GC24-5673-01

Virtual Machine/
Enterprise Systems Architecture

Release 2.2 Performance Report

Virtual Machine/
Enterprise Systems Architecture

GC24-5673-01

Release 2.2 Performance Report

Note!

Before using this information and the product it supports, be sure to read the general information under “Notices” on page v.

This book is also provided as an online book that can be viewed with the IBM* BookManager* READ and IBM Library Reader* licensed programs.

Second Edition (June 1994)

This edition, GC24-5673-01, applies to Virtual Machine/Enterprise Systems Architecture (VM/ESA), Release 2.2, program number 5684-112, and to all subsequent releases of this product until otherwise indicated in new editions or Technical Newsletters. Changes are made periodically to the information herein.

This report includes extensive performance data changes from release to release. For this reason, change bars are not shown.

Publications are *not* stocked at the address given below. U.S. customers can order publications by calling the IBM Software Manufacturing Company at 1-800-879-2755. Outside the U.S., customers should contact the IBM branch office serving their locality.

A form for readers′ comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, Information Development, Dept. G60, 1701 North Street, Endicott, NY, 13760-5553 USA. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Contents

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.

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 product, program, or service that does not infringe any of the intellectual property rights of IBM may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you license to these patents. You can send inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 208 Harbor Drive, Stamford, CT, 06904-2501 USA.

Programming Information

This publication is intended to help the customer understand the performance of VM/ESA Release 2.2 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 2.2. See the IBM Programming Announcement for VM/ESA Release 2.2 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
BookManager
CICS
CICS/VSE
ECKD
Enterprise System/9000
ES/9000
IBM
IBMLink
Library Reader
OfficeVision
PR/SM
Processor Resource/Systems Manager
RACF
SAA
SQL/DS
Virtual Machine/Enterprise Systems Architecture
Virtual Machine/Extended Architecture
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

Sending Comments for This Publication

You can send comments using the form in the back of this publication, or use electronic mail or FAX. Please include:

- Publication title
- Order number (including suffix)
- Page, section title, or topic

The electronic addresses are:

Format	Address
Internet	pubrcf@gdlvme.vnet.ibm.com
IBM Mail	USIB2L8Z@IBMMAIL
IBMLink*	GDLVME(PUBRCF)
FAX	(607)752-2327

Your comments will help improve our information. Thank you.

Acknowledgements

The following people contributed to this report:

Glendale Programming Laboratory

Bill Bitner
Paul Duncan
Wes Ernsberger
Greg Gasper
Bill Guzior
Gary Hine
Greg Kudamik
Steve Lyons
Scott Marcotte
Virg Meredith
Linda Quinn
Jennifer Thrall

Washington Systems Center

Al Coley
Chuck Morse

Abstract

The *VM/ESA Release 2.2 Performance Report* summarizes the performance evaluation of VM/ESA Release 2.2. Measurements were obtained for the CMS-intensive, 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 2.2 performance evaluations conducted by the Glendale Programming Laboratory.

Discussion concentrates on the performance changes in VM/ESA Release 2.2, the performance effects of migrating from VM/ESA Release 2.1 to VM/ESA Release 2.2, and the performance of new functions provided in VM/ESA Release 2.2. These new functions include enhanced minidisk caching, scheduler enhancements, SPXTAPE, ISFC changes, and improved CMS block allocation for virtual disks in storage. A number of additional evaluations are also included.

Referenced Publications

The following VM/ESA publications are referred to in this report.

- *VM/ESA: Performance,* SC24-5642
- *VM/ESA: Planning and Administration,* SC24-5521
- *VM/ESA: SFS and CRR Planning, Administration, and Operation,* SC24-5649
- *VM/ESA: CP Command and Utility Reference,* SC24-5519

VM/ESA Release 1.0 Performance Report, ZZ05-0469, is classified as IBM Internal Use Only. The information contained in this document may be discussed with customers, but the document may not be given to customers. They may ask their IBM representative for access to the information contained in this publication.

The following performance reports are unclassified and may be ordered by customers.

- *VM/ESA Release 1.1 Performance Report,* GG66-3236
- *VM/ESA Release 2 Performance Report,* GG66-3245
- *VM/ESA Release 2.1 Performance Report,* GC24-5673-00
- *VM/ESA Release 2.2 Performance Report,* GC24-5673-01 (this report)

Summary of Key Findings

This report summarizes the performance evaluation of Virtual Machine/Enterprise Systems Architecture* (VM/ESA*) Release 2.2. Measurements were obtained for the CMS-intensive, VSE guest, and VMSES/E environments on various Enterprise System/9000* (ES/9000*) processors. This section summarizes the key findings. For further information on any given topic, refer to the page indicated in parentheses.

Performance Changes: VM/ESA Release 2.2 includes a number of performance enhancements (page 5). These include enhanced minidisk caching, scheduler improvements, the SPXTAPE command, and improved CMS block allocation for virtual disk in storage. Some changes have the potential to adversely affect performance (page 10). Lastly, a number of changes were made that benefit VM/ESA performance management (page 13).

Migration from VM/ESA Release 2.1: Benchmark measurements showed the following performance results for VM/ESA Release 2.2 relative to VM/ESA Release 2.1:

CMS-intensive Internal throughput rate (ITR) improved 2.0% to 3.7%, while external response time ranged from equivalence to a 14% improvement (page 22). Enhanced minidisk caching was the most important factor that resulted in these improvements.

VSE guest ITR ranged from equivalence to a 1.5% decrease for the PACEX8 I/O-intensive batch workload. Elapsed times were equivalent (page 48). ITR and response times were equivalent for the VSECICS transaction processing workload (page 58).

Since this environment cannot use minidisk caching in VM/ESA Release 2.1, minidisk caching was turned off in VM/ESA Release 2.2 for these regression comparisons.

VMSES/E Overall elapsed time improved by 11% to 22% for dedicated measurements (page 64). Much larger improvements are possible in environments with high processor utilization due to concurrent activity (page 73). In addition, VMSES/E now automates further 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 (page 64).

Migration from Other VM Releases: The performance measurement data in this report can be used in conjunction with similar data in the four previous VM/ESA performance reports to get a general understanding of the performance aspects of migrating from earlier VM releases to VM/ESA Release 1.5 (370 Feature) or VM/ESA Release 2.2 (page 77).

New Functions: Major enhancements have been made to minidisk caching in VM/ESA Release 2.2. These enhancements broaden MDC's scope of applicability, improve MDC performance, and provide better caching controls (page 5).

Key Findings

Guest operating systems can now benefit from minidisk caching. The degree of benefit will depend upon how frequently the system workload does DASD read I/Os, data reference patterns, and how much storage is available for the minidisk cache. For VSE guests, elapsed time reductions of up to 51% and ITR improvements of up to 7.2% were observed for the I/O-intensive PACEX8 batch workload (page 90). With PACEX8, the use of minidisk caching resulted in elapsed times that were substantially better than VSE native. Processor usage was higher, however. Response time reductions of up to 10% and ITR improvements of up to 0.8% were observed for the VSECICS transaction processing workload (page 107).

Most of the CMS-intensive performance improvement from VM/ESA Release 2.1 comes from enhanced MDC (page 22). The degree of improvement varies depending on how it is used (page 111). The minidisk cache implementation has changed significantly. Because of this, MDC usage guidelines should be consulted before migrating to VM/ESA Release 2.2 (page 87).

VM/ESA Release 2.2 includes a number of scheduler improvements. Excess processor resources are now distributed proportionally based on share settings. In addition, you can now set an upper bound to a given virtual machine′s processor share. These, and other scheduler changes, should result in more efficient and consistent use of system resources, easier tuning, and fewer scheduler problems (page 139).

SPXTAPE is a new CP command that provides greatly improved performance and function relative to SPTAPE. For the cases measured, SPXTAPE DUMP was about 10 times faster, while SPXTAPE LOAD was about 2 times faster. SPXTAPE also uses fewer tapes. For these same measurements, SPXTAPE used 33% to 42% fewer tapes relative to SPTAPE (page 150).

There were a number of changes to ISFC this release. These increased ISFC processor usage by about 12%. ISFC multiprocessor locking improvements resulted in an increase in maximum capacity. Throughput increases ranging from 8% to 19% were observed for the measured cases (page 158).

The block allocation algorithm used by CMS was improved for the case where the file resides in a virtual disk in storage. With this change, measurement results show that substitution of virtual disks in storage for T-disks can now result in equivalent or improved system performance, even when they are held long-term (page 162).

Additional Evaluations: CMS-intensive measurements taken on the 9121-941 with 1, 2, 3, and 4 processors online show ITR scaling properly as the number of processors is increased (page 170).

Measurement results are provided that illustrate the kinds of performance effects that can occur on VM systems that are highly processor-constrained (page 177).

CMS-intensive measurements showed equivalent to improved performance when migrating from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 1.5 (370 Feature) (page 183). Other measurements showed a 26% elapsed time improvement for applying service with VMSES/E on VM/ESA Release 1.5 (370 Feature) relative to using VMSES on VM/ESA Release 1.0 (page 188).

Measurement results show that system responsiveness can be significantly improved by increasing the CP scheduler's interactive bias (page 192).

VM/VTAM* 3.4.1 showed about a 0.5% decrease in ITR relative to VM/VTAM 3.4.0. This is because certain VTAM internal trace options are now mandatory. This was done to improve VTAM serviceability (page 196).

Changes That Affect Performance

This chapter contains descriptions of various changes to VM/ESA Release 2.2 that affect performance. This information is equivalent to the information on VM/ESA Release 2.2 performance changes found in Appendix E of *VM/ESA Performance*, with additional detail plus information that has become available since its publication.

