load the text file (NOMAP option)
Run the module 2 more times
Erase extraneous files from A-disk

ASM627A: Assemble (F-Assembler) and Run

QUERY reader and printer
SPOOL PRT CLASS D
GLOBAL appropriate MACLIBs
LISTFILE assembler file
XEDIT assembler file and QQUIT
Assemble the file (NOLIST option)
Erase the text deck
Reset GLOBAL MACLIBs
Load the TEXT file (NOMAP option)
Generate a module (ALL and NOMAP options)
Run the module
Load the text file (NOMAP option)
Run the module
Load the text file (NOMAP option)
Run the module
Erase extraneous files from A-disk
QUERY DISK, USERS, and TIME

XED117: Edit a VS BASIC Program

XEDIT the program
Get into input mode
Enter 29 input lines
Quit without saving file (QQUIT)

XED127: Edit a VS BASIC Program

Do a FILELIST
XEDIT the program
Issue a GET command
Issue a LOCATE command
Change 6 lines on the screen
Issue a TOP and BOTTOM command
Quit without saving file
Quit FILELIST
Repeat all of the above statements, changing 9 lines instead of 6 and
without issuing the TOP and BOTTOM commands

XED137: Edit a COBOL Program

Do a FILELIST
XEDIT the program
Issue a mixture of 26 XEDIT file manipulation commands
Quit without saving file
Quit FILELIST

XED147: Edit a COBOL Program

Do a FILELIST
XEDIT the program
Issue a mixture of 3 XEDIT file manipulation commands
Enter 19 XEDIT input lines
Quit without saving file
Quit FILELIST

COB217: Compile a COBOL Program

Set ready message short
LINK and ACCESS a disk
QUERY link and disk
LISTFILE the COBOL program
Invoke the COBOL compiler
Erase the compiler output
RELEASE and DETACH the linked disk
Set ready message long
SET MSG OFF
QUERY SET
SET MSG ON
Set ready message short
LINK and ACCESS a disk
LISTFILE the COBOL program
Run the COBOL compiler
Erase the compiler output
RELEASE and DETACH the linked disk
QUERY TERM and RDYMSG
Set ready message long
SET MSG OFF
QUERY set
SET MSG ON
PURGE printer

COB417: Run a COBOL Program

Define temporary disk space for 2 disks using an EXEC
QUERY DASD and format both temporary disks
Establish 4 FILEDEFS for input and output files
QUERY FILEDEFS
GLOBAL TXTLIB
Load the program
Set PER Instruction
Start the program
Display registers
End PER
Issue the BEGIN command
QUERY search of minidisks
RELEASE the temporary disks
Define one temporary disk as another
DETACH the temporary disks
Reset the GLOBALs and clear the FILEDEFS

FOR217: Compile 6 VS FORTRAN Programs

NUCXDROP NAMEFIND using an EXEC
 QUERY and PURGE the reader
 Compile a FORTRAN program
 Issue INDICATE commands
 Compile another FORTRAN program
 Issue INDICATE commands
 Compile another FORTRAN program
 Issue INDICATE commands
 Repeat the above 6 statements
 Erase extraneous files from A-disk
 PURGE the printer

FOR417: Run 2 FORTRAN Programs

SPOOL PRT CLASS D
 GLOBAL appropriate text libraries
 Issue 2 FILEDEFS for output
 Load and start a program
 Rename output file and PURGE printer
 Repeat above 5 statements for two other programs, except
 erase the output file for one and don't issue spool printer
 List and erase output files
 Reset GLOBALs and clear FILEDEFS

PRD517: Productivity Aids Session

Run an EXEC to set up names file for user
 Issue NAMES command and add operator
 Locate a user in names file and quit
 Issue the SENDFILE command
 Send a file to yourself
 Issue the SENDFILE command
 Send a file to yourself
 Issue the SENDFILE command
 Send a file to yourself
 Issue RDRLIST command, PEEK and DISCARD a file
 Refresh RDRLIST screen, RECEIVE an EXEC on B-disk, and quit
 TRANSFER all reader files to punch
 PURGE reader and punch
 Run a REXX EXEC that generates 175 random numbers
 Run a REXX EXEC that reads multiple files of various sizes from
 both the A-disk and C-disk
 Erase EXEC off B-disk
 Erase extraneous files from A-disk

DCF517: Edit and SCRIPT a File

XEDIT a SCRIPT file
Input 25 lines
File the results
Invoke SCRIPT processor to the terminal
Erase SCRIPT file from A-disk

PLI317: Edit and Compile a PL/I Optimizer Program

Do a GLOBAL TXTLIB
Perform a FILELIST
XEDIT the PL/I program
Run 15 XEDIT subcommands
File the results on A-disk with a new name
Quit FILELIST
Enter 2 FILEDEFS for compile
Compile PL/I program using PLIOPT
Erase the PL/I program
Reset the GLOBALS and clear the FILEDEFS
QUERY virtual devices
TELL yourself one pass of script run

PLI717: Edit, Compile, and Run a PL/I Optimizer Program

Copy and rename the PL/I program and data file from C-disk
XEDIT data file and QQUIT
XEDIT a PL/I file
Issue RIGHT 20, LEFT 20, and SET VERIFY ON
Change two lines
Change filename and file the result
Compile PL/I program using PLIOPT
Set two FILEDEFS and QUERY the settings
Issue GLOBAL for PL/I transient library
Load the PL/I program (NOMAP option)
Start the program
Type 8 lines of one data file
Erase extraneous files from A-disk
Erase extra files on B-disk
Reset the GLOBALS and clear the FILEDEFS
TELL yourself one pass of script run

WND517A: Use Windows

SET FULLSCREEN ON
TELL yourself a message to create window
QUERY DASD and reader
Forward 1 screen
TELL yourself a message to create window
Drop window message
Scroll to top and clear window
Backward 1 screen
Issue a HELP WINDOW and choose Change Window Size
QUERY WINDOW
Quit HELP WINDOWS
Change size of window message
Forward 1 screen
Display window message
TELL yourself a message to create window
Issue forward and backward border commands in window message
Position window message to another location
Drop window message
Scroll to top and clear window
Display window message
Erase MESSAGE LOGFILE
IPL CMS
SET AUTOREAD ON
SET REMOTE ON

WND517AL: Use Windows with LOGON, LOGOFF

SET FULLSCREEN ON
TELL yourself a message to create window
QUERY DASD and reader
Forward 1 screen
TELL yourself a message to create window
Drop window message
Scroll to top and clear window
Backward 1 screen
Issue a help window and choose Change Window Size
QUERY WINDOW
Quit help windows
Change size of window message
Forward 1 screen
Display window message
TELL yourself a message to create window
Issue forward and backward border commands in window message
Position window message to another location
Drop window message
Scroll to top and clear window
Display window message
Erase MESSAGE LOGFILE
LOGOFF user and wait 60 seconds
LOGON user on original GRAF-ID
SET AUTOREAD ON
SET REMOTE ON

HLP517: Use HELP and Miscellaneous Commands

Issue HELP command
Choose HELP CMS
Issue HELP HELP
Get full description and forward 1 screen
Quit HELP HELP
Choose CMSQUERY menu
Choose QUERY menu
Choose AUTOSAVE command
Go forward and backward 1 screen
Quit all the layers of HELP
RELEASE Z-disk
Compare file on A-disk to C-disk 4 times
Send a file to yourself
Change reader copies to two
Issue RDRLIST command
RECEIVE file on B-disk and quit RDRLIST
Erase extra files on B-disk
Erase extraneous files from A-disk

IBM Office Benchmark (IOB V2.1)

Workload Description

The IBM Office Benchmark (IOB) Version 2.1 is a corporate-wide benchmark designed to measure generic office system performance. It consists of a definition of the office user; databases for calendars, documents, and mail; and the work the office users do. This workload was developed in Dallas.

Unless otherwise noted in the measurement tables, the IOB measurements included in this report use the DisplayWrite/370 2.1.0 and the OfficeVision/VM 1.2.0 GA level licensed programs. For the 9121-480 IOB measurements that list "OV/VM 1.1.0" as the OfficeVision release, the DisplayWrite/370 2.1.0 and the OfficeVision/VM 1.1.0 Service Level 101 licensed programs were used.

Measurement Methodology

The general methodology was to logon as many users as possible until the processor utilization reached the desired level.

The IOB workload does not aim for a specific think time or use a certain think time distribution. Instead, the think time is dictated by the IOB workload. The think time includes an average two second delay between commands issued by TPNS, the built-in think times that are part of the IOB scripts, and the IOB script scheduling algorithm. When users finish running a script, the script scheduling algorithm calculates how much time was spent running the script, subtracts this number from ten minutes, and delays the user for the resulting amount of time. Thus, if a script completed in 7.9 minutes, the user would be delayed for 2.1 minutes before starting the next script and this time would be included in the user's think time.

Getting a valid measurement takes several steps as follows:

- All of the users are logged on via TPNS and reach the OfficeVision main menu (the A00 screen).
- After a random delay of up to 10 minutes, each user selects a script and starts. The random delay prevents all users from starting at once.
- A stabilization period (typically 35 minutes) is allowed to elapse so that startup anomalies and user synchronization are eliminated.
- RTM/ESA and the internal XXTRANS tool are used to ensure that the system has properly stabilized.
- At the end of stabilization, measurement tools are started simultaneously to gather data for the 30-minute measurement interval.
- At the end of the measurement interval, the measurement tools are stopped. The resulting performance data are then reduced and analyzed.

After the run data are analyzed and qualified for IOB certification, the run data are sent to Dallas for certification. All of the runs in this report were certified as valid IOB runs.

IOB Script Descriptions

The IOB workload consists of nine scripts (scenarios). These scripts are listed in Table 56 with their defined usage factor.

Script Name	% Used	Script Description
VMB2LML	17	Send Note and Process Light Mail
VMB2HML	17	Send Note and Process Heavy Mail
VMB2VCAL	13	View Individual Calendar
VMB2UCAL	13	Update Individual Calendar
VMB2DIR	20	View User Directory
VMB2CDOC	7	Create Small Text Document
VMB2UDOC	7	Revise Small Text Document
VMB2EB	3	End/Begin Office
VMB2ONOF	3	Logoff/Logon System

Table 56. IOB Workload Script Summary

The following is the list of tasks in each script within the IOB workload.

Send Note and Process Light Mail

- Create a note and send the note to two users (self and Ennnn).
- View the note log.
- View the first item, a note.
- Delete the first item, a note.
- Open Mail and View the In-Basket (old and new mail).
- View the first item, a note.
- Delete the first item, a note.