Most of the changes are performance improvements and are listed under "Performance Improvements" on page 5. However, some have the potential to adversely affect performance. These are listed under "Performance Considerations" on page 10. 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 when discussing the measurement results. These results serve to further illustrate where these changes apply and how they may affect performance.

"Performance Management" on page 13 is the third section of this chapter. It discusses changes that affect VM/ESA performance management.

Performance Improvements

The following items improve the performance of VM/ESA:

- Enhanced Minidisk Cache
- Scheduler Share Capping and Proportional Distribution
- SPXTAPE Command
- IUCV Pathlength Reduction
- Global Lock Removal for ISFC I/O Buffer
- Asynchronous Data Mover Facility (ADMF)
- TRSOURCE Selectivity
- Clean Start
- CMS - Larger Preallocated Stack for DMSITS (SVC Interrupt Handler)
- CMS - Improved Block Allocation for Virtual Disk in Storage
- SFS - Improved Handling of Released File Blocks
- SFS - Revoke Performance
- VMSES/E

Enhanced Minidisk Cache

Major enhancements have been made to minidisk caching in VM/ESA Release 2.2. These enhancements broaden MDC′s scope of applicability, improve MDC performance, and provide better controls over the caching process.

Broader Applicability: Prior to VM/ESA Release 2.2, minidisk caching was subject to a number of restrictions:

- CMS 4K-formatted minidisks only
- access via diagnose or *BLOCKIO interfaces only
- the cache must reside in expanded storage
- the minidisk cannot be on shared DASD
- dedicated devices are not supported
- FBA minidisks must be on page boundaries

With VM/ESA Release 2.2, there are fewer restrictions. The last three of the above restrictions still apply and the minidisk must be on 3380, 3390, 9345, or FBA DASD. However:

- The minidisk can be in any format.
- In addition to the diagnose and *BLOCKIO interfaces, minidisk caching now also applies to DASD accesses that are done using SSCH, SIO, or SIOF.
- The cache can reside in real storage, expanded storage, or a combination of both.

By lifting these restrictions, the benefits of minidisk caching are now available to a much broader range of computing situations. Guest operating systems are a prime example. Other examples include CMS minidisks that are not 4K-formatted, applications that use SSCH, SIO, or SIOF to access minidisks, and systems that do not have expanded storage.

Improved Performance Characteristics: Even in situations where minidisk caching is already being used, enhanced minidisk caching can provide performance benefits. This is primarily due the following factors:

- The new minidisk cache reads and caches whole tracks instead of individual blocks. The entire track is read in with one I/O, resulting in improved efficiency and reduced average access times.

  When the track satisfies certain restrictions (primarily, all DASD records on the track have the same length and an integral number of these records fit into a 4KB page), the contents of that track are managed in the minidisk cache on a page-by-page basis. For example, infrequently referenced pages can be moved to expanded storage or discarded from the cache. Tracks that do not meet these restrictions are managed on a track basis. The whole track is either in main storage, expanded storage, or is removed from the cache based on average frequency of reference.

- When the minidisk cache is placed in real memory, performance is improved relative to having the cache in expanded storage.

Better Controls: Commands have been added that allow you to:

- set and query the size of the cache
- set and query cache settings for a real device or for a minidisk
- purge the cache of data from a specific real device or minidisk
- change a user's ability to insert data into the cache
- bias the arbiter for or against minidisk caching

For additional background information on enhanced minidisk caching, see *VM/ESA Planning and Administration*. See "Enhanced Minidisk Cache" on page 87 for usage guidelines. See "VSE Guest Environment" on page 90 and "CMS Environment" on page 111 for performance measurement results.

Scheduler Share Capping and Proportional Distribution

The CP scheduler has been improved in two significant ways in VM/ESA Release 2.2. First, surplus processing from users who are not using their entire share will be given out to other users in a manner that is proportional to their shares. In prior releases, surplus processing was given to the ready-to-run user having the highest share.

Second, an installation is now able to specify a limit on the amount of processing resource a user may receive. Three types of limits are supported: NOLIMIT, LIMITSOFT, and LIMITHARD.

- NOLIMIT (the default) means that the user will not be limited. This results in scheduling that is equivalent to prior VM/ESA releases.
- LIMITSOFT means that the user will not get more than its share if there is any other user that can use the surplus without exceeding its own limit. However, if no other user can use the processing time, the limit will be overridden to give the user more than its share rather than letting the processor time go to waste.

- LIMITHARD means that the user will not get more than its share, even if no one else can use the surplus.

See “Scheduler Share Capping and Proportional Distribution” on page 139 for performance measurement results.

SPXTAPE Command

SPXTAPE is a new implementation of the SPTAPE command. It upgrades and extends the SPTAPE function to be able to handle the large number of spool files used by VM/ESA customers.

In most cases, the elapsed time required to dump spool files is an order of magnitude less as compared to SPTAPE. The performance is good enough to make it practical to consider backing up all spool files on a nightly basis if desired.

The large reduction in elapsed time is primarily due to the following factors:

- Spool files are written to tape in 32K blocks.
- Tape I/O is overlapped with DASD I/O.
- Many spool file blocks are read at one time using the paging subsystem. This results in reduced DASD access time and overlap of DASD I/Os that are done to different spool volumes.
- Multiple tape drives are supported. These are used to eliminate tape mount delays and increase the amount of overlap between tape I/O and DASD I/O.

SPTAPE writes a tape mark between each backed up spool file. SPXTAPE writes the spool file data as one tape file consisting of 32K blocks. This reduces the number of tape volumes required to hold the spool files. Relative to SPTAPE, reductions ranging from 30% to 60% have been observed. The smaller the average spool file size, the larger the reduction.

See “SPXTAPE” on page 150 for performance measurement results.

IUCV Pathlength Reduction

IUCV and APPC/VM processor usage was reduced substantially in VM/ESA Release 1.1 and VM/ESA Release 2.1. Processor usage has been further reduced in VM/ESA 2.2. Instruction trace results show that the pathlength of a 1-byte asynchronous SEND was reduced by about 15%, while the pathlength of a 10KB asynchronous SEND was reduced by about 12%.

All of the improvements were in the base pathlength that is independent of message length. As a result, SENDs with large message lengths experience a smaller percentage improvement relative to SENDs with short message lengths.

Most of the improvement was in the area of IUCV external interrupt reflection. As a result, the best improvements are for asynchronous requests because an IUCV external interrupt is used to notify the requesting virtual machine when the request has completed. Instruction trace results indicate that the pathlength of a 1-byte synchronous SEND is reduced by about 9% (versus 15% for a 1-byte asynchronous SEND).

Global Lock Removal for ISFC I/O Buffer

Prior to VM/ESA Release 2.2, all ISFC activity was serialized by the ISFC global lock. With VM/ESA Release 2.2, each active link has a link lock associated with it and the I/O-related functions of ISFC are now serialized by this link lock instead of by the ISFC global lock. This change reduces contention for the ISFC global lock, thus improving responsiveness and increasing the maximum amount of message traffic that ISFC can handle when there are multiple active links. See “ISFC Changes” on page 158 for performance measurement results.

Asynchronous Data Mover Facility (ADMF)

VM/ESA′s support of the ADMF hardware feature provides an extension to the existing channel subsystem which is capable of off-loading page move activity onto the I/O processor, freeing the instruction processor for other work while the page movement is performed. No external control or intervention is necessary. ADMF is made available for qualified guest use provided they:

- Have VM/ESA Release 2.2 or above running natively on the hardware
- Have the Dynamic Relocation Facility (DRF) available
- Are a preferred guest (V=R, V=F)

TRSOURCE Selectivity

This enhancement provides conditional statements for traces of type “data” so that the user can determine what data (if any) needs to be collected when a trace point is executed. If, on a given trace point occurrence, no data is to be collected, no trace record is cut. This capability allows high frequency code paths to be traced with minimal impact to system performance.

Clean Start

A clean start IPLs the system without attempting to recover spool files and system data files that existed prior to system shutdown. A clean start will typically take less time than a cold start because cold start recovers the system data files.

CMS - Larger Preallocated Stack for DMSITS

This change reduces the number of dynamic storage requests made by the CMS SVC interrupt handler. This results in a reduction in processor requirements for a broad range of CMS functions.

CMS - Improved Block Allocation for Virtual Disk in Storage

In VM/ESA Release 2.1, block allocation was always done using a moving cursor algorithm. This method continues its scan for free blocks from the point where it left off, wrapping around when the end of the minidisk is reached. For normal minidisks on DASD, this algorithm is advantageous because it helps to keep a given file′s blocks near to each other (and often contiguous). This reduces the number of I/Os and DASD access time. However, this algorithm is not as well suited to virtual disks in storage.

With VM/ESA Release 2.2, if the minidisk is a virtual disk in storage, the block allocation algorithm is changed so as to scan for blocks starting at the first available block on the minidisk. In this way, blocks that become available as files are erased are more quickly reused. As a result, the virtual disk in storage

will tend to generate less paging activity and have fewer page frames associated with it.

See “Improved CMS Block Allocation for Virtual Disk in Storage” on page 162 for performance measurement results.

SFS - Improved Handling of Released File Blocks

When one or more file blocks are released (from erasing a file, for example), those blocks normally become immediately available for use by other SFS requests once the change is committed. However, when SFS is required to maintain a consistent image of that file, or the directory or file space that file resides in, the released blocks cannot be made available for reuse until that requirement goes away. For example, if some other user currently has that file open for read, that file′s blocks cannot be made available until that other user has closed the file. Other examples of this read consistency requirement include DIRCONTROL directories (which have ACCESS to RELEASE read consistency) and the DMSOPCAT interface.

Prior to VM/ESA Release 2.2, the way in which these blocks were managed could cause instances of high processor utilization in the SFS server in cases where very large numbers of deferred file blocks were involved. A change has been included in VM/ESA Release 2.2 that eliminates most of the processing time used to manage blocks that require deferred availability.

SFS - Revoke Performance

The processing associated with revoking authority from an SFS file or directory has been changed to reduce the likelihood of catalog I/O having to occur when there are still other users who have individual authorizations to that object.

VMSES/E

The performance of the VMFBLD function has been improved. This will generally result in a decrease in elapsed time required to build a nucleus. See “VMSES/E” on page 64 for performance measurement results.

The automation of more service processing in VMSES/E R2.2 eliminates certain manual tasks. Therefore, the overall time required to do these tasks will decrease. See “VMSES/E” on page 64 for a list of tasks that have been automated by VMSES/E R2.2.

Performance Considerations

These items warrant consideration since they have potential for a negative impact to performance.

- Minidisk Caching
- ISFC Pathlength
- Proportional Distribution of Excess Share
- STORBUF Default Changes
- Favoring of Short Transactions
- CP SNAPDUMP Storage Dumping Enhancement
- GCS Data Space Support

Minidisk Caching

The enhanced minidisk caching provided by VM/ESA Release 2.2 will normally result in improved performance. However, since it substantially different from the minidisk caching in previous releases, you should review how minidisk caching is being used on your system before migrating to VM/ESA Release 2.2. Refer to the MDC migration considerations in “Enhanced Minidisk Cache” on page 87.

ISFC Pathlength

ISFC pathlengths have increased by about 12% in VM/ESA Release 2.2. This will ordinarily have no significant effect on overall system performance. However, applications that make heavy use of ISFC may experience some decrease in performance. See “ISFC Changes” on page 158 for performance measurement results.

Proportional Distribution of Excess Share

In prior releases, surplus processing was given to the ready-to-run user having the highest share. This has been changed in VM/ESA Release 2.2. Surplus processing from users who are not using their entire share is now given out to other users in a manner that is proportional to their shares.

For most installations, this change will either have no significant effect or result in improved performance characteristics. However, there may be cases where an installation′s desired performance characteristics have an implicit dependency upon the old method of allocating excess share.

For example, consider a VM/ESA system where most of the users run at the default relative share setting of 100, an important server machine that does large amounts of processing has a relative share of 120, and there are several other virtual machines that have very large relative shares. Prior to VM/ESA Release 2.2, the server machine may have provided excellent performance, but only because it was preferentially receiving large amounts of unused share. With VM/ESA Release 2.2, that server machine′s allocation of the excess share can become much smaller as a result of the new proportional distribution method, possibly resulting in periods of unacceptable server performance.

Before migrating to VM/ESA Release 2.2, check your virtual machine share allocations for situations like this. If you find any such case, increase that virtual

machine's share allocation to more properly reflect that virtual machine's true processing requirements.

STORBUF Default Changes

The default STORBUF settings have been increased from 100%, 85%, 75% to 125%, 105%, 95%. If your system is currently running with the default settings, you can continue to run with the old defaults by issuing SET SRM STORBUF 100 85 75.

Experience has shown that most VM/ESA systems run best with some degree of storage overcommitment. The new defaults are a reasonable starting point for systems that do not use expanded storage for paging. Systems that do use expanded storage for paging often run best with even larger overcommitment. You can do this either by specifying larger STORBUF values or by using SET SRM XSTORE to tell the scheduler what percentage of expanded storage to include when determining the amount of available storage for dispatching purposes. See "System Scheduling Control Options" in *VM/ESA Performance* for more information.

Favoring of Short Transactions

One side effect of the scheduler changes that were made to implement the proportional distribution of excess share is that there now tends to be somewhat less favoring of short transactions over long-running transactions. This shows up as increased trivial response time and somewhat decreased non-trivial response time. This can be observed in the CMS-intensive measurement results shown in "CMS-Intensive" on page 22. This effect is generally small, but is more significant on small processors where processing time is a larger percentage of overall response time.

The SET SRM IABIAS command can be used, if desired, to increase the extent to which the CP scheduler favors short transactions over longer ones. Doing so can result in improved system responsiveness, as illustrated by the measurement results in "Effect of IABIAS on Response Times" on page 192.

CP SNAPDUMP Storage Dumping Enhancement

VM/ESA Release 2.2 provides a new SNAPDUMP command that can be used to generate a system dump, identical to a hard abend dump, without bringing the system down. When using this command, bear in mind that:

- All activity in the system is stopped while the dump is in progress.
- The elapsed time required to take the dump is similar to the elapsed time required to obtain a system abend dump.
- The elapsed time can be reduced by using the SET DUMP command to restrict the dump to just those areas that are required.

GCS Data Space Support

This change will allow applications running on GCS to obtain performance benefits by using VM Data Spaces through use of the existing CP macros. When running in an XC mode virtual machine, this support requires some additional processing by GCS. For example, the access registers must be saved, along with the general registers, whenever a GCS supervisor call is made. To avoid this additional processing, do not run GCS in an XC mode virtual machine unless you are running a GCS application that makes use of data spaces.

Performance Management

These changes affect the performance management of VM/ESA.

- Monitor Enhancements
- INDICATE Command
- Comparisons between Old and New Minidisk Cache
- MONVIEW
- Install Process for S-Disk Modules in Logical Segments
- Effects on Accounting Data
- Realtime Monitor VM/ESA Enhancements

Monitor Enhancements

A number of new monitor records and fields have been added. Some of the more significant changes are summarized below. For a complete list of changes, see the MONITOR LIST1403 file (on MAINT′s 194 disk) for VM/ESA Release 2.2.

- Scheduler Monitor Enhancements

  The monitor has been enhanced to provide data on the new maximum share feature of the scheduler. This includes the maximum share setting in user configuration and scheduler records, a system count of current users in the limit list, and the rate that users are being added to the limit list. A state counter for user in the limit list has been added to high frequency user sampling and system sampling.

  Other data related to scheduler features that are not new in this release have also been added.

  The amount of total storage considered available when making scheduler decisions.

  The sum of the working sets for users in the various dispatch classes.

  The percentage of expanded storage to use in available memory calculations as set by the SET SRM XSTORE command.

- Minidisk Cache Monitor Changes

  The new minidisk cache required that several changes be made to related monitor data. Several existing fields contain minidisk cache information. The source (control block field) for monitor fields has been changed to maintain compatibility where possible. However, some of the existing fields are no longer meaningful because of the different design of the new minidisk cache. These obsolete fields have been made reserved fields. In addition, some new information has been added to monitor:

  The system-wide setting for minidisk cache.

  Real storage usage by minidisk cache.

  Expanded storage usage by minidisk cache.

  Related settings for individual virtual machines (NOMDCFS and NOINSERT).

  Cache eligibility settings for each real device.

- Improved User State Sampling

The accuracy of high frequency state sampling for virtual machines has been improved. Previously, the “running” state was skewed low. In fact, on uni-processor machines, the percentage of time spent in the “running” state was shown as zero. When virtual machines are being sampled, it is CP (the monitor) that is running. While the skewing is less on n-ways as n increases, it is still skewed. This has been corrected in VM/ESA Release 2.2 by checking to see if a user virtual machine has been displaced by the monitor. If so, that virtual machine is marked as running.

- Other Monitor Additions

  Processor spin time information for formal spin locks.

  User domain information for the new “logon by” feature.

  Indication of Asynchronous Data Mover installation.

INDICATE Command

Two of the INDICATE commands have been extended to accommodate the scheduler and minidisk caching changes.

- Another line is added to the response from INDICATE LOAD to show the number of users who are currently in the limit list. The limit list is introduced in VM/ESA Release 2.2 by the new maximum share scheduler function. This list represents the subset of users on the dispatch list who are currently being prevented from running because they would exceed their maximum share setting.

  The response from INDICATE LOAD has been changed slightly to reflect the fact that the minidisk cache can reside in real storage as well as expanded storage. When minidisk caching is being done in both real and expanded storage, the numbers shown reflect the combined benefits of both caches.

  In VM/ESA Release 2.1, the MDC hit ratio is computed on a block basis. In VM/ESA Release 2.2, it is computed on an I/O basis.

- In the INDICATE QUEUES response, a user who is currently in the limit list will be designated as L0, L1, L2, or L3. The current Q0, Q1, Q2, and Q3 designations will be shown for users who are in the dispatch list and not on the limit list.

Comparisons between Old and New Minidisk Cache

As mentioned in the discussion of monitor changes, the new minidisk cache design means that some of MDC performance measures no longer apply, while others have a somewhat different meaning. Because of this, you should exercise caution when comparing the performance of minidisk caching on VM/ESA Release 2.2 with the performance of minidisk caching when the same system was running an earlier VM/ESA release.

The MDC hit ratio is especially unsuitable for comparing minidisk caching performance between VM/ESA Release 2.2 and a prior VM/ESA release.

- Because the enhanced minidisk cache lifts a number of restrictions, it is likely that a significant amount of data that was ineligible will now start participating in the minidisk cache. This is likely to affect the MDC hit ratio. If there is some constraint on the MDC size, this additional data may well cause the hit ratio to go down. At the same time, however, the number of

real I/Os that are avoided is likely to go up because these additional DASD areas now benefit from minidisk caching.