Send Note and Process Heavy Mail

- Create a note and send the note to two users (self and Ennnn).
- View the note log.
- View the first item, a note.
- Delete the first item, a note.
- Open Mail and View the In-Basket (old and new mail).
- View the first item, a note.
- Forward the first item to another user (Ennnn) with an attachment.
- Delete the original first item, a note.
- View the eighth item in the mail list, a two page document.
- Print the document (to PRNTEAT1).

View Individual Calendar

- View the user's calendar for Wednesday of a defined week.

Update Individual Calendar

- View the user's calendar for Wednesday of a defined week.
- Delete a meeting.
- Add a meeting.

View User Directory

- Search the user directory for the user's name and view the person's telephone number.

Create Small Text Document

- Get a pre-stored document format.
- Key in a two-page document.
- Save the document.
- Print the document (to PRNTEAT1).
- Delete the document.

Revise Small Text Document

- Open a two-page document for revision.
- Move one paragraph.
- Delete one paragraph.
- Insert one paragraph.
- Save the altered document.
- Send the document to three users (all are Ennnn).

End/Begin Office

- End or exit the office software program or environment.
- Begin or enter the office software program or environment.

Logon/Logoff System

- Take the option to log off completely from the system.
- Log back onto the system and enter the office environment.

VSE Guest (PACE)

PACE is a synthetic VSE batch workload consisting of 7 unique jobs representing the commercial environment. This set of jobs is replicated c times, producing the PACE Xc workload, in order to scale the load to the environment being tested. The most commonly-used multiples are PACE $X4$, PACE $X6$, and PACE $X8$.

The seven jobs are as follows:

- Y n DL/1
- Y n SORT
- Y n COBOL
- Y n BILL
- Y n STOCK
- Y n PAY
- Y n FORT

There are c copies of these jobs used in the PACE Xc workload. They are differentiated by the n digit in the name (n having a value from 1 to c).

The programs, data, and work space for the jobs are all maintained by VSAM on separate volumes.

The VSE system is configured with the full complement of 12 static partitions (BG, and F1 through FB). F4 through FB are the partitions used to run the workload batch jobs for PACE $X8$; a subset of these partitions are used for smaller PACE variations.

The partitions are configured identically except for the job classes. The jobs and the partition job classes are configured so that the jobs are equally distributed over the partitions and so that, at any one time, the jobs currently running are a mixed representation of the 7 jobs.

When a workload is ready to run, the following preparatory steps are taken:

- CICS/ICCF is shut down
- VTAM is shut down
- The LST queue is emptied (PDELETE LST,ALL)

Once performance data gathering is initiated for the system (hardware instrumentation, CP MONITOR, RTM), the workload is started by releasing all of the batch jobs into the partitions simultaneously using the POWER command, PRELEASE RDR,*Y. The start time is noted.

As the workload nears completion, various partitions will finish the work allotted to them. The finish time for both the first and last partitions is noted. The difference between these two times should not be more than about 10% of the total elapsed time. If it is more, the jobs and partitions have to be adjusted to get a more even work distribution.

At workload completion, the ITR can be calculated by dividing the number of batch jobs by processor busy time. The processor busy time is calculated as elapsed (wall clock) time multiplied by processor busy percent divided by 100.

VSE/ESA 1.2.0 or VSE/ESA 1.2.2 were used in the measurements to gather data for the VSE guest measurements in this report.

VSE Guest (VSECICS)

Workload Description

The VSECICS workload consists of seven applications, written in COBOL and assembler, which include order entry, receiving and stock control, inventory tracking, production specification, banking, and hotel reservations. These applications invoke a total of 17 transactions averaging approximately 6 VSAM calls and 2 communication calls per transaction.

Two independent CICS partitions are run to effectively utilize the measured processor. The storage configuration for this workload is 128MB primary storage and no expanded storage. Each of the two CICS/VSE* partitions accesses 16 VSAM KSDS files. Measurements are taken at the 70% and the 90% processor utilization points.

CICS is measured by logging on a predefined number of users, each of which starts running commands from 1 of 12 possible scripts. Once the system reaches a steady state condition, the think time is adjusted to provide a transaction rate that will cause the processor to reach the target utilization level (70% or 90%). CICS is measured as a steady state system, over a period deemed to be a repeatable sample of work.

Software products used by the CICS workload include VSE/ESA 1.2.0, CICS/VSE 2.1.0, ACF/VTAM 3.3.0, and VSE/VSAM 2.1.1. POWER and ACF/VTAM run in their own individual address spaces. This allows, among other things, virtual storage constraint relief. Access methods used include the Sequential Access Method (SAM) and the Virtual Storage Access Method (VSAM). CMF data is logged and then processed by the CICSPARS post-processing facility. Internal response time and total transaction counts are gathered from the CICSPARS report. Legent's EXPLORE** is used to gather system performance data including processor utilization, CHPID utilization, and DASD utilization.

The workload produces an 80% read and 20% write mixture. Each application uses several transactions that employ differing sets of CICS functions. Table 57 on page 240 indicates the number of transactions for each application and the frequency of specific CICS functions within each.

Transaction Type	VSAM Calls	Read	Read Next	Add	Update	Delete	Trans. Data	Temp. Stor.	% Mix
Banking	3 10	2 8		2	1	1 2			8 8
Hotel Reservations	2 2	1 1		1	1				3 3
Inventory Control	0 17 4	1 5	16 14	3	2			1	3 6 8
Order Entry	3 3 9 22	1 1 9 9	2	1 4	1 9		1 2 2 1		5 5 5 5
Product Specification	18 34	8 2	10 32						10 9
Stock Control	18 9 3	1 1	8 1	9	9	1			5 3 10
Teller System	0								4

Table 57. CICS/VSE Transaction Characteristics

Measurement Methodology

Twenty DASD volumes (including DOSRES and SYSWK1) are required to run this workload. Each CICS (CICST1 and CICST2) has its own set of 8 dedicated volumes for VSAM data files. There should be two CHPIDS going to each string of data volumes. The VSAM data file strings should be attached to separate 3880 or 3990 noncached control units. Any 3380 model DASD may be used for the VSAM data files; however, all volumes must be the same model.

At every measurement point, a CICSPARS report is generated for each of the two CICS workload systems. To determine the total transaction count, which is used to calculate the ITR, add the TOTAL TASKS SELECTED fields from both CICSPARS reports.

When a native measurement is performed, the %CPU Busy used to calculate the ITR is selected from the CPU %Active field of the EXPLORE report. However, if a VSE guest measurement is performed, the %CPU Busy is selected from the System Summary by Time VMPRF report.

The ITR is calculated as

$$\frac{\text{transactions}}{\text{processor busy seconds}}$$

MVS Guest

Workload Description

The MVS batch workload is designed to represent a “typical” commercial batch workload. It is run under RACF control with the data base 100% managed by DFSMS*. Job submission is controlled by the IBM job scheduling package OPC/A. A printer simulator is used to simulate output processing for output in the print class.

The job stream consists of user applications using a variety of licensed programs and DB2 batch jobs. Each copy of the batch workload uses 138 permanent data sets and consists of 32 jobs containing a total of 132 steps.

One copy of this workload contains the following 32 jobs (including 5 in-stream procedures):

- 8 PL/I Go jobs
- 4 COBOL Compile, Link, and Go jobs
- 4 DB/2 batch application jobs
- 4 DB/2 utility jobs
- 3 SORT jobs
- 2 ADRDSSU jobs
- 2 SLR jobs
- 1 COBOL Compile job
- 1 C Compile, Link, and Go job
- 1 FORTRAN compile job
- 1 FORTRAN Go job
- 1 LKED job

Measurement Methodology

To obtain the data in this report, the batch workload was run on a 9021-520 in three distinct ways:

1. Native, on a machine with 128MB of central storage, 256MB of expanded storage, and one processor.
2. As a V=R guest without a dedicated processor. The guest had 128MB of central storage and 256MB of expanded storage, but it shared a single real processor with VM.
3. As a V=R guest with a dedicated processor. Again, the guest had 128MB of central storage and 256MB of expanded storage.

After each MVS system IPL and before any workload measurements were performed, a priming run was done to load modules into the virtual lookaside facility (VLF). All runs used five initiators and three simulated printers.

The Resource Measurement Facility (RMF) was run to collect MVS-specific performance information. For the V=R runs, the VM/ESA Monitor and RTM/ESA were run on VM to collect system performance information. All performance monitors were stopped when the last job finished running. The ITR was calculated as

$$\frac{32}{\text{Elapsed seconds} \times \text{CPU \% busy} \times 0.01}$$

The "CPU % busy" in this formula was calculated in three different ways corresponding to the three distinct ways the workload was run:

1. For the native runs, this was simply the "CPU BUSY PERCENTAGE" from the RMF CPU Activity report.
2. For the V=R non-dedicated runs, this was the processor busy percentage as reported by VMPRF in the PROCESSORS_BY_TIME report.
3. For the V=R dedicated runs, this was the RMF CPU BUSY PERCENTAGE plus the average VMPRF processor busy percentage for the *other* processor. To VM, a dedicated processor always appears busy, so RMF's estimate of processor activity for the dedicated processor and VMPRF's estimate of processor activity for the other processor were used.

The ETR is simply

$$\frac{32}{\text{Elapsed seconds}}$$

Appendix C. Configuration Details

Named Saved Segments / Systems

CMS allows the use of saved segments for shared code. Using saved segments can greatly improve performance by reducing end users' working set sizes and thereby decreasing paging. The environments in this report used the following saved segments:

CMS	Contains the CMS nucleus and File Status Tables (FSTs) for the S and Y disks.
CMSFILES	Contains the SFS server code modules DMSDAC and DMSSAC
CMSINST	Contains the EXECs-in-storage segment
VMLIB	Contains the CSL code
VMMTLIB	Contains the CMS multitasking code (VM/ESA Release 2)
HELP	Contains FSTs for the HELP disk
GOODSEG	Contains FSTs for the C disk
FORTRAN	This segment space has 2 members: DSSVFORT for the FORTRAN compiler and FTNLIB20 for the Library composite modules.
DSMSEG4B	Contains DCF (Document Composition Facility)
OFSSEG	Contains OV/VM user functions
EPUYSSEG	Contains OV/VM mailbox manager code
DW370210	Contains the DW370 module
DDDCL210	Contains the DW370 compiled CLISTS
DW362	Contains FSTs for the DW/370 362 disk
ADM399	Contains FSTs for the OV/VM 399 disk
GCSXA	Contains the GCS nucleus
VTAMXA	Contains the VTAM code

Server Options

SFS DMSPARMS

This section lists the start-up parameter settings used by each of the SFS servers. The start-up parameters determine the operational characteristics of the file pool server. The SFS servers used the following DMSPARMS file:

```
ADMIN OPERATOR MAINT U3 MARK
FILEPOOLID fp_name
NOBACKUP
FORMAT
USERS 780
FULLDUMP
SAVESEGID CMSFILES
MSGS
ACCOUNT
```

For all SFS measurements, the SAVESEGID is specified to identify the segment containing the file pool server runnable code. The USERS parameter is used by the SFS server to configure itself with the appropriate number of user agents and buffers. It is recommended that USERS be set to the administrator's best estimate of the maximum number of logged-on virtual machines that will be using the file pool during peak usage. The ratio of logged-on users to active users varies greatly on actual production machines.

For more information on SFS and SFS tuning parameters, see the *SFS and CRR Planning, Administration, and Operation* manual.

CRR DMSPARMS

This section lists the start-up parameter settings used by the CRR recovery server. The start-up parameters determine the operational characteristics of the CRR recovery server. The CRR server uses the following DMSPARMS file:

```
ADMIN OPERATOR MAINT U3 MARK
NOBACKUP
FULLDUMP
SAVESEGID CMSFILES
FILEPOOLID CRRSERV1
ACCOUNT
NOFORMAT
CRR
LUNAME CRRSERV1.PERFTEST
USERS 10
MSGS
```

For more information on CRR and CRR tuning parameters, see the *SFS and CRR Planning, Administration, and Operation* manual.

Glossary of Performance Terms

Many of the performance terms use postscripts to reflect the sources of the data described in this document. In all cases, the terms presented here are taken directly as written in the text to allow them to be found quickly. Often there will be multiple definitions of the same data field, differing only in the postscript. This allows the precise definition of each data field in terms of its origins. The postscripts are:

<none>. No postscript indicates that the data are obtained from the VM Realtime Monitor for VM/XA or VM/ESA. For VM/SP and VM/ESA Release 1.0 (370 Feature), it denotes data from the VM Monitor Analysis Program.

(C). Denotes data obtained from the VSE CICS Transaction File.

(H). Denotes data obtained from the internal processor instrumentation tools.

(Q). Denotes data from the SFS QUERY FILEPOOL STATUS command.

Server. Indicates that the data are for specific virtual machines, (for example SFS, CRR, or VTAM/VSCS). If there is more than one virtual machine of the same type, these data fields are for all the virtual machines of that type.

(T). Identifies data from the licensed program, Teleprocessing Network Simulator (TPNS).

(V). Denotes data from the licensed program VM Performance Reporting Facility for VM/XA or VM/ESA. For VM/SP and VM/ESA Release 1.0 (370 Feature), it denotes data from the VM Monitor Analysis Program.

(H,V). Denotes data obtained from the internal processor instrumentation tools for VM/SP and VM/ESA Release 1.0 (370 Feature). For VM/XA and VM/ESA it denotes data obtained from the licensed program, VM Performance Reporting Facility.

The formulas used to derive the various statistics are also shown here. If a term in a formula is in italics, such as *Total_Transmits*, then a description of how its value is derived is provided underneath the formula. If a term is not in italics, such as SFSTIME, then it has an entry in the glossary describing its derivation.

Agent. The unit of sub-dispatching within a CRR or SFS file pool server.

Agents Held. The average number of agents that are in a Logical Unit of Work (LUW). This is calculated by:

$$\frac{1}{1000} \times \sum_{f \in \text{filepools}} \frac{\textit{Agent_Holding_Time}_f}{\textit{SFSTIME}_f}$$

Agent_Holding_Time is from the QUERY FILEPOOL STATUS command.

Agents In Call. The average number of agents that are currently processing SFS server requests. This is calculated by:

$$\frac{1}{1000} \times \sum_{f \in \text{filepools}} \frac{\textit{Filepool_Request_Service_Time}_f}{\textit{SFSTIME}_f}$$

Filepool_Request_Service_Time is from the QUERY FILEPOOL STATUS command.

AVG ELIST SIZE (V). The average number of virtual machines in the eligible list queues (E0, E1, E2, E3) calculated by VMPRF.

This is taken from the average value of the <Elist> field in the VMPRF System Summary By Time report.

Avg Filepool Request Time (ms). The average time it takes for a request to the SFS file pool server machine to complete. This is calculated by:

$$\sum_{f \in \text{filepools}} \frac{\textit{Agents In Call}}{\frac{\textit{Total_Filepool_Requests}_f}{\textit{SFSTIME}_f}}$$

Total_Filepool_Requests is from the QUERY FILEPOOL STATUS command.

AVG FIRST (T). The average response time in seconds for the first reply that returns to the screen. For non-fullscreen commands this is the command reflect on the screen. This is calculated by:

$$\frac{1}{\textit{ETR (T)}} \times \sum_{t \in \text{TPNS machines}} \frac{\textit{First_Response}_t \times \textit{Total_Transmits}_t}{\textit{TPNS_Time}_t}$$

First_Response is the average first response given in the RSPRPT section of the TPNS reports. *Total_Transmits* is the total TPNS transmits and *TPNS_Time* is the run interval log time found in the Summary of Elapsed Time and Times Executed section of the TPNS reports.

AVG LAST (T). The average response time in seconds for the last response to the screen. If there is more than one TPNS this is calculated by:

$$\frac{1}{\textit{ETR (T)}} \times \sum_{t \in \text{TPNS machines}} \frac{\textit{Last_Response}_t \times \textit{Total_Transmits}_t}{\textit{TPNS_Time}_t}$$

Last_Response is the average last response given in the RSPRPT section of the TPNS reports. *Total_Transmits* is the total TPNS transmits and *TPNS_Time* is the run interval log time found in the Summary of Elapsed Time and Times Executed section of the TPNS reports.

AVG Lock Wait Time (ms). The average time it takes for an SFS lock conflict to be resolved. This is calculated by:

$$\frac{\sum_{f \in \text{filepools}} \frac{\text{Lock_Wait_Time}_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{\text{Total_Lock_Conflicts}_f}{\text{SFSTIME}_f}}$$

Lock_Wait_Time and *Total_Lock_Conflicts* are both from the QUERY FILEPOOL STATUS command.

AVG LUW Time (ms). The average duration of an SFS logical unit of work. This is calculated by:

$$\frac{\sum_{f \in \text{filepools}} \frac{\text{Agent_Holding_Time}_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{\text{Begin_LUWs}_f}{\text{SFSTIME}_f}}$$

Agent_Holding_Time and *Begin_LUWs* are both from the QUERY FILEPOOL STATUS command.

AVG RESP (C). The average response time in seconds for a VSE CICS transaction. This is calculated by:

$$\sum_{t \in \text{CICSPARfiles}} \frac{\text{Last_Response}_t \times \text{Total_Transmits}_t}{\text{Total_Transmits}_t}$$

Last_Response is taken from the AVG TASK RESPONSE TIME line and *Total_Transmits* is from the TOTAL TASKS SELECTED line the CICSPAR reports. *CICS_Time* is the run interval time, which is 900 seconds for all measurements.

AVG THINK (T). Average think time in seconds. The average think time determined by TPNS for all users. This is calculated by :

$$\frac{1}{\text{ETR (T)}} \times \sum_{t \in \text{TPNS machines}} \frac{\text{Think_Time}_t \times \text{Total_Transmits}_t}{\text{TPNS_Time}_t}$$

Think_Time is the average think time given in the RSPRPT section of the TPNS reports. *Total_Transmits* is the total TPNS transmits and *TPNS_Time* is the run interval log time found in the Summary of Elapsed Time and Times Executed section of the TPNS reports.

Bactrian. A two-humped curve used to represent the think times for both active users and users who are logged on but inactive. The distribution includes those long think times that occur when a user is not actively issuing commands. Actual user data were collected and used as input to the creation of the Bactrian distribution.

BIO Request Time (ms). Average time required to process a block I/O request in milliseconds. This is calculated by:

$$\frac{\sum_{f \in \text{filepools}} \frac{\text{Total_BIO_Request_Time}_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{\text{Total_BIO_Requests}_f}{\text{SFSTIME}_f}}$$

Total_BIO_Request_Time and *Total_BIO_Requests* are both from the QUERY FILEPOOL STATUS command.

Blocking Factor (Blocks/BIO). The average number of blocks read or written per Block I/O Request. This is calculated by:

$$\frac{\sum_{f \in \text{filepools}} \frac{\text{Total_DASD_Block_Transfers}_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{\text{Total_BIO_Requests}_f}{\text{SFSTIME}_f}}$$

Total_DASD_Block_Transfers and *Total_BIO_Requests* are both from the QUERY FILEPOOL STATUS command.

Chaining Factor (Blocks/IO). The average number of blocks read or written per I/O request. This is calculated by:

$$\frac{\sum_{f \in \text{filepools}} \frac{\text{Total_DASD_Block_Transfers}_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{\text{Total_IO_Requests}_f}{\text{SFSTIME}_f}}$$

Total_DASD_Block_Transfers and *Total_IO_Requests* are both from the QUERY FILEPOOL STATUS command.

Checkpoint. 1) In an SFS file pool server, the periodic processing that records a consistent state of the file pool on DASD. 2) In a CRR recovery server, the process used to maintain the log disks. All active syncpoint information is written to the logs.

Checkpoint Duration. The average time, in seconds, required to process an SFS checkpoint. This is calculated by:

$$\frac{1}{1000} \times \frac{\sum_{f \in \text{filepools}} \text{Checkpoint_Time}_f}{\sum_{f \in \text{filepools}} \text{Checkpoints_Taken}_f}$$

Checkpoint_Time and *Checkpoints_Taken* are from the QUERY FILEPOOL STATUS command.

Checkpoint Utilization. The percentage of time an SFS file pool server spends performing checkpoints. This is calculated by:

$$\frac{1}{10} \times \sum_{f \in \text{filepools}} \frac{\text{Checkpoint_Time}_f}{\text{SFSTIME}_f}$$

Checkpoint_Time is from the QUERY FILEPOOL STATUS command.

Checkpoints Taken (delta). The number of checkpoints taken by all file pools on the system. This is calculated by:

$$\sum_{f \in \text{filepools}} \text{Checkpoints_Taken}_f$$

Checkpoints_Taken is from the QUERY FILEPOOL STATUS command.

CMS BLOCKSIZE. The block size, in bytes, of the users' CMS minidisks.

Command. In the context of reporting performance results, any user interaction with the system being measured.