- In VM/ESA Release 2.1, the MDC hit ratio is computed on a block basis. In VM/ESA Release 2.2, it is computed on an I/O basis. This difference can sometimes result in a significant difference in the computed hit ratio.
- There is another important difference if you are looking at RTM data. In VM/ESA Release 2.1, the MDC hits are only divided by those DASD reads that are eligible for minidisk caching. In VM/ESA Release 2.2, the MDC hits are divided by *all* DASD reads (except for page, spool, and virtual disk in storage I/O). This can lead to MDC hit ratios that appear lower on VM/ESA Release 2.2 than were experienced on earlier releases, even though minidisk caching may actually be more effective.

To avoid these problems, look at I/Os avoided instead. I/Os avoided is a bottom-line measure of how effective MDC is at reducing DASD I/Os. Further, this measure is very similar[1] in meaning between VM/ESA Release 2.2 and prior VM/ESA releases. RTM VM/ESA provides I/Os avoided on a system basis. (For VM/ESA Release 2.1, look at MDC_IA on the SYSTEM screen. For VM/ESA Release 2.2, look at VIO_AVOID on the new MDCACHE screen.) VMPRF's DASD_BY_ACTIVITY report shows I/Os avoided on a device basis.

MONVIEW

The MONVIEW package is a set of tools that can assist you when looking at raw Virtual Machine/Extended Architecture* (VM/XA*) or VM/ESA monitor data. It accepts monitor data from tape or disk and creates a CMS file with a single record for each monitor record. Options exist to translate the header of monitor data for domain/record/timestamp.

MONVIEW is provided on an "as is" basis, and is installed as samples on the MAINT 193 disk.

Install Process for S-Disk Modules in Logical Segments

The real storage requirements of a CMS-intensive environment can often be reduced by placing the frequently executed S-disk modules into a logical segment so that one copy is shared by all users. This used to be done as an extra step following VM/ESA installation. With VM/ESA Release 2.2, this has now been integrated into the VM/ESA installation process. Two logical segments are used: one for modules that can run above the 16 meg line and one for modules than cannot. A discussion of how to manually create a logical segment for modules has been retained in Chapter 3 of *VM/ESA Performance* for reference by those who wish to customize this step.

---

[1] The difference is with read requests involving multiple blocks. Such requests are normally handled with one real I/O. In that case, VM/ESA Release 2.1 and VM/ESA Release 2.2 will both count one I/O avoided if all the blocks are found in the MDC. Occasionally, however, CP will have to handle that request by splitting it up into two or more separate real I/Os. Then, VM/ESA Release 2.1 will count one I/O avoided for each of these real I/Os that can be avoided because all of its blocks are in the MDC. However, VM/ESA Release 2.2 will count a total of one I/O avoided if all the blocks are in the MDC. In most cases, this difference should have a negligible effect on the overall count of I/Os avoided.

Effects on Accounting Data

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 2.2. 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. For most workloads, this should not change much as a result of the changes made in VM/ESA Release 2.2. Most CMS-intensive workloads are expected to experience little change in virtual processor time and a slight decrease in CP processor time. I/O-intensive environments that are set up to use the enhanced minidisk cache and were not using minidisk caching prior to VM/ESA Release 2.2 can experience larger decreases in total processor time (up to 6%).

Milliseconds of Virtual processor time (37-40) This is the virtual time charged to a user. As mentioned above, little change is expected for most workloads.

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 should see little change in most cases.

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 should see little change in most cases.

Realtime Monitor VM/ESA Enhancements

RTM VM/ESA 1.5.2 has been updated to include performance data for the new minidisk cache in VM/ESA Release 2.2. Most of this information is provided in two new screens — MDCACHE and MDCLOG. The remaining information is provided as updates to the existing SYSTEM and XTLOG screens.

The calculation of the minidisk cache hit ratio (MDHR) has been changed in two ways.

1. In VM/ESA Release 2.1, the MDC hit ratio is computed on a block basis. In VM/ESA Release 2.2, it is computed on an I/O basis. This difference is also the case for the MDC hit ratio reported by the INDICATE LOAD command.
2. In VM/ESA Release 2.1, the MDC hits are only divided by those DASD reads that are eligible for minidisk caching. In VM/ESA Release 2.2, the MDC hits are divided by *all* DASD reads (except for page, spool, and virtual disk in storage I/O). This difference only applies to the MDC hit ratio as reported by RTM. The MDC hit ratio reported by the INDICATE USER command continues to be MDC hits divided by MDC-eligible DASD reads.

Measurement Information

This chapter discusses the types of processors used for measurements in the report, the level of software used, the configuration details 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 were used for this report.

- 9021-941

  This processor was used with 1, 2, 3, and 4 processors online. Any processors that were not being used were varied offline.

- 9121-742

  This processor was used for the 9121-742 and 9121-621 measurements. To run as a 9121-621, the 9121-742 was physically partitioned.

- 9121-480

  This processor was used for the 9121-480 and 9121-320 measurements. To run as a 9121-320, one processor was varied offline.

- 9221-170
- 9221-120

Software Used

Unless otherwise noted, a pre-general-availability level of VM/ESA Release 2.2 was used for the measurements in this report. Not all of the VM/ESA Release 2.2 measurements in this report were made with the same level of code. As the product developed, newer code levels were made that supplanted the level that had been in use. In any evaluation section that compares VM/ESA Release 2.2 to itself, the same level of code was maintained. Keep this in mind when trying to compare results that are taken from different sections.

Other releases of VM were used in the report. VM/ESA Release 2.1 was at the GA+first-RSU level (General Availability, Recommended Service Upgrade Tape). The service that was part of VM/ESA Release 2.1 after the first RSU level and forwarded to the pre-general-availability code can account for some of the difference between VM/ESA Release 2.1 and VM/ESA Release 2.2.

See the appropriate workload section in Appendix A, “Workloads” on page 208 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:

Measurement Information

1. Workload: This specifies the name of the workload associated with the measurement. For more detail on the workload, see Appendix A, "Workloads" on page 208.
2. Hardware Configuration: This summarizes the hardware configuration and 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.

       Any real storage not defined for the specific measurement was configured as expanded storage and attached to an idle user.
     - Expanded: The amount of expanded storage used on the processor.
   - Tape: The type of tape drive and the tape's purpose.
   - DASD: The DASD configuration used during the measurement.

     The table indicates the type of DASD used during the measurement, type of control units that connect these volumes to the system, the number of paths between the processor and the DASD, and the distribution of the DASD volumes for PAGE, SPOOL, TDSK, USER, SERVER and SYSTEM. R or W next to the DASD counts means Read or Write caching enabled respectively.
   - Communications: The type of control unit, number of communication control units, number of lines per control unit, and the line speed.
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.
   - 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, share of the system resources scheduled, number of pages reserved, and any other options that were set.
4. Measurement Discussion: This contains an analysis of the performance data in the table and gives the overall performance findings.

5. Measurement Data: This contains the table of performance results. These data were obtained or derived from the tools listed in “Tools Description” on page 20.

   There are several cases where the same information is reported from two sources because the sources calculate the value in a slightly different manner. For 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 can directly count the command rate as it runs the commands in the scripts. RTM, on the other hand, reports the command (transaction) rate that is determined by the CP scheduler, which 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 is principally based on the TPNS command rate. Furthermore, some values in the table (like TOT INT ADJ) are normalized to the TPNS command rate in an effort to get the most accurate performance measures possible.

   There are instances in these tables where two variables are equal yet there appears a non-zero number for their difference or percent difference. This indicates that the variables are only equal when they were rounded off to the significant digits that appear in the table.

   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 and internal tools were used to evaluate the performance measurements. The programs used in the measurements are listed below.

RTM Real Time Monitor records and reports performance data for VM systems. The appropriate level of RTM VM/ESA was used for the VM/ESA Release 2.1 and VM/ESA Release 2.2 systems.

VMPRF VM Performance Reporting Facility is the VM monitor reduction program.

TPNS Teleprocessing Network Simulator is a terminal and network simulation tool.

TPNS Reduction Program Reduces the TPNS log data to provide performance, load, and response time information.

CICSPARS CICS* Performance Analysis Reporting System, provides CICS response time and transaction information.

EXPLORE** monitors and reports performance data for VSE systems.

FSTTAPE Reduces hardware monitor data for the 9121 processors.

Hardware Monitor Collects branch, event, and timing data.

MONFAST Collects and reports branch, event, and timing data on a 9221 processor.

REDFP Consolidates the QUERY FILEPOOL STATUS data.

Migration from VM/ESA Release 2.1

This chapter explores the performance effects of migrating from VM/ESA Release 2.1 to VM/ESA Release 2.2. The following environments were measured: CMS-intensive, VSE guest, and VMSES/E.

CMS-Intensive

VM/ESA Release 2.2 had improved internal throughput rates and had improved or equivalent response times for the CMS-intensive environments measured. The ITR improvements resulting from decreased processor use can be attributed to the following functional enhancements:

- Enhanced minidisk cache
- IUCV pathlength reduction
- CMS - larger preallocated stack for DMSITS (SVC interrupt handler)
- Virtual disk in storage block allocation improvement

For more information on these and other performance-related enhancements in VM/ESA Release 2.2, see “Changes That Affect Performance” on page 4.

The internal throughput rates and response times for these measurements are shown in Figure 1 and Figure 2 on page 23.