CP/CMD. For the VSECICS, FS7B, and IOB workloads this is the average amount of CP processor time used per command in milliseconds. For the PACE workload this is the average CP processor time per job in seconds. This is calculated by:

For VM/XA and VM/ESA:

For VSECICS, FS7B, and IOB workloads:

$$10 \times \frac{(\text{TOTAL} - \text{TOTAL EMUL})}{\text{ETR (T) or ETR (C)}}$$

For PACE workload:

$$\text{PBT/CMD} - \text{EMUL/CMD}$$

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

For FS7B and IOB workloads:

$$10 \times \frac{\text{CP_CPU}}{\text{ETR (T)}}$$

For PACE workload:

$$6000 \times \frac{\text{CP_CPU}}{\text{ETR (H)}}$$

CP_CPU comes from the CPCPU field in the VMMAP Statistical Summary report.

CP/CMD (H). See CP/CMD. This is the hardware based measure. This is calculated by:

For 9221 Processors:

For VSECICS, FS7B, and IOB workloads:

$$\frac{\text{CP_CPU_PCT} \times \text{TOTAL (H)}}{10 \times \text{ETR (T) or ETR (C)}}$$

For PACE workload:

$$6000 \times \frac{\text{CP_CPU_PCT} \times \text{TOTAL (H)}}{\text{ETR (H)}}$$

CP_CPU_PCT is taken from the SUPV CPU Busy (VM/SP and VM/ESA Release 1.0 (370 Feature)) or the Host CPU Busy line (VM/XA or VM/ESA) in the CPU Busy/MIPs section of the RE0 report.

For all workloads running on 9121 and 9021 Processors:

$$\text{PBT/CMD (H)} - \text{EMUL/CMD (H)}$$

CP CPU/CMD (V) Server. CP processor time, in milliseconds, run in the designated server machine per command. This is calculated by:

$$\left(\frac{1}{V_Time \times \text{ETR (T)}} \right) \times \sum_{s \in \text{server class}} (\text{TCPU}_s - \text{VCPUs}_s)$$

For VM/XA or VM/ESA:

TCPU is Total CPU busy seconds, *VCPU* is Virtual CPU seconds, and *V_Time* is the VMPRF time interval obtained from the Resource Utilization by User Class section of the VMPRF report.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

TCPU is Total CPU busy seconds, *VCPU* is Virtual CPU seconds, both from the VMMAP User Resource Utilization report for each server user ID. *V_Time* is from the TOTAL SECONDS ANALYZED field at the beginning of the VMMAP report.

CPU Time (estimated). Estimated processor times in milliseconds, based on opcode usage, provided by the STARS tool. These are not real processor times, but estimates used only in comparisons of similar traces.

CPU/*RPI (V). Processor time used per successful data transfer to the *RPI system service. This is calculated by:

$$\frac{1}{*\text{RPI}/\text{CPU sec (V)}}$$

DASD IO/CMD (V). The number of real SSCH or RSCH instructions issued to DASD, per job, used by the VSE guest in a PACE measurement. This is calculated by:

$$60 \times \frac{\text{DASD IO RATE (V)}}{\text{ETR (H)}}$$

DASD IO RATE (V). The number of real SSCH or RSCH instructions issued to DASD, per second, used by the VSE guest in a PACE measurement. This is calculated by:

$$\frac{\text{DASD IO TOTAL (V)}}{V_Time}$$

For VM/XA or VM/ESA:

V_Time is taken from the time stamps at the beginning of the VMPRF DASD Activity Ordered by Activity report.

For VM/SP or VM/ESA Release 1.0 (370 Feature):

V_Time is taken from the TOTAL SECONDS ANALYZED line at the beginning of the VMMAP report.

DASD IO TOTAL (V). The number of real SSCH or RSCH instructions issued to DASD used by the VSE guest in a PACE measurement. This is calculated by:

$$\sum_{d \in \text{VSE Guest DASD}} \text{Total}_d$$

For VM/XA or VM/ESA:

Total is taken from the Count column in the VMPRF DASD Activity Ordered by Activity report for the individual DASD volumes used by the VSE guest.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

Total is taken from the TOTAL I/O ISSUED column in the VMMAP Disk and Tape I/O Activity report for the individual DASD volumes used by the VSE guest.

Deadlocks (delta). The total number of SFS file pool deadlocks that occurred during the measurement interval summed over all production file pools. A deadlock occurs when two users each request a resource that the other currently owns. This is calculated by:

$$\sum_{f \in \text{filepools}} \text{Deadlocks}_f$$

Deadlocks is from the QUERY FILEPOOL STATUS command.

DIAGNOSE. An instruction that is used to request CP services by a virtual machine. This instruction causes a SIE interception and returns control to CP.

DIAG 08/CMD. The number of DIAGNOSE code Xc08c instructions used per command. DIAGNOSE code Xc08c is the CP function call to issue CP commands from an application. This is calculated by:

$$\frac{\text{DIAG_08_SEC}}{\text{ETR (T)}}$$

DIAG_08_SEC is taken from the NSEC column for the total smart interval time on the RTM PRIVOPS Screen.

Note: Because of the way this number is calculated, it can be inaccurate, particularly for small values.

DIAG 10/CMD. The number of DIAGNOSE code Xc10c instructions used per command. DIAGNOSE code Xc10c is the CP function call to release pages of virtual storage. This is calculated by:

$$\frac{\text{DIAG_10_SEC}}{\text{ETR (T)}}$$

DIAG_10_SEC is taken from the NSEC column for the total smart interval time on the RTM PRIVOPS Screen.

Note: Because of the way this number is calculated, it can be inaccurate, particularly for small values.

DIAG 14/CMD. The number of DIAGNOSE code Xc14c instructions used per command. DIAGNOSE code Xc14c is the CP function call to perform virtual spool I/O. This is calculated by:

$$\frac{\text{DIAG_14_SEC}}{\text{ETR (T)}}$$

DIAG_14_SEC is taken from the NSEC column for the total smart interval time on the RTM PRIVOPS Screen.

Note: Because of the way this number is calculated, it can be inaccurate, particularly for small values.

DIAG 58/CMD. The number of DIAGNOSE code Xc58c instructions used per command. DIAGNOSE code Xc58c is the CP function call that enables a virtual machine to communicate with 3270 virtual consoles. This is calculated by:

$$\frac{\text{DIAG_58_SEC}}{\text{ETR (T)}}$$

DIAG_58_SEC is taken from the NSEC column for the total smart interval time on the RTM PRIVOPS Screen.

Note: Because of the way this number is calculated, it can be inaccurate, particularly for small values.

DIAG 98/CMD. The number of DIAGNOSE code Xc98c instructions used per command. This allows a specified virtual machine to lock and unlock virtual pages and to run its own channel program. See the *VM/ESA System Facilities for Programming* for more information. This is calculated by:

$$\frac{\text{DIAG_98_SEC}}{\text{ETR (T)}}$$

DIAG_98_SEC is taken from the NSEC column for the total smart interval time on the RTM PRIVOPS Screen.

Note: Because of the way this number is calculated, it can be inaccurate, particularly for small values.

DIAG 98/CMD (V) VTAM Servers. See DIAG 98/CMD for a description of this instruction. This represents the sum of all DIAGNOSE code Xc98c instructions per command for all VTAM and VSCS servers. This is calculated by:

For VM/XA and VM/ESA:

$$\frac{\text{DIAG_98_VTAM} + \text{DIAG_98_VSCS}}{\text{ETR (T)}}$$

DIAG_98_VTAM and *DIAG_98_VSCS* are taken from the VM/PRF Virtual Machine Communication by User Class report for the VTAM and VSCS server classes respectively.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

$$\frac{1000}{\text{ETR (T)} \times \text{VMMAP_Time}} \times \sum_{s \in \text{server class}} \text{Non_Spool_IO}_s$$

Non_Spool_IO is from the NON SPOOLED I/O column for each server user ID in the VMMAP User Resource Utilization Part 2 report. *VMMAP_Time* is from the TOTAL SECONDS ANALYZED field at the beginning of the VMMAP report.

DIAG A4/CMD. The number of DIAGNOSE code XcA4c instructions used per command. DIAGNOSE code XcA4c is the CP function call that supports synchronous I/O to supported DASD. This is calculated by:

$$\frac{\text{DIAG_A4_SEC}}{\text{ETR (T)}}$$

DIAG_A4_SEC is taken from the NSEC column for the total smart interval time on the RTM PRIVOPS Screen.

Note: Because of the way this number is calculated, it can be inaccurate, particularly for small values.

DIAG A8/CMD. The number of DIAGNOSE code XcA8c instructions used per command. DIAGNOSE code XcA8c is the CP function call that supports synchronous general I/O to fully supported devices. This is calculated by:

$$\frac{\text{DIAG_A8_SEC}}{\text{ETR (T)}}$$

DIAG_A8_SEC is taken from the NSEC column for the total smart interval time on the RTM PRIVOPS Screen.

Note: Because of the way this number is calculated, it can be inaccurate, particularly for small values.

DIAG 214/CMD. The number of DIAGNOSE code Xc214c instructions used per command. DIAGNOSE code Xc214c is used by the Pending Page Release function. This is calculated by:

$$\frac{DIAG_214_SEC}{ETR (T)}$$

DIAG_214_SEC is taken from the NSEC column for the total smart interval time on the RTM PRIVOPS Screen.

Note: Because of the way this number is calculated, it can be inaccurate, particularly for small values.

DIAG/CMD. The total number of DIAGNOSE instructions used per command or job. This is calculated by:

For VM/XA or VM/ESA:

For VSECICS, FS7B, and IOB workloads:

$$\frac{1}{(ETR (T) \text{ or } ETR (C) \times RTM_Time)} \times \sum_{x \in \text{DIAGNOSE}} TOTALCNT_x$$

For PACE workload:

$$\frac{60}{(ETR (H) \times RTM_Time)} \times \sum_{x \in \text{DIAGNOSE}} TOTALCNT_x$$

TOTALCNT is the count for the individual DIAGNOSE codes taken over the total RTM time interval on the RTM PRIVOPS Screen. *RTM_Time* is the total RTM time interval taken from the RTM PRIVOPS screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

For FS7B and IOB workloads:

$$\frac{DIAG_Total}{RTM_Time \times ETR (T)}$$

For PACE workload:

$$60 \times \frac{DIAG_Total}{RTM_Time \times ETR (H)}$$

DIAG_Total is the value of DIAGNOSE for the total RTM time interval on the RTM PRIVOPS Screen. *RTM_Time* is the total RTM time interval taken from the RTM PRIVOPS screen.

DISPATCH LIST. The average over time of the number of virtual machines (including loading virtual machines) in any of the dispatch list queues (Q0, Q1, Q2 and Q3).

For VM/XA and VM/ESA this is calculated by:

$$\frac{1}{Num_Entries} \times \sum_{t \in \text{SCLOG entries}} Q0_t + Q0L_t + Q1_t + Q1L_t + Q2_t + Q2L_t + Q3_t + Q3L_t$$

Q0_t, *Q0L_t* .. are from the Q0CT, Q0L ... columns in the RTM SCLOG screen. *Num_Entries* is the total number of entries in the RTM SCLOG screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature) this is calculated by:

$$Q1 + Q2 + Q3$$

Q1, *Q2*, and *Q3* are taken from the AVG column of the *Q1*, *Q2* and *Q3* fields in the VM MAP Statistical Summary Report.

DPA (Dynamic Paging Area). The area of real storage used by CP to hold virtual machine pages, pageable CP modules and control blocks.

EDF. Enhanced Disk Format. This refers to the CMS minidisk file system.

Elapsed Time (H). The total time in seconds of a given measurement. This is the hardware estimate.

For 9221 Processors:

This is the Elapsed Time line in the RE0 file.

For 9121 and 9021 Processors this is calculated by:

$$Stop_Time - Start_Time$$

Stop_time and *Start_Time* are on the first line of the HISTDATA file.

ELIGIBLE LIST. The average over time of the number of virtual machines (including loading virtual machines) in any of the eligible list queues (E0, E1, E2 and E3).

For VM/XA and VM/ESA this is calculated by:

$$\frac{1}{Num_Entries} \times \sum_{t \in \text{SCLOG entries}} E0_t + E0L_t + E1_t + E1L_t + E2_t + E2L_t + E3_t + E3L_t$$

E0_t, *E0L_t* .. are from the E0CT, E0L ... columns in the RTM SCLOG screen. *Num_Entries* is the total number of entries in the RTM SCLOG screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature) this is calculated by:

$$E1 + E2$$

E1, and *E2* are taken from the AVG column of the *E1* and *E2* fields in the VM MAP Statistical Summary Report

EMUL ITR. Emulation Internal Throughput Rate. The average number of transactions completed per second of emulation time.

For VM/XA or VM/ESA:

This is from the EM_ITR field under TOTALITR of the RTM TRANSACT screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature), this is calculated by:

$$100 \times \frac{ETR}{\frac{TotEmul}{PROCESSORS}}$$

TotEmul is from the VIRTCPU field in the VM MAP Statistical Summary report.

EMUL/CMD. For the VSECICS, FS7B, and IOB workloads this is the amount of processor time spent in emulation mode per command in milliseconds. For the PACE workload this is the emulation processor time per job in seconds.

For the VSECICS, FS7B, and IOB workloads this is calculated by:

$$10 \times \frac{TOTAL_EMUL}{ETR (T) \text{ or } ETR (C)}$$

For the PACE workload this is calculated by:

$$6000 \times \frac{TOTAL_EMUL}{ETR (H)}$$

EMUL/CMD (H). See EMUL/CMD. This is the hardware based measurement.

For the VSECICS, FS7B, and IOB workloads this is calculated by:

$$10 \times \frac{\text{TOTAL EMUL (H)}}{\text{ETR (T) or ETR (C)}}$$

For the PACE workload this is calculated by:

$$6000 \times \frac{\text{TOTAL EMUL (H)}}{\text{ETR (H)}}$$

ETR. External Throughput Rate. The number of commands completed per second, computed by RTM.

For VM/XA or VM/ESA:

This is found in the NSEC column for ALL_TRANS for the total RTM interval time on the RTM Transaction screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This is the value of the CMNDRATE field in the VMMAP Statistical Summary.

ETR (C). See ETR. CICS-based calculation of ETR. It is calculated by:

$$\frac{1}{\text{CICS_Time}} \times \sum_{t \in \text{CICSPARfiles}} \text{Total_Transmits}_t$$

Total_Transmits is from the TOTAL TASKS SELECTED line the CICSPAR reports. *CICS_Time* is the run interval time, which is 900 seconds for all measurements.

ETR (H). The number of jobs completed per minute determined by the hardware instrumentation tools. It is calculated by:

$$60 \times \frac{\text{Jobs}}{\text{Elapsed Time (H)}}$$

Jobs is the number of jobs run in the workload. The values of *Jobs* are 28, 42 and 56 for the PACEX4, PACEX6 and PACEX8 workloads respectively.

ETR (T). See ETR. TPNS-based calculation of ETR. It is calculated by:

$$\sum_{t \in \text{TPNS machines}} \frac{\text{Total_Transmits}_t}{\text{TPNS_Time}_t}$$

Total_Transmits is found in the Summary of Elapsed Time and Times Executed section of TPNS report (TOTALS for XMITs by TPNS). *TPNS_Time* is the last time in requested (reduction) period minus the first time in requested (reduction) period. These times follow the Summary of Elapsed Time... in the TPNS report.

ETR RATIO. This is the ratio of the RTM-based ETR calculation (VM/XA or VM/ESA), or the VMMAP based ETR calculation (VM/SP and VM/ESA Release 1.0 (370 Feature)), and the TPNS based ETR calculation. This is calculated by :

$$\frac{\text{ETR}}{\text{ETR (T)}}$$

Expanded Storage. An optional integrated high-speed storage facility, available on certain processors, that allows for the rapid transfer of 4KB blocks between itself and real storage.

EXP. STORAGE. The amount of expanded storage used during a given run. See expanded storage.

External Response Time. The average response time, in seconds, for the last response to the screen. See AVG LAST (T).

FAST CLR/CMD. The number of fast path clears of real storage per command or job. This includes V=R and regular guests and is calculated by:

For VSECICS, FS7B, and IOB workloads:

$$\frac{\text{Fast_Clear_Sec}}{\text{ETR (T) or ETR (C)}}$$

For PACE workload:

$$60 \times \frac{\text{Fast_Clear_Sec}}{\text{ETR (H)}}$$

Fast_Clear_Sec is taken from the NSEC column for the total RTM time interval for the FAST_CLR entry on the RTM SYSTEM screen.

File Pool. In SFS, a collection of minidisks managed by a server machine.

FP REQ/CMD (Q). Total file pool requests per command. This is calculated by:

$$\sum_{f \in \text{filepools}} \frac{\text{Total_Filepool_Requests}_f}{\text{SFSTIME}_f}$$

Total_Filepool_Requests is from the QUERY FILEPOOL STATUS command.

FREE TOTL/CMD. The number of requests for free storage per command or job. This includes V=R and regular guests and is calculated by:

For VSECICS, FS7B, and IOB workloads:

$$\frac{\text{Free_Total_Sec}}{\text{ETR (T) or ETR (C)}}$$

For PACE workload:

$$60 \times \frac{\text{Free_Total_Sec}}{\text{ETR (H)}}$$

Free_Total_Sec is taken from the NSEC column for the total RTM time interval on the RTM SYSTEM screen.

Note: Because of the way this number is calculated, it can be inaccurate, particularly for small values.

FREE UTIL. The proportion of the amount of available free storage actually used. This is calculated by :

$$\frac{\text{Free_Size}}{\text{FREEPGS} \times 4096}$$

Free_Size is found in the FREE column for the total RTM time interval (<..>) on the RTM SYSTEM screen.

FREEPGS. The total number of pages used for FREE storage (CP control blocks).

For VM/XA or VM/ESA:

This is found in the FPGS column for the total RTM time interval (<..) on the RTM SYSTEM screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature), this is calculated by:

$$\frac{Free_Size}{4096} + Free_Xtend$$

Free_Size is from the FREE STORAGE SIZE field at the beginning of the VM MAP report. *Free_Xtend* is from the FREXTNDPG field in the VM MAP Statistical Summary report.

FST (File Status Table). CMS control block that contains information about a file belonging to a minidisk or SFS directory.

GB. Gigabytes. 1024 megabytes.

GRAF-ID. Display terminal address.

GUEST SETTING. This field represents the type of VSE guest virtual machine in a PACE measurement. This fields possible values are V=V, V=F or V=R.

GUESTWT/CMD. The number of entries into guest enabled wait state per job. This is calculated by:

$$60 \times \frac{GUESTWT/SEC}{ETR (H)}$$

GUESTWT/SEC. The number of entries into guest enabled wait state per second.

This field is taken from the NESC column for the RTM total count since last reset, for the GUESTWT field in the RTM SYSTEM screen.

Hardware Instrumentation.. See Processor Instrumentation

IML MODE. This is the hardware IML mode used in VSE guest measurements. The possible values for this field are 370, ESA, or LPAR.

Instruction Path Length. The number of machine instructions used to run a given command, function or piece of code.

Internal Response Time. The response time as seen by CP. This does not include line or terminal delays.

IO TIME/CMD (Q). Total elapsed time in seconds spent doing SFS file I/Os per command. This is calculated by:

$$\frac{1}{(1000 \times ETR (T))} \times \sum_{f \in \text{filepools}} \frac{Total_BIO_Request_Time_f}{SFSTIME_f}$$

Total_BIO_Request_Time is from the QUERY FILEPOOL STATUS command.

IO/CMD (Q). SFS file I/Os per command. This is calculated by:

$$\frac{1}{ETR (T)} \times \sum_{f \in \text{filepools}} \frac{Total_IO_Requests_f}{SFSTIME_f}$$

Total_IO_Requests is from the QUERY FILEPOOL STATUS command.

ITR. Internal Throughput Rate. This is the number of units of work accomplished per unit of processor busy time in a nonconstrained environment. For the FS7B and IOB workloads this is represented as commands per processor second. For the PACE workload this is represented as jobs per processor minute. This is calculated by:

For FS7B and IOB workloads running on VM/XA or VM/ESA:

This is found from the TOTALITR for SYS_ITR on the RTM TRANSACT screen.

For FS7B and IOB workloads running on VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

$$100 \times \frac{ETR}{UTIL/PROC}$$

For PACE workload:

$$100 \times \frac{ETR (H)}{UTIL/PROC}$$

ITR (H). See ITR. This is the hardware based measure. In this case, ITR is measured in external commands per unit of processor busy time. For the FS7B and IOB workloads this is represented as commands per processor second, while for the PACE workload this is represented in jobs per processor minute. This is calculated by:

For the VSECICS, FS7B, and IOB workloads:

$$100 \times \frac{ETR (C) \text{ or } ETR (T)}{TOTAL (H)}$$

For the PACE workloads:

$$6000 \times \frac{Jobs}{Elapsed \text{ time } (H) \times TOTAL (H)}$$

Jobs is the number of jobs run in the workload. The values of *Jobs* are 28, 42 and 56 for the PACEX4, PACEX6 and PACEX8 workloads respectively.