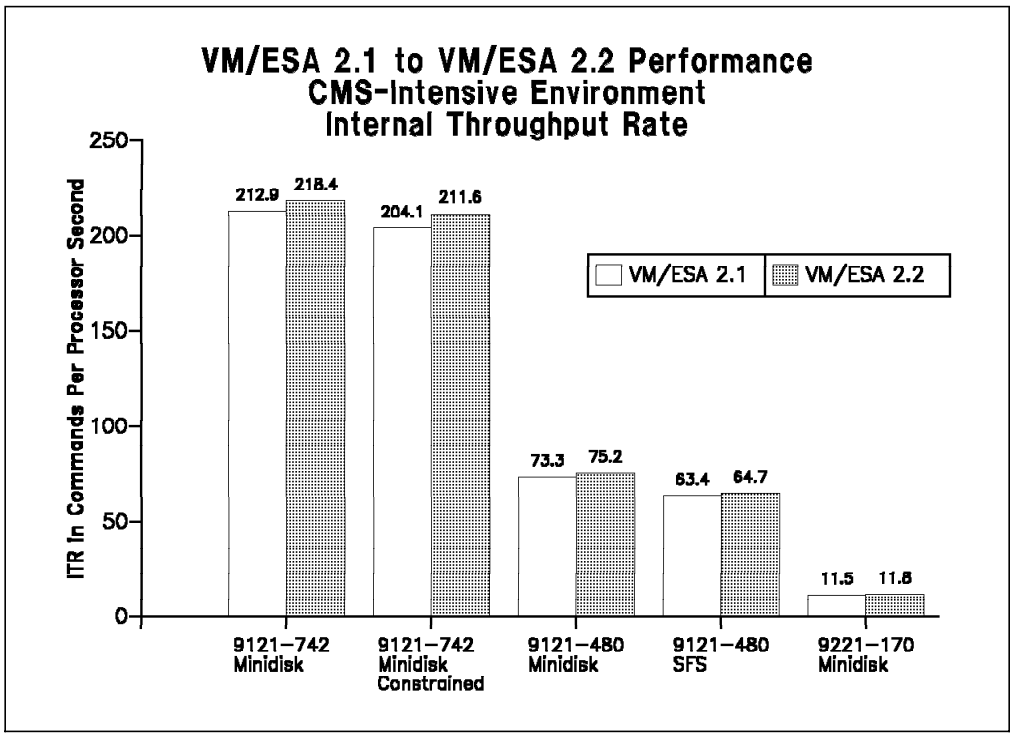

Figure 1. Internal throughput rate for the various CMS-intensive environments

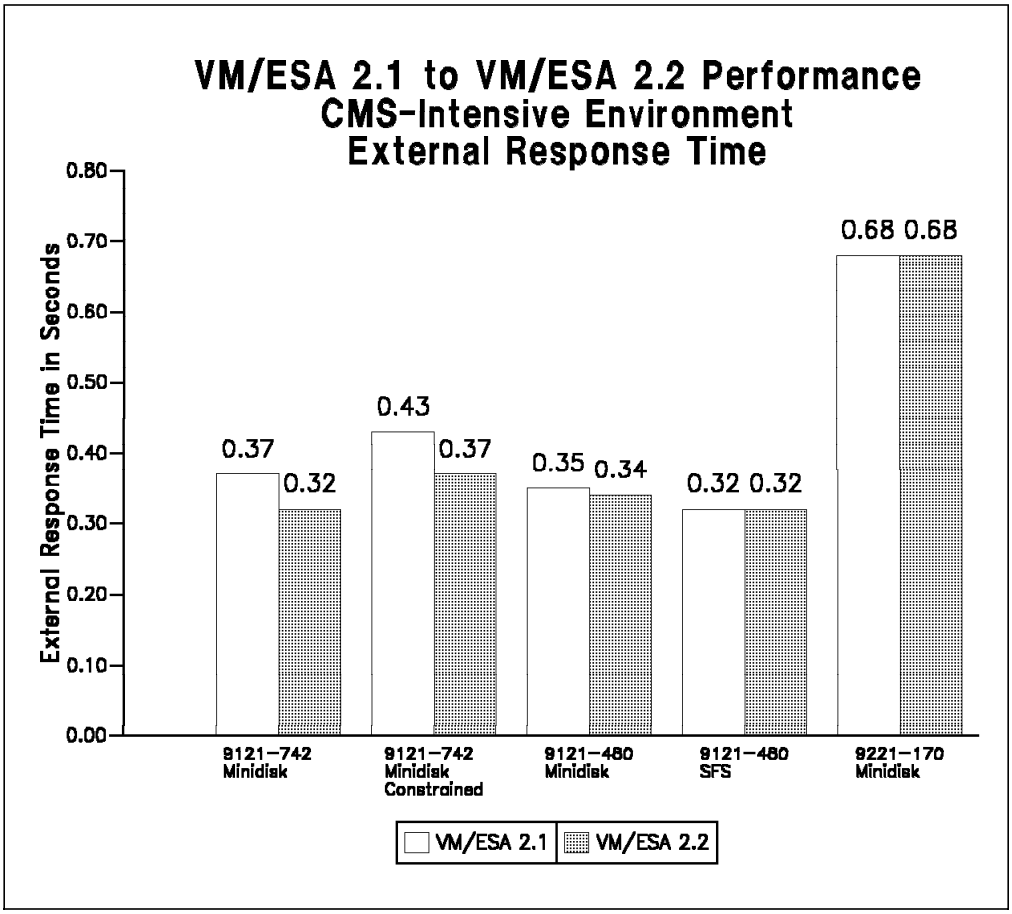

Figure 2. External response time for the various CMS-intensive environments

Migration: CMS-Intensive

9121-742 / Minidisk

Workload: FS8F0R

Hardware Configuration

Processor model:	9121-742
Processors used:	4
Storage, VM/ESA 2.1:	
Real:	1024MB
Expanded:	1024MB (64MB to 1024MB for MDC)
Storage, VM/ESA 2.2:	
Real:	1024MB (default MDC)
Expanded:	1024MB (BIAS 0.1)
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*User*	*Server*	*System*
3390-2	3990-3	4	6	7	7	32 R		2 R
3390-2	3990-2	4	16	6	6			

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	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	500	QUICKDSP ON
VSCSn	3	VSCS	64MB/XA	10000	1200	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	550	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	5500	Users	3MB/XC	100		

Measurement Discussion: Note that in the above configuration description, minidisk caching (MDC) is set up differently for the two releases. For VM/ESA Release 2.1, only expanded storage could be used for MDC. The arbiter (the code that controls how much expanded storage to use for paging versus MDC) was allowed to dynamically adjust based on system activity with a 64MB minimum size. This method had been determined to perform well for this environment.

For VM/ESA Release 2.2, the new MDC support was tuned differently based on experiments performed in this environment. (For a description of this enhanced MDC support refer to "Enhanced Minidisk Cache" on page 87.) In VM/ESA Release 2.2, there are two arbiters, one for expanded storage and one for real storage. These arbiters were again allowed to determine the proper amounts of storage to use for paging and MDC. However, it was determined that setting BIAS to 0.1 for expanded storage (using the SET MDCACHE command) improved overall system performance[2]. This setting biases the expanded storage arbiter against MDC (that is, uses less expanded storage for MDC than the default algorithm). See "Biasing the Expanded Storage Arbiter" on page 122 for further information on the BIAS option.

Because of the enhancements made to MDC, the MDC data in Table 1 on page 26 has changed. MDC READS (blks) and MDC READS (I/Os) have very different meanings and should not be compared. Similarly, MDC WRITES (blks) and MDC WRITES (I/Os) should not be compared. There is no equivalent to MDC MODS in VM/ESA Release 2.2. Also, the MDC HIT RATIO is computed differently for the two releases. Although, for this set of data, the hit ratios came out the same, a difference does not necessarily indicate an increase or decrease in MDC efficiency. For more discussion of MDC HIT RATIO, refer to "Comparisons between Old and New Minidisk Cache" on page 14. For specific definitions, refer to "Glossary of Performance Terms" on page 228.

The following table shows that VM/ESA Release 2.2 compared to VM/ESA Release 2.1 has improved its performance characteristics. The key indicators of external response time (AVG LAST(T)) and internal throughput rate (ITR(H)) both improved. The external response time improved by 13.5% and the internal throughput improved by 2.6%.

[2] A similar improvement could have been achieved by using the SET MDCACHE command or the RETAIN XSTORE command to specify a maximum expanded storage MDC size (something like 50MB would be suitable in this case). However, the BIAS option is the preferred method in most situations because it gives the system more flexibility. For example, at times when the demand for MDC is especially high and the demand for paging is especially low, use of the BIAS option allows the system to temporarily expand the expanded storage MDC size beyond 50MB.

Migration: CMS-Intensive

Table 1 (Page 1 of 3). Minidisk-only CMS-intensive migration from VM/ESA Release 2.1 on the 9121-742

Release Run ID	VM/ESA 2.1 S46E5505	VM/ESA 2.2 S47E550A	Difference	%Difference
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Response Time				
TRIV INT	0.102	0.108	0.006	5.88%
NONTRIV INT	0.348	0.339	-0.009	-2.59%
TOT INT	0.245	0.249	0.004	1.63%
TOT INT ADJ	0.245	0.236	-0.009	-3.61%
AVG FIRST (T)	0.281	0.240	-0.041	-14.57%
AVG LAST (T)	0.370	0.320	-0.050	-13.51%
Throughput				
AVG THINK (T)	26.11	26.09	-0.02	-0.06%
ETR	192.71	182.96	-9.75	-5.06%
ETR (T)	192.92	193.13	0.20	0.10%
ETR RATIO	0.999	0.947	-0.052	-5.16%
ITR (H)	212.95	218.41	5.46	2.57%
ITR	53.23	51.78	-1.45	-2.71%
EMUL ITR	81.20	77.56	-3.65	-4.49%
ITRR (H)	1.000	1.026	0.026	2.57%
ITRR	1.000	0.973	-0.027	-2.71%
Proc. Usage				
PBT/CMD (H)	18.784	18.314	-0.470	-2.50%
PBT/CMD	18.764	18.278	-0.486	-2.59%
CP/CMD (H)	6.865	6.480	-0.385	-5.61%
CP/CMD	6.479	6.058	-0.421	-6.50%
EMUL/CMD (H)	11.919	11.834	-0.085	-0.71%
EMUL/CMD	12.285	12.220	-0.065	-0.53%
Processor Util.				
TOTAL (H)	362.38	353.69	-8.70	-2.40%
TOTAL	362.00	353.00	-9.00	-2.49%
UTIL/PROC (H)	90.60	88.42	-2.17	-2.40%
UTIL/PROC	90.50	88.25	-2.25	-2.49%
TOTAL EMUL (H)	229.95	228.55	-1.40	-0.61%
TOTAL EMUL	237.00	236.00	-1.00	-0.42%
MASTER TOTAL (H)	92.77	90.81	-1.96	-2.11%
MASTER TOTAL	93.00	91.00	-2.00	-2.15%
MASTER EMUL (H)	39.69	39.08	-0.61	-1.53%
MASTER EMUL	41.00	41.00	0.00	0.00%
TVR(H)	1.58	1.55	-0.03	-1.80%
TVR	1.53	1.50	-0.03	-2.07%
Storage				
NUCLEUS SIZE (V)	2368KB	2556KB	188KB	7.94%
TRACE TABLE (V)	800KB	800KB	0KB	0.00%
WKSET (V)	72	72	0	0.00%
PGBLPGS	233K	233K	0K	0.00%
PGBLPGS/USER	42.4	42.4	0.0	0.00%
FREEPGS	15397	15848	451	2.93%
FREE UTIL	0.91	0.92	0.01	0.69%
SHRPGS	1777	1760	-17	-0.96%

Table 1 (Page 2 of 3). Minidisk-only CMS-intensive migration from VM/ESA Release 2.1 on the 9121-742

Release Run ID	VM/ESA 2.1 S46E5505	VM/ESA 2.2 S47E550A	Difference	%Difference
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Paging				
READS/SEC	580	604	24	4.14%
WRITES/SEC	430	448	18	4.19%
PAGE/CMD	5.235	5.447	0.212	4.05%
PAGE IO RATE (V)	154.200	158.400	4.200	2.72%
PAGE IO/CMD (V)	0.799	0.820	0.021	2.62%
XSTOR IN/SEC	843	844	1	0.12%
XSTOR OUT/SEC	1394	1413	19	1.36%
XSTOR/CMD	11.595	11.687	0.091	0.79%
FAST CLR/CMD	6.676	8.761	2.085	31.23%
Queues				
DISPATCH LIST	108.24	98.85	-9.38	-8.67%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	1807	1802	-5	-0.28%
VIO/CMD	9.366	9.331	-0.036	-0.38%
RIO RATE (V)	544	542	-2	-0.37%
RIO/CMD (V)	2.820	2.806	-0.013	-0.47%
NONPAGE RIO/CMD (V)	2.020	1.986	-0.034	-1.69%
DASD RESP TIME (V)	19.300	20.000	0.700	3.63%
MDC REAL SIZE (MB)	na	33.0		
MDC XSTOR SIZE (MB)	64.0	44.3	-17.7	-30.78%
MDC READS (blks)	1076	na		
MDC READS (I/Os)	na	550		
MDC WRITES (blks)	411	na		
MDC WRITES (I/Os)	na	27		
MDC MODS	313	na		
MDC AVOID	513	508	-5	-0.97%
MDC HIT RATIO	0.91	0.91		
PRIVOPs				
PRIVOP/CMD	20.295	20.563	0.268	1.32%
DIAG/CMD	25.397	25.217	-0.181	-0.71%
DIAG 04/CMD	0.932	0.931	-0.001	-0.07%
DIAG 08/CMD	0.738	0.739	0.001	0.14%
DIAG 0C/CMD	1.320	1.125	-0.195	-14.79%
DIAG 14/CMD	0.025	0.025	0.000	0.56%
DIAG 58/CMD	1.248	1.248	0.000	-0.01%
DIAG 98/CMD	0.314	0.344	0.030	9.58%
DIAG A4/CMD	3.628	3.567	-0.061	-1.69%
DIAG A8/CMD	2.830	2.829	-0.001	-0.04%
DIAG 214/CMD	13.218	13.261	0.044	0.33%
SIE/CMD	51.834	56.958	5.124	9.89%
SIE INTCPT/CMD	35.765	38.162	2.396	6.70%
FREE TOTL/CMD	62.201	44.919	-17.282	-27.78%

Migration: CMS-Intensive

Table 1 (Page 3 of 3). Minidisk-only CMS-intensive migration from VM/ESA Release 2.1 on the 9121-742

Release Run ID	VM/ESA 2.1 S46E5505	VM/ESA 2.2 S47E550A	Difference	%Difference
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
VTAM Machines				
WKSET (V)	4139	4144	5	0.12%
TOT CPU/CMD (V)	2.7472	2.7673	0.0201	0.73%
CP CPU/CMD (V)	1.2613	1.2312	-0.0301	-2.39%
VIRT CPU/CMD (V)	1.4859	1.5361	0.0502	3.38%
DIAG 98/CMD (V)	0.314	0.344	0.030	9.65%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

9121-742 / Minidisk / Storage-Constrained

Workload: FS8F0R

Hardware Configuration

Processor model:	9121-742
Processors used:	4
Storage, VM/ESA 2.1:	
Real:	320MB
Expanded:	64MB (All for MDC)
Storage, VM/ESA 2.2:	
Real:	384MB (Default MDC)
Expanded:	0MB
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-3	4	8	5	7	32 R		2 R
3390-2	3990-2	4	18	4	6			

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	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	500	QUICKDSP ON
VSCSn	3	VSCS	64MB/XA	10000	1200	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	550	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	5200	Users	3MB/XC	100		

Migration: CMS-Intensive

Measurement Discussion: As in the previous section, the above configuration is set up differently for the two releases. For VM/ESA Release 2.1, 64MB of expanded storage was carved out of the available real storage so that minidisk caching (MDC) could be used. With the enhanced MDC provided with VM/ESA Release 2.2, this was not necessary. MDC can now directly use real storage for caching. Note that the total available storage is the same for both releases. No adjustments to the VM/ESA Release 2.2 real storage arbiter settings were made, allowing it to determine the amount of storage to use for MDC. For a discussion of enhanced minidisk caching, refer to "Enhanced Minidisk Cache" on page 87.

The following table shows that VM/ESA Release 2.2 compared to VM/ESA Release 2.1 has improved its performance characteristics. The key indicators of external response time (AVG LAST(T)) and internal throughput rate (ITR(H)) both improved. The external response time improved by 13.7% and the internal throughput improved by 3.7%.

For this storage constrained environment, the enhanced MDC proved to be even more beneficial than the unconstrained environment in the previous section. It is more efficient to use real storage for MDC than to use expanded storage configured out of real storage. See "Minidisk Caching in Real Storage" on page 111 for further information.

Table 2 (Page 1 of 3). Minidisk-only storage constrained CMS-intensive migration from VM/ESA Release 2.1 on the 9121-742

Release **Run ID**	**VM/ESA 2.1** **S46E5202**	**VM/ESA 2.2** **S47E5201**	**Difference**	**%Difference**
Environment				
Real Storage	**320MB**	**384MB**		
Exp. Storage	**64MB**	**0MB**		
Users	**5200**	**5200**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Response Time				
TRIV INT	0.178	0.163	-0.015	-8.43%
NONTRIV INT	0.482	0.437	-0.045	-9.34%
TOT INT	0.369	0.339	-0.030	-8.13%
TOT INT ADJ	0.342	0.306	-0.036	-10.64%
AVG FIRST (T)	0.329	0.280	-0.049	-14.91%
AVG LAST (T)	0.430	0.371	-0.059	-13.71%
Throughput				
AVG THINK (T)	26.14	26.11	-0.03	-0.13%
ETR	169.11	164.68	-4.43	-2.62%
ETR (T)	182.24	182.46	0.22	0.12%
ETR RATIO	0.928	0.903	-0.025	-2.74%
ITR (H)	204.12	211.58	7.46	3.65%
ITR	47.41	47.81	0.40	0.84%
EMUL ITR	74.02	72.75	-1.27	-1.72%
ITRR (H)	1.000	1.037	0.037	3.65%
ITRR	1.000	1.008	0.008	0.84%
Proc. Usage				
PBT/CMD (H)	19.596	18.906	-0.691	-3.52%
PBT/CMD	19.589	18.908	-0.681	-3.48%
CP/CMD (H)	7.501	6.908	-0.593	-7.91%
CP/CMD	7.024	6.467	-0.557	-7.92%
EMUL/CMD (H)	12.095	11.998	-0.097	-0.80%
EMUL/CMD	12.566	12.441	-0.125	-0.99%

Table 2 (Page 2 of 3). Minidisk-only storage constrained CMS-intensive migration from VM/ESA Release 2.1 on the 9121-742