ITRR. Internal Throughput Rate Ratio. This is the RTM based ITR (VM/XA or VM/ESA), or the VM MAP based ITR (VM/SP, VM/SP HPO, and VM/ESA Release 1.0 (370 Feature)), normalized to a specific run. This is calculated by:

$$\frac{ITR}{ITR_1}$$

ITR₁ is the ITR of the first run in a given table.

ITRR (H). See ITRR. This is the ITR (H) normalized to a specific run. This is calculated by:

$$\frac{ITR (H)}{ITR (H)_1}$$

ITR (H)₁ is the ITR (H) of the first run in a given table.

Inter-user Communication Vehicle (IUCV). A VM generalized CP interface that helps the transfer of messages either among virtual machines or between CP and a virtual machine.

k. Multiple of 1000.

Kb. Kilobits. One kilobit is 1024 bits.

KB. Kilobytes. One kilobyte is 1024 bytes.

LUW Rollbacks (delta). The total number of SFS logical units of work that were backed out during the measurement interval, summed over all production file pools. This is calculated by:

$$\sum_{f \in \text{filepools}} LUW_Rollbacks_f$$

LUW_Rollbacks is from the QUERY FILEPOOL STATUS command.

MASTER EMUL. Total emulation state utilization for the master processor. For uniprocessors this is the same as TOTAL EMUL and is generally not shown.

For VM/XA or VM/ESA:

This is taken from the %EM column for the first processor listed in the LOGICAL CPU STATISTICS section of the RTM CPU screen. The total RTM interval time value is used (<..).

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This is taken from the VIRTCPU line in the VMMAP Statistical Summary report.

MASTER EMUL (H). Total emulation state utilization for the master processor. For uniprocessors this is the same as TOTAL EMUL and is generally not shown. This is the hardware based calculation.

For VM/XA or VM/ESA:

This is taken from the %CPU column of the GUES-CPU section of the HISTRPT file for the master processor number as shown by RTM. In RTM, the first processor listed on the CPU screen is the master processor.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

All VM/SP or VM/ESA Release 1.0 (370 Feature) measurements in this report are uniprocessor measurements; hence, this is equal to TOTAL EMUL (H)

MASTER TOTAL. Total utilization of the master processor. For uniprocessor this is the same as TOTAL and is generally not shown.

For VM/XA or VM/ESA:

This is taken from the %CPU column for the first processor listed in the LOGICAL CPU STATISTICS section of the RTM CPU screen. The total RTM interval time value is used (<..).

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This is taken from the TOTCPU line in the VMMAP Statistical Summary report.

MASTER TOTAL (H). Total utilization of the master processor. For uniprocessor this is the same as

TOTAL (H) and is generally not shown. This is the hardware based calculation.

For VM/XA or VM/ESA:

This is taken from the %CPU column of the Syst-CPU section of the HISTRPT file for the master processor number as shown by RTM. In RTM, the first processor listed on the CPU screen is the master processor.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

All VM/SP and VM/ESA Release 1.0 (370 Feature) measurements in this report are uniprocessor measurements; hence, this is equal to TOTAL (H).

MB. Megabytes. One megabyte is 1,048,576 bytes.

MDC Hit Ratio. Minidisk Cache Hit Ratio. The ratio of requested blocks that are found in expanded storage for a read operation to the total number of blocks read.

This is from the RTM MDHR field for the total RTM time interval (<..) on the RTM SYSTEM screen.

MDC Mods. Minidisk Cache Modifications. The number of times per second blocks were written in the cache, excluding the writes that occurred as a result of minidisk cache misses.

This is taken from the NSEC column for the RTM MDC_MO field for the total RTM time interval on the RTM SYSTEM screen.

MDC Reads. Minidisk Cache Reads. The number of times per second blocks were found in the cache as the result of a read operation.

This is taken from the NSEC column for the RTM MDC_HT field for the total RTM time interval on the RTM SYSTEM screen.

MDC Writes. Minidisk Cache Writes. The number of CMS Blocks moved per second from main storage to expanded storage.

This is taken from the NSEC column for the RTM MDC_PW field for the total RTM time interval on the RTM SYSTEM screen.

MDSK/CMD. The average number of DIAGNOSE minidisk I/O requests per command for all users in the system. This is calculated by:

$$\frac{MDSK_SEC}{ETR (T)}$$

MDSK_SEC is found in the average line for the DIAG DISK I/O column in the VMMAP Summary of Privileged Instruction Simulation Summary - Part 2.

Millisecond. One one-thousandth of a second.

Minidisk Caching. Refers to the use of a portion of expanded storage as a write-through, (data written to minidisks is also written to the cache), minidisk cache for CMS 4KB block size data. It is used to help elimi-

nate I/O bottlenecks and improve system response time by replacing actual DASD I/Os with faster expanded storage I/O.

ms. Millisecond.

Non-dormant Time. The percent of time that a user is not on the dormant list or SVM wait.

Non-shared Storage. The portion of a virtual machine's storage that is unique to that virtual machine, (as opposed to shared storage such as a saved segment that is shared among virtual machines). This is usually represented in pages.

NONTRIV INT. Non-trivial Internal response time in seconds. The number of transactions that completed with more than one drop from Q1 or one or more drops from Q0, Q2, or Q3 per second.

For VM/XA or VM/ESA:

This is from TOTALTTM for the RTM NTRIV field on the RTM TRANSACT screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature), this is calculated by:

$$\frac{\text{Minor_Resp} \times \text{Minor_Rate} + \text{Major_Resp} \times \text{Major_Rate}}{\text{Minor_Rate} + \text{Major_Rate}}$$

Minor_Resp and *Major_Resp* are the MINORESP and MAJORESP fields in the VMMAP Statistical Summary report respectively. *Minor_Rate* and *Major_Rate* are the MINORATE and MAJORATE fields in the VMMAP Statistical Summary report respectively.

NUCLEUS SIZE (V). The resident CP nucleus size in kilobytes.

For VM/XA or VM/ESA:

This is from the <K bytes> column on the Total Resident Nucleus line in the VMPRF System Configuration Report.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This is from the NUCLEUS SIZE line on the first page of the VMMAP report.

PAGE/CMD. The number of pages moved between real storage and DASD per command or job. This is calculated by:

For VSECICS, FS7B, and IOB workloads:

$$\frac{\text{READS/SEC} + \text{WRITES/SEC}}{\text{ETR (T) or ETR(C)}}$$

For PACE workload:

$$60 \times \frac{\text{READS/SEC} + \text{WRITES/SEC}}{\text{ETR (H)}}$$

PAGE IO RATE (V). The number of real SSCH or RSCH instructions issued on behalf of system paging.

For VM/XA or VM/ESA:

This is the sum of all the entries in the SSCH+RSCH column for Page devices listed in the VMPRF DASD System Areas by Type report.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This is the sum of all entries in the Total I/O Issued column for Paging devices (plus swapping devices for VM/SP HPO) listed in the VMMAP Disk and Tape I/O Activity Summary report divided by *VMMAP_time*. *VMMAP_time* is from the TOTAL SECONDS ANALYZED field at the beginning of the VMMAP report.

PAGE IO/CMD (V). The number of real SSCH and RSCH instructions issued per command on behalf of system paging. This is calculated by:

$$\frac{\text{PAGE IO RATE (V)}}{\text{ETR (T)}}$$

Path length. See Instruction Path Length

PBT/CMD. For the VSECICS, FS7B, and IOB workloads, this is the number of milliseconds of processor activity per command. For the PACE workload, this is the number of seconds of processor activity per job. This is calculated by:

For the VSECICS, FS7B, and IOB workloads:

$$10 \times \frac{\text{TOTAL}}{\text{ETR (T) or ETR(C)}}$$

For the PACE workload:

$$6000 \times \frac{\text{TOTAL}}{\text{ETR (H)}}$$

PBT/CMD (H). See PBT/CMD. This is the hardware based measure.

For the VSECICS, FS7B, and IOB workloads:

$$10 \times \frac{\text{TOTAL (H)}}{\text{ETR (T) or ETR(C)}}$$

For the PACE workload:

$$6000 \times \frac{\text{TOTAL (H)}}{\text{ETR (H)}}$$

PD3. One of the CMS-intensive workloads used in the Large Systems Performance Reference (LSPR) to evaluate relative processor performance.

PGBLPGS. The number of system pageable pages available.

For VM/XA or VM/ESA:

This is from the PPAG field for the total RTM time interval (<-) on the RTM SYSTEM screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature), this is calculated by:

$$\frac{\text{DPA}}{4096}$$

DPA is from the DYNAMIC PAGING AREA field at the beginning of the VMMAP report.

PGBLPGS/USER. The number of system pageable pages available per user. This is calculated by:

$$\frac{\text{PGBLPGS}}{\text{USERS}}$$

Privileged Operation. Any instruction that must be run in supervisor state.

PRIVOP/CMD. The number of virtual machine privileged instructions simulated per command or job. This does not include DIAGNOSE instructions. This is calculated by:

For VSECICS, FS7B, and IOB workloads:

$$\frac{1}{(\text{ETR (T) or ETR (C)} \times \text{RTM_Time})} \times \sum_{x \in \text{privops}} \text{TOTALCNT}_x$$

For PACE workload:

$$\frac{60}{(\text{ETR (H)} \times \text{RTM_Time})} \times \sum_{x \in \text{privops}} \text{TOTALCNT}_x$$

TOTALCNT is the count for the individual privop taken over the total RTM time interval on the RTM PRIVOPS Screen. *RTM_Time* is the total RTM time interval taken from the RTM PRIVOPS screen. **Note:** PRIVOPS are recorded differently in 370 and XA modes.

PRIVOPS (Privileged Operations). See Privileged Operation.

Processor Instrumentation. An IBM internal tool used to obtain hardware-related data such as processor utilizations and throughput rates.

Processor Utilization. The percent of time that a processor is not idle.

PROCESSORS. The data field denoting the number of processors that were active during a measurement.

For VM/XA or VM/ESA:

This is from the NC field under CPU statistics on the RTM CPU screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This value is input by hand, and has a value of 1 for all measurements in this report except for the VM/SP HPO Release 5.0 measurements, for which the value is 2.

Production File Pool. An SFS file pool in which users are enrolled with space. All SFS read/write activity is to production file pools.

QUICKDSP ON. When a virtual machine is assigned this option, it bypasses the normal scheduler algorithm and is placed on the dispatch list immediately when it has work to do. It does not spend time in the eligible lists. QUICKDSP can be specified either via a CP command or in the CP directory entry.

READS/SEC. The number of pages read per second done for system paging.

For VM/XA or VM/ESA:

This is taken from the NSEC column for the PAGREAD field for the total RTM time interval on the RTM SYSTEM screen.