Release Run ID	VM/ESA 2.1 S46E5202	VM/ESA 2.2 S47E5201	Difference	%Difference
Environment				
Real Storage	**320MB**	**384MB**		
Exp. Storage	**64MB**	**0MB**		
Users	**5200**	**5200**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Processor Util.				
TOTAL (H)	357.13	344.96	-12.17	-3.41%
TOTAL	357.00	345.00	-12.00	-3.36%
UTIL/PROC (H)	89.28	86.24	-3.04	-3.41%
UTIL/PROC	89.25	86.25	-3.00	-3.36%
TOTAL EMUL (H)	220.42	218.91	-1.51	-0.68%
TOTAL EMUL	229.00	227.00	-2.00	-0.87%
MASTER TOTAL (H)	90.97	88.13	-2.84	-3.12%
MASTER TOTAL	91.00	88.00	-3.00	-3.30%
MASTER EMUL (H)	37.79	37.12	-0.67	-1.78%
MASTER EMUL	39.00	39.00	0.00	0.00%
TVR(H)	1.62	1.58	-0.04	-2.74%
TVR	1.56	1.52	-0.04	-2.51%
Storage				
NUCLEUS SIZE (V)	2368KB	2540KB	172KB	7.26%
TRACE TABLE (V)	800KB	800KB	0KB	0.00%
WKSET (V)	71	76	5	7.04%
PGBLPGS	55462	71684	16222	29.25%
PGBLPGS/USER	10.7	13.8	3.1	29.25%
FREEPGS	14536	14726	190	1.31%
FREE UTIL	0.92	0.92	0.01	0.61%
SHRPGS	1713	1724	11	0.64%
Paging				
READS/SEC	2050	1834	-216	-10.54%
WRITES/SEC	1646	1466	-180	-10.94%
PAGE/CMD	20.281	18.086	-2.195	-10.82%
PAGE IO RATE (V)	697.800	626.600	-71.200	-10.20%
PAGE IO/CMD (V)	3.829	3.434	-0.395	-10.31%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	6.722	8.906	2.184	32.49%
Queues				
DISPATCH LIST	115.55	107.41	-8.13	-7.04%
ELIGIBLE LIST	0.00	0.00	0.00	na

Migration: CMS-Intensive

Table 2 (Page 3 of 3). Minidisk-only storage constrained CMS-intensive migration from VM/ESA Release 2.1 on the 9121-742

Release Run ID	VM/ESA 2.1 S46E5202	VM/ESA 2.2 S47E5201	Difference	%Difference
Environment				
Real Storage	**320MB**	**384MB**		
Exp. Storage	**64MB**	**0MB**		
Users	**5200**	**5200**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
I/O				
VIO RATE	1715	1721	6	0.35%
VIO/CMD	9.410	9.432	0.022	0.23%
RIO RATE (V)	1077	1010	-67	-6.22%
RIO/CMD (V)	5.910	5.535	-0.374	-6.33%
NONPAGE RIO/CMD (V)	2.081	2.101	0.021	0.99%
DASD RESP TIME (V)	22.000	22.300	0.300	1.36%
MDC REAL SIZE (MB)	na	45.8		
MDC XSTOR SIZE (MB)	64.0	na		
MDC READS (blks)	1015	na		
MDC READS (I/Os)	na	529		
MDC WRITES (blks)	388	na		
MDC WRITES (I/Os)	na	25		
MDC MODS	295	na		
MDC AVOID	483	498	15	3.11%
MDC HIT RATIO	0.91	0.93		
PRIVOPs				
PRIVOP/CMD	20.944	21.196	0.252	1.20%
DIAG/CMD	25.446	25.457	0.011	0.04%
DIAG 04/CMD	0.983	1.000	0.017	1.72%
DIAG 08/CMD	0.736	0.737	0.001	0.19%
DIAG 0C/CMD	1.320	1.125	-0.194	-14.72%
DIAG 14/CMD	0.024	0.025	0.000	0.54%
DIAG 58/CMD	1.249	1.249	0.000	0.00%
DIAG 98/CMD	0.373	0.420	0.047	12.58%
DIAG A4/CMD	3.628	3.593	-0.035	-0.96%
DIAG A8/CMD	2.822	2.824	0.002	0.09%
DIAG 214/CMD	13.169	13.336	0.167	1.27%
SIE/CMD	53.824	54.806	0.982	1.82%
SIE INTCPT/CMD	37.138	37.268	0.130	0.35%
FREE TOTL/CMD	60.359	44.847	-15.511	-25.70%
VTAM Machines				
WKSET (V)	4008	4015	7	0.17%
TOT CPU/CMD (V)	2.9340	2.9473	0.0133	0.45%
CP CPU/CMD (V)	1.3467	1.3092	-0.0375	-2.78%
VIRT CPU/CMD (V)	1.5873	1.6381	0.0508	3.20%
DIAG 98/CMD (V)	0.373	0.420	0.047	12.65%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

9121-480 / Minidisk

Workload: FS8F0R

Hardware Configuration

Processor model:	9121-480
Processors used:	2
Storage, VM/ESA 2.1:	
Real:	224MB
Expanded:	32MB (all for MDC)
Storage, VM/ESA 2.2:	
Real:	256MB (default MDC)
Expanded:	0MB
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2	16 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*
3088-08	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XC	10000	512	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1860	Users	3MB/XC	100		

Migration: CMS-Intensive

Measurement Discussion: As in the previous section, storage used by the enhanced minidisk caching (MDC) for the VM/ESA Release 2.2 measurement was taken directly from real storage.

The following table shows that VM/ESA Release 2.2 compared to VM/ESA Release 2.1 has improved its performance characteristics. The key indicators of external response time (AVG LAST(T)) and internal throughput rate (ITR(H)) both improved. The external response time improved by 5.0% and the internal throughput improved by 2.6%.

Table 3 (Page 1 of 3). Minidisk-only CMS-intensive migration from VM/ESA Release 2.1 on the 9121-480

Release **Run ID**	**VM/ESA 2.1** **L26E186I**	**VM/ESA 2.2** **L27E186J**	**Difference**	**%Difference**
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.123	0.127	0.004	3.25%
NONTRIV INT	0.428	0.396	-0.032	-7.48%
TOT INT	0.325	0.305	-0.020	-6.15%
TOT INT ADJ	0.287	0.269	-0.018	-6.18%
AVG FIRST (T)	0.240	0.245	0.005	2.08%
AVG LAST (T)	0.354	0.336	-0.018	-4.95%
Throughput				
AVG THINK (T)	26.16	26.19	0.02	0.10%
ETR	57.71	57.66	-0.05	-0.09%
ETR (T)	65.40	65.36	-0.04	-0.06%
ETR RATIO	0.882	0.882	0.000	-0.03%
ITR (H)	73.27	75.15	1.87	2.56%
ITR	32.34	33.16	0.82	2.53%
EMUL ITR	48.01	48.09	0.09	0.18%
ITRR (H)	1.000	1.026	0.026	2.56%
ITRR	1.000	1.025	0.025	2.53%
Proc. Usage				
PBT/CMD (H)	27.296	26.615	-0.681	-2.49%
PBT/CMD	27.216	26.620	-0.596	-2.19%
CP/CMD (H)	9.477	8.849	-0.627	-6.62%
CP/CMD	8.868	8.261	-0.607	-6.84%
EMUL/CMD (H)	17.819	17.766	-0.053	-0.30%
EMUL/CMD	18.348	18.359	0.011	0.06%
Processor Util.				
TOTAL (H)	178.52	173.97	-4.56	-2.55%
TOTAL	178.00	174.00	-4.00	-2.25%
UTIL/PROC (H)	89.26	86.98	-2.28	-2.55%
UTIL/PROC	89.00	87.00	-2.00	-2.25%
TOTAL EMUL (H)	116.54	116.12	-0.42	-0.36%
TOTAL EMUL	120.00	120.00	0.00	0.00%
MASTER TOTAL (H)	89.03	87.50	-1.53	-1.72%
MASTER TOTAL	89.00	87.00	-2.00	-2.25%
MASTER EMUL (H)	51.86	64.55	12.69	24.47%
MASTER EMUL	54.00	67.00	13.00	24.07%
TVR(H)	1.53	1.50	-0.03	-2.20%
TVR	1.48	1.45	-0.03	-2.25%

Table 3 (Page 2 of 3). Minidisk-only CMS-intensive migration from VM/ESA Release 2.1 on the 9121-480

Release Run ID	VM/ESA 2.1 L26E186I	VM/ESA 2.2 L27E186J	Difference	%Difference
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Storage				
NUCLEUS SIZE (V)	2364KB	2556KB	192KB	8.12%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	82	84	2	2.44%
PGBLPGS	47502	55460	7958	16.75%
PGBLPGS/USER	25.5	29.8	4.3	16.75%
FREEPGS	5097	5282	185	3.63%
FREE UTIL	0.95	0.92	-0.03	-3.50%
SHRPGS	1230	1231	1	0.08%
Paging				
READS/SEC	612	629	17	2.78%
WRITES/SEC	432	436	4	0.93%
PAGE/CMD	15.963	16.293	0.331	2.07%
PAGE IO RATE (V)	171.800	176.700	4.900	2.85%
PAGE IO/CMD (V)	2.627	2.703	0.077	2.91%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	6.666	8.720	2.054	30.81%
Queues				
DISPATCH LIST	33.43	36.13	2.70	8.07%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	671	663	-8	-1.19%
VIO/CMD	10.259	10.143	-0.116	-1.13%
RIO RATE (V)	368	364	-4	-1.09%
RIO/CMD (V)	5.627	5.569	-0.058	-1.03%
NONPAGE RIO/CMD (V)	3.000	2.865	-0.134	-4.48%
DASD RESP TIME (V)	19.100	19.500	0.400	2.09%
MDC REAL SIZE (MB)	na	40.9		
MDC XSTOR SIZE (MB)	32.0	na		
MDC READS (blks)	364	na		
MDC READS (I/Os)	na	187		
MDC WRITES (blks)	139	na		
MDC WRITES (I/Os)	na	9.38		
MDC MODS	106	na		
MDC AVOID	173	176	3	1.73%
MDC HIT RATIO	0.91	0.94		