For VM/SP or VM/ESA Release 1.0 (370 Feature):

This is from the PREAD field in the VM MAP Statistical Summary report.

For VM/SP HPO:

This is the sum of the PREAD and SWPGREAD fields in the VM MAP Statistical Summary report.

REAL STORAGE. The amount of real storage used for a particular measurement.

Relative Share. For CP/ESA, determines the virtual machine's scheduling priority.

RELSHARE. The virtual machine's relative share.

RESERVE. See SET RESERVED

RIO/CMD (H). The number of real SSCH (or SIO for 370) and RSCH (or RIO for 370) instructions issued per command or job. This is calculated by:

For VSECICS, FS7B, and IOB workloads:

$$\frac{\text{RIO RATE (H)}}{\text{ETR (T) or ETR (C)}}$$

For PACE workload:

$$60 \times \frac{\text{RIO RATE (H)}}{\text{ETR (H)}}$$

RIO/CMD (V). The number of real SSCH (or SIO for 370) and RSCH (or RIO for 370) instructions issued per command. This is calculated by:

For FS7B and IOB workloads:

$$\frac{\text{RIO RATE (V)}}{\text{ETR (T)}}$$

For PACE workload:

$$60 \times \frac{\text{RIO RATE (V)}}{\text{ETR (H)}}$$

RIO RATE (H). The number of real SSCH (or SIO for 370) and RSCH (or RIO for 370) instructions issued per second.

For VM/XA or VM/ESA:

This is taken from the VM PRF report. See RIO RATE (V).

For VM/SP or VM/ESA Release 1.0 (370 Feature) this is calculated by:

$$\frac{\text{RIO} + \text{SIO}}{\text{H_Time}}$$

RIO is taken from RIO or RSCH line and *SIO* is taken from the SIOF or SSCH line in the RE0 report. *H_Time* is taken from the Elapsed Time line in the RE0 report.

RIO RATE (V). The number of real SSCH (or SIO for 370) and RSCH (or RIO for 370) instructions issued per second.

For VM/XA or VM/ESA:

This is taken from the I/O Rate column for the overall average on the VM PRF System Performance Summary by Time report; the value reported does not include assisted I/Os.

For VM/SP or VM/ESA Release 1.0 (370 Feature):

This field is taken from the hardware report. See RIO RATE (H).

For VM/SP HPO:

This is the sum of RATE/SEC in VM MAP's Channel Activity Summary report.

Rollback Requests (delta). The total number of SFS rollback requests made during a measurement. This is calculated by:

$$\sum_{f \in \text{filepools}} \text{Rollback_Requests}_f$$

Rollback_Requests is from the QUERY FILEPOOL STATUS command.

Rollbacks Due to Deadlock (delta). The total number of LUW rollbacks due to deadlock that occurred during the measurement interval over all production file pools. A rollback occurs whenever a deadlock condition cannot be resolved by the SFS server. This is calculated by:

$$\sum_{f \in \text{filepools}} \text{Rollbacks_Due_to_Deadlock}_f$$

Rollbacks_Due_to_Deadlock is from the QUERY FILEPOOL STATUS command.

***RPI/CPU sec (V).** Data transfers, per CPU second, for successful IUCV data transfers to the CP system service, *RPI.

$$\frac{\text{*RPI/sec Total (V)} \times \text{VM PRF_time}}{\sum_{f \in \text{RACF machines}} \text{Total_CPU_Secs}_f}$$

VM PRF_time is the total elapsed time, in seconds, for which the *RPI/sec Total (V) is calculated. *Total_CPU_Secs* is from the Resource Utilization by User Class section of the VM PRF report.

***RPI/CMD (V).** Data transfer rate, per command, for successful IUCV data transfers to the CP system service *RPI. This is calculated by:

$$\frac{\text{RPI/sec Total (V)}}{\text{ETR (T)}}$$

***RPI/sec Total (V).** Data transfer rate, per second, for successful IUCV data transfers by a user to the CP system service, *RPI.

This is from the Data Transfer Rates, CP Services, *RPI (summary) field in the VM PRF Rates of Data Transfers by Users to CP System Services report.

***RPI/sec RACFVMn (V).** (*n* is blank, 1, 2, or 3) Data transfer rate, per second, for successful IUCV data transfers sent by the indicated RACF virtual machine connected to the *RPI system service.

This is from the Message Rate, IUCV Send field of the Virtual Machine Communications - User to User Samples VM PRF report.

RSTOR RATIO. This is the amount of REAL STORAGE normalized to a specific run. This is calculated by:

$$\frac{\text{REAL STORAGE}}{\text{REAL STORAGE}_1}$$

REAL STORAGE₁ is the REAL STORAGE of the first run in a given table.

RTM. Realtime Monitor. A licensed program realtime monitor and diagnostic tool for performance monitoring, analysis, and problem solving.

RTM/ESA. See RTM.

Run ID. An internal use only name used to identify a performance measurement.

SAC Calls / FP Request. The average number of calls within the SFS server to its Storage Access Component (SAC) per file pool request. In environments where there are multiple file pools, this average is taken over all file pool servers. This is calculated by:

$$\frac{\sum_{f \in \text{filepools}} \frac{\text{Sac_Calls}_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{\text{Total_Filepool_Requests}_f}{\text{SFSTIME}_f}}$$

Sac_Calls and *Total_Filepool_Requests* are from the QUERY FILEPOOL STATUS command.

Seconds Between Checkpoints. The average number of seconds between SFS file pool checkpoints in the average file pool. This is calculated by:

$$\sum_{f \in \text{filepools}} \frac{1}{\frac{\text{Checkpoints_Taken}_f}{\text{SFSTIME}_f}}$$

Checkpoints_Taken is from the QUERY FILEPOOL STATUS command.

SET RESERVED (Option). This is a CP command that can be used to allow a V=V virtual machine to have a specified minimum number of pages resident in real storage. It is used to reduce paging and improve performance for a given virtual machine.

SFSTIME. The elapsed time in seconds between QUERY FILEPOOL STATUS invocations for a given file pool done at the beginning and end of a measurement.

SFS TIME/CMD (Q). Total elapsed time per command, in seconds, required to process SFS server requests. This is calculated by:

$$\frac{1}{\text{ETR (T)}} \times \sum_{f \in \text{filepools}} \frac{\text{Filepool_Request_Service_Time}_f}{\text{SFSTIME}_f}$$

Filepool_Request_Service_Time is from the QUERY FILEPOOL STATUS command.

Shared Storage. The portion of a virtual machines storage that is shared among other virtual machines (such as saved segments). This is usually represented in pages.

SHRPGS. The number of shared frames currently resident.

For VM/XA or VM/ESA:

This is from the SHAR field for the total RTM time interval (<-) on the RTM SYSTEM screen.

For VM/SP or VM/ESA Release 1.0 (370 Feature):

This is from the SHRPGS field in the VMMAP Statistical Summary report.

SIE. ESA Architecture instruction to Start Interpretive Execution. This instruction is used to run a virtual machine in emulation mode.

SIE INTCPT/CMD. The number of exits from SIE which are SIE interceptions per command or job. SIE is exited either by interception or interruption. An intercept is caused by any condition that requires CP interaction such as I/O or an instruction that has to be simulated by CP. This is calculated by:

$$\frac{\text{Percent_Intercept} \times \text{SIE/CMD}}{100}$$

Percent_Intercept is taken from the %SC field for average of all processors for the total RTM time interval (<-..) on the RTM CPU screen.

SIE/CMD. SIE instructions used per command or job. This is calculated by:

For VSECICS, FS7B, and IOB workloads:

$$\frac{\text{SIE_SEC}}{\text{ETR (T) or ETR (C)}}$$

For PACE workload:

$$60 \times \frac{\text{SIE_SEC}}{\text{ETR (H)}}$$

SIE_SEC is taken from the XSI field for the total for all processors for the total RTM time interval (<-..) on the RTM CPU screen.

Special Operations. Instructions that cause either CP or the hardware licensed internal code to provide services for the virtual machine. This includes privileged instructions, normally invalid opcodes (DIAGNOSE instructions for example), and the first reference to a non-shared page.

T/V Ratio. See TVR

TOT CPU/CMD (V) Server. The total amount of processor time, in milliseconds, for the server virtual machine(s). This is calculated by:

$$\frac{1}{(V_Time \times \text{ETR (T)})} \times \sum_{s \in \text{server class}} \text{Total_CPU_Secs}_s$$

For VM/XA or VM/ESA:

Total_CPU_Secs and *V_Time* are from the Resource Utilization by User Class section of the VMPRF reports.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

Total_CPU_Secs is from the VMMAP User Resource Utilization report for each server user ID. *V_Time* is from the TOTAL SECONDS ANALYZED field at the beginning of the VMMAP report.

TOT INT. Total Internal Response Time in seconds. Internal response time averaged over all trivial and non-trivial transactions.

For VM/XA and VM/ESA:

This is the value for TOTALTTM for ALL_TRANS on the RTM TRANSACT screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This is the value of the AVGRES field in the VMMAP Statistical Summary report.

TOT INT ADJ. Total internal response time (TOT INT) reported by RTM (for VM/XA or VM/ESA), or by VMMAP (For VM/ESA Release 1.0 (370 Feature)), adjusted to reflect what the response time would have been had CP seen the actual command rate (as recorded by TPNS). This is a more accurate measure of internal response time than TOT INT. In addition, TOT INT ADJ can be directly compared to external response time (AVG LAST (T)) as they are both based on the same, TPNS-based measure of command rate. This is calculated by:

$$\text{TOT INT} \times \text{ETR RATIO}$$

TOTAL. The total processor utilization for a given measurement summed over all processors.

For VM/XA or VM/ESA:

This comes from the %CPU column for all processors for the total RTM interval time (<-..) on the RTM CPU screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This is the SYSTCPU field in the VMMAP Statistical Summary report.

TOTAL (H). See TOTAL. This is the hardware based measurement.

For 9221 Processors:

This is taken from the Total CPU Busy line in the CPU Busy/Mips section of the RE0 report.

For 9121 and 9021 Processors, this is calculated by:

$$\text{UTIL/PROC (H)} \times \text{PROCESSORS}$$

TOTAL EMUL. The total emulation state time for all users across all online processors. This indicates the percentage of time the processors are in emulation state.

For VM/XA or VM/ESA:

This comes from the %EM column for all processors for the total RTM interval time (<-.) on the RTM CPU screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This is the SYSVCPU field in the VMMAP Statistical Summary report.

TOTAL EMUL (H). The total emulation state time for all users across all online processors. This indicates the percentage of time the processors are in emulation state. This is calculated by:

For 9221 Processors:

This comes from the SIE CPU Busy / Total CPU Busy (PCT) line in the RE0 report (VM/XA and VM/ESA) and the 370 PP CPU Busy / Total CPU Busy (PCT) line (VM/SP and VM/ESA Release 1.0 (370 Feature)).

For 9121 and 9021 Processors:

This is taken from the %CPU column for the GUES-CPU section of the HISTRPT file.

TPNS. Teleprocessing Network Simulator. A licensed program terminal and network simulation tool that provides system performance and response time information.

Transaction. A user/system interaction as counted by CP. For a single-user virtual machine a transaction should roughly correspond to a command. It does not include network or transmission delays and may include false transactions. False transactions can be those that wait for an external event, causing them to be counted as multiple transactions, or those that process more than one command without dropping from queue, causing multiple transactions to be counted as one.

TRACE TABLE (V). The size in kilobytes of the CP trace table.

For VM/XA or VM/ESA:

This is the value of the <K bytes> column on the Trace Table line in the VMPRF System Configuration Report.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This is from the INTERNAL TRACE TABLE line on the first page of the VMMA report.

Transaction (T). This is the interval from the time the command is issued until the last receive prior to the next send. This includes clear screens as a result of an intervening MORE... or HOLDING condition.

TRIV INT. Trivial Internal Response Time in seconds. The rate per second for transactions that complete with one and only one drop from Q1 and no drops from Q0, Q2, and Q3.

For VM/ESA or VM/XA:

This is from TOTALTTM for the TRIV field on the RTM TRANSACT screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This is from the TRIVRESP field in the VMMAP Statistical Summary report.

TVR. Total to Virtual Ratio. This is the ratio of total processor utilization to virtual processor utilization. This is calculated by:

$$\frac{\text{TOTAL}}{\text{TOTAL EMUL}}$$

TVR (H). See TVR. Total to Virtual Ratio measured by the hardware monitor. This is calculated by:

$$\frac{\text{TOTAL (H)}}{\text{TOTAL EMUL (H)}}$$

USERS. The number of virtual machines logged on to the system during a measurement interval that are associated with simulated end users. This includes active and inactive virtual machines but does not include service machines.

User Priority. Priority of the virtual machine specified in the VM directory that is used by VM/SP scheduler when placing virtual machines in the run list.

UTIL/PROC. Per processor utilization. This is calculated by:

$$\frac{\text{TOTAL}}{\text{PROCESSORS}}$$

UTIL/PROC (H). Processor utilization reported by the hardware.

For 9221 Processors, this is calculated by:

$$\frac{\text{TOTAL (H)}}{\text{PROCESSORS}}$$

For 9121 and 9021 Processors:

This is taken from the %CPU column in the SYST-CPU section of the HISTRPT file.

VIO RATE. The total number of all virtual I/O requests per second for all users in the system.

For VM/XA or VM/ESA:

This is from the ISEC field for the total RTM time interval (<-) on the RTM SYSTEM screen.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

This is from the VIO field in the VMMAP Statistical Summary report.

VIO/CMD. The average number of virtual I/O requests per command or job for all users in the system. This is calculated by:

For VSECICS, FS7B, and IOB workloads:

$$\frac{\text{VIO RATE}}{\text{ETR (T) or ETR (C)}}$$

For PACE workload:

$$60 \times \frac{\text{VIO RATE}}{\text{ETR (H)}}$$

VIRT CPU/CMD (V) Server. Virtual processor time, in milliseconds, run in the designated server(s) machine per command. This is calculated by:

$$\frac{1}{(V_Time \times ETR (T))} \times \sum_{s \in \text{server class}} Virt_CPU_Secs_s$$

For VM/XA or VM/ESA

Virt_CPU_Secs and *V_Time* are from the Resource Utilization by User Class section of the VMPRF reports.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

Virt_CPU_Secs is from the VMMAP User Resource Utilization report for each server user ID. *VMMAP_Time* is from the TOTAL SECONDS ANALYZED field at the beginning of the VMMAP report.

VM MODE. This field is the virtual machine setting (370, XA or ESA) of the VSE guest virtual machine in PACE and VSECICS measurements.

VM SIZE. This field is the virtual machine storage size of the VSE guest virtual machine in PACE and VSECICS measurements.

VMMAP. Virtual Machine Monitor Analysis Program. A licensed program that produces performance reports and history files for VM/SP and VM/ESA Release 1.0 (370 Feature) data.

VMPAF. Virtual Machine Performance Analysis Facility. A tool used for performance analysis of VM systems.

VMPRF. VM Performance Reporting Facility. A licensed program that produces performance reports and history files from VM/XA or VM/ESA monitor data.

VSCSs. The number of virtual machines running VSCS external to VTAM during a measurement interval.

VSE SUPERVISOR. This field is the VSE supervisor mode used in a PACE or VSECICS measurement. This field's possible values are 370 and ESA.

VSE WKSET (V). The average working set size in pages of the VSE guest machine in the VSECICS or PACE workload.

For VM/XA or VM/ESA:

This is found in the WSS column in the VMPRF Resource Utilization by User Class report for the VSE user class.

For VM/SP or VM/ESA Release 1.0 (370 Feature):

$$\frac{WSS}{4}$$

WSS is found in the WSS column in the VMMAP User Resource Utilization report for the VSE guest user ID.

VTAMs. The number of virtual machines running VTAM during a measurement interval.

VUMAPC. Virtual Machine UP/MP Analysis Program. An internal tool used to reduce the processor instrumentation data.

V=F. Virtual equals fixed machine. A virtual machine that has a fixed, contiguous area of real storage. Unlike V=R, storage does not begin at page 0. For guests running V=F, CP does not page this area. Requires the PR/SM hardware feature to be installed.

V=R. Virtual equals real machine. Virtual machine that has fixed, contiguous area of real storage starting at page 0. CP does not page this area.

V=V. Virtual equals virtual machine. Default storage processing. CP pages the storage of a V=V machine in and out of real storage.

WKSET (V). The average working set size. This is the scheduler's estimate of the amount of storage the average user will require, in pages.

For VM/XA or VM/ESA:

This is the average of the values for WSS in the VMPRF Resource Utilization by User report, (found in the Sum/Avg line).

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature) this is calculated by:

$$\frac{WKSET}{4}$$

WKSET is in the VMMAP Statistical Summary report.

WKSET (V) Server. Total working set of a related group of server virtual machine(s). This is calculated by:

For VM/XA or VM/ESA:

$$\sum_{s \in \text{server Logged Users}} Avg_WSS_s$$

Avg_WSS is found in the Avg WSS column in the VMPRF Resource Utilization by User Class report for each class of server.

For VM/SP, VM/SP HPO, or VM/ESA Release 1.0 (370 Feature):

$$\frac{1}{4} \times \sum_{s \in \text{server class}} Working_Set_Size_s$$

Working_Set_Size is found in the WSS column in the VMMAP User Resource Utilization report for each server user ID.

WRITES/SEC. The number of page writes per second done for system paging.

For VM/XA or VM/ESA:

This is taken from the NSEC column for the PAWRIT field for the total RTM time interval on the RTM SYSTEM screen.

For VM/SP or VM/ESA Release 1.0 (370 Feature):

This is from the PWRITE field in the VMMAP Statistical Summary report.

For VM/SP HPO:

This is the sum of PWRITE and SWPGWRITE from the VMMAP Statistical Summary report.

XSTOR IN/SEC. The number of pages per second read into main storage from expanded storage. This includes fastpath and non-fastpath pages and is calculated by:

$$\text{Fastpath_In} + \text{NonFastpath_In}$$

Fastpath_In and *NonFastpath_In* are taken from the NSEC column for the XST_PGIF and XST_PGIS fields for the total RTM time interval on the RTM SYSTEM screen.

For VM/SP HPO:

This is the sum of PPPAGEIN and SWPAGEIN fields in VMMAP's Statistical Summary report.

XSTOR OUT/SEC. The number of pages per second written from main storage into expanded storage.

This is taken from the NSEC column for the XST_PGO field for the total RTM time interval on the RTM SYSTEM screen.

For VM/SP HPO:

This is the sum of PPPAGEOUT and SWPAGEOUT fields in VMMAP's Statistical Summary report.

XSTOR RATIO. This is the amount of EXP. STORAGE normalized to a specific run. This is calculated by:

$$\frac{\text{EXP. STORAGE}}{\text{EXP. STORAGE}_1}$$

EXP. STORAGE₁ is the EXP. STORAGE of the first run in a given table.

XSTOR/CMD. The number of pages read into main storage from expanded storage and written to expanded storage from main storage per command or job. This is calculated by:

For VSECICS, FS7B, and IOB workloads:

$$\frac{\text{XSTOR IN/SEC} + \text{XSTOR OUT/SEC}}{\text{ETR (T) or ETR (C)}}$$

For PACE workload:

$$60 \times \frac{\text{XSTOR IN/SEC} + \text{XSTOR OUT/SEC}}{\text{ETR (H)}}$$

Communicating Your Comments to IBM

VM/ESA Release 2 Performance Report
Enterprise Systems
Services and Support
Washington Systems Center
Technical Bulletin
Publication No. GG66-3245-00

If you especially like or dislike anything about this book, please use one of the methods listed below to send your comments to IBM. Whichever method you choose, make sure you send your name, address, and telephone number if you would like a reply.

Feel free to comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. However, the comments you send should pertain to only the information in this manual and the way in which the information is presented. To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

If you are mailing a readers' comment form (RCF) from a country other than the United States, you can give the RCF to the local IBM branch office or IBM representative for postage-paid mailing.

- If you prefer to send comments by mail, use the RCF at the back of this book.
- If you prefer to send comments by FAX, use this number:
 - United States & Canada: 301+240-8836
 - Other countries: (+1)+301+240-8836

Make sure to include the following in your note:

- Title and publication number of this book
- Page number or topic to which your comment applies.

Readers' Comments

**VM/ESA Release 2 Performance Report
Enterprise Systems
Services and Support
Washington Systems Center
Technical Bulletin
Publication No. GG66-3245-00**

Use this form to tell us what you think about this manual. If you have found errors in it, or if you want to express your opinion about it (such as organization, subject matter, appearance) or make suggestions for improvement, this is the form to use.

To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer. This form is provided for comments about the information in this manual and the way it is presented.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Be sure to print your name and address below if you would like a reply.

Name

Address

Company or Organization

Phone No.




Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES



BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

Chuck Morse
International Business Machines Corporation
Washington Systems Center
800 N FREDERICK AVENUE
GAITHERSBURG MD 20879-3395



Fold and Tape

Please do not staple

Fold and Tape



IBM Internal Use Only - This copy not for customer use

Printed in U.S.A.

GG66-3245-00