Migration: CMS-Intensive

Table 3 (Page 3 of 3). Minidisk-only CMS-intensive migration from VM/ESA Release 2.1 on the 9121-480

Release Run ID	VM/ESA 2.1 L26E186I	VM/ESA 2.2 L27E186J	Difference	%Difference
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
PRIVOPs				
PRIVOP/CMD	13.967	13.905	-0.062	-0.44%
DIAG/CMD	27.810	27.608	-0.202	-0.72%
DIAG 04/CMD	2.496	2.472	-0.024	-0.97%
DIAG 08/CMD	0.738	0.739	0.002	0.22%
DIAG 0C/CMD	1.319	1.126	-0.192	-14.57%
DIAG 14/CMD	0.025	0.025	0.000	0.17%
DIAG 58/CMD	1.247	1.249	0.002	0.13%
DIAG 98/CMD	1.209	1.163	-0.046	-3.81%
DIAG A4/CMD	3.646	3.585	-0.061	-1.68%
DIAG A8/CMD	2.831	2.834	0.003	0.10%
DIAG 214/CMD	13.159	13.266	0.107	0.82%
SIE/CMD	51.649	53.592	1.943	3.76%
SIE INTCPT/CMD	34.605	34.835	0.230	0.66%
FREE TOTL/CMD	67.092	49.538	-17.554	-26.16%
VTAM Machines				
WKSET (V)	525	508	-17	-3.24%
TOT CPU/CMD (V)	3.9669	3.8587	-0.1082	-2.73%
CP CPU/CMD (V)	1.5630	1.4619	-0.1011	-6.47%
VIRT CPU/CMD (V)	2.4039	2.3968	-0.0071	-0.30%
DIAG 98/CMD (V)	1.208	1.162	-0.046	-3.80%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

9121-480 / SFS

Workload: FS8FMAXR

Hardware Configuration

Processor model: 9121-480
Processors used: 2
Storage, VM/ESA 2.1:
  Real: 224MB
  Expanded: 32MB (all for MDC)
Storage, VM/ESA 2.2:
  Real: 256MB (default MDC)
  Expanded: 0MB
Tape: 3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	5	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2		16 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*
3088-08	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*	*SHARE*	*RESERVED*	*Other Options*
CRRSERV1	1	SFS	16MB/XC	100		
ROSERV1	1	SFS	32MB/XC	100		QUICKDSP ON
RWSERVn	2	SFS	64MB/XC	1500	1300	QUICKDSP ON
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	512	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1620	Users	3MB/XC	100		

Migration: CMS-Intensive

Measurement Discussion: Internal throughput (ITR(H)) improved by 2.0%, while external response time (AVG LAST(T)) was essentially unchanged. The ITR improvement was primarily due to the enhanced minidisk cache. A reduction in IUCV pathlength was also a signficant factor.

Minidisk caching (MDC) was set up based on the performance recommendations that apply to that VM/ESA release. For VM/ESA Release 2.1, MDC requires expanded storage, so the system was configured to have 32MB of expanded storage and this was dedicated to MDC. Earlier measurements using 32MB and 64MB of expanded storage indicated that 32MB provided good performance. For VM/ESA Release 2.2, the system was configured to have no expanded storage and MDC was run in real storage. The amount of real storage used by MDC was determined by the real storage arbiter (default).

The real storage arbiter appears to have made a good trade-off between using real storage for MDC versus using it for demand paging. It chose to use 67MB of real storage for MDC — more than twice the (fixed) amount used in the VM/ESA Release 2.1 measurement. Relative to VM/ESA Release 2.1, this resulted in a decrease of 0.28 non-paging I/Os (NONPAGE RIO/CMD(V)) and an increase of 0.13 paging I/Os (PAGE IO/CMD(V)), for a net decrease of 0.15 I/Os per command.

Table 4 (Page 1 of 3). SFS CMS-intensive migration from VM/ESA Release 2.1 on the 9121-480

Release **Run ID**	**VM/ESA 2.1** **L26S1621**	**VM/ESA 2.2** **L27S1622**	**Difference**	**%Difference**
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.113	0.119	0.006	5.31%
NONTRIV INT	0.413	0.412	-0.001	-0.24%
TOT INT	0.313	0.314	0.001	0.32%
TOT INT ADJ	0.274	0.275	0.001	0.32%
AVG FIRST (T)	0.217	0.230	0.013	5.98%
AVG LAST (T)	0.323	0.323	0.001	0.16%
Throughput				
AVG THINK (T)	26.20	26.19	0.00	-0.02%
ETR	49.84	49.92	0.08	0.16%
ETR (T)	56.88	56.97	0.09	0.16%
ETR RATIO	0.876	0.876	0.000	0.00%
ITR (H)	63.42	64.70	1.28	2.03%
ITR	27.80	28.35	0.55	1.99%
EMUL ITR	41.82	41.95	0.14	0.33%
ITRR (H)	1.000	1.020	0.020	2.03%
ITRR	1.000	1.020	0.020	1.99%
Proc. Usage				
PBT/CMD (H)	31.538	30.911	-0.626	-1.99%
PBT/CMD	31.471	30.895	-0.576	-1.83%
CP/CMD (H)	11.195	10.653	-0.542	-4.84%
CP/CMD	10.549	10.006	-0.543	-5.15%
EMUL/CMD (H)	20.343	20.259	-0.084	-0.41%
EMUL/CMD	20.922	20.889	-0.033	-0.16%

Table 4 (Page 2 of 3). SFS CMS-intensive migration from VM/ESA Release 2.1 on the 9121-480

Release Run ID	VM/ESA 2.1 L26S1621	VM/ESA 2.2 L27S1622	Difference	%Difference
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Processor Util.				
TOTAL (H)	179.38	176.09	-3.28	-1.83%
TOTAL	179.00	176.00	-3.00	-1.68%
UTIL/PROC (H)	89.69	88.05	-1.64	-1.83%
UTIL/PROC	89.50	88.00	-1.50	-1.68%
TOTAL EMUL (H)	115.70	115.41	-0.30	-0.26%
TOTAL EMUL	119.00	119.00	0.00	0.00%
MASTER TOTAL (H)	89.52	87.82	-1.70	-1.90%
MASTER TOTAL	90.00	88.00	-2.00	-2.22%
MASTER EMUL (H)	52.36	51.99	-0.38	-0.72%
MASTER EMUL	54.00	54.00	0.00	0.00%
TVR(H)	1.55	1.53	-0.02	-1.58%
TVR	1.50	1.48	-0.03	-1.68%
Storage				
NUCLEUS SIZE (V)	2364KB	2556KB	192KB	8.12%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	79	81	2	2.53%
PGBLPGS	48294	55721	7427	15.38%
PGBLPGS/USER	29.8	34.4	4.6	15.38%
FREEPGS	4581	4816	235	5.13%
FREE UTIL	0.95	0.96	0.01	0.72%
SHRPGS	1483	1454	-29	-1.96%
Paging				
READS/SEC	517	527	10	1.93%
WRITES/SEC	358	362	4	1.12%
PAGE/CMD	15.384	15.606	0.222	1.44%
PAGE IO RATE (V)	130.900	138.500	7.600	5.81%
PAGE IO/CMD (V)	2.301	2.431	0.130	5.64%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	6.980	8.444	1.464	20.97%
Queues				
DISPATCH LIST	33.81	36.26	2.46	7.27%
ELIGIBLE LIST	0.00	0.00	0.00	na

Migration: CMS-Intensive

Table 4 (Page 3 of 3). SFS CMS-intensive migration from VM/ESA Release 2.1 on the 9121-480

Release Run ID	VM/ESA 2.1 L26S1621	VM/ESA 2.2 L27S1622	Difference	%Difference
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
I/O				
VIO RATE	588	578	-10	-1.70%
VIO/CMD	10.338	10.146	-0.192	-1.86%
RIO RATE (V)	340	332	-8	-2.35%
RIO/CMD (V)	5.978	5.828	-0.150	-2.51%
NONPAGE RIO/CMD (V)	3.676	3.397	-0.280	-7.61%
DASD RESP TIME (V)	17.600	18.300	0.700	3.98%
MDC REAL SIZE (MB)	na	67.0		
MDC XSTOR SIZE (MB)	32.0	na		
MDC READS (blks)	274	na		
MDC READS (I/Os)	na	150		
MDC WRITES (blks)	106	na		
MDC WRITES (I/Os)	na	15		
MDC MODS	59	na		
MDC AVOID	121	126	5	4.13%
MDC HIT RATIO	0.85	0.83		
PRIVOPs				
PRIVOP/CMD	20.669	20.404	-0.265	-1.28%
DIAG/CMD	25.772	25.686	-0.086	-0.33%
DIAG 04/CMD	2.682	2.727	0.045	1.68%
DIAG 08/CMD	0.737	0.735	-0.002	-0.25%
DIAG 0C/CMD	1.158	1.147	-0.012	-0.99%
DIAG 14/CMD	0.025	0.024	0.000	-0.38%
DIAG 58/CMD	1.248	1.248	0.000	-0.03%
DIAG 98/CMD	1.437	1.315	-0.123	-8.54%
DIAG A4/CMD	2.022	1.984	-0.038	-1.90%
DIAG A8/CMD	2.617	2.606	-0.012	-0.44%
DIAG 214/CMD	12.712	12.761	0.049	0.39%
SIE/CMD	59.268	60.579	1.311	2.21%
SIE INTCPT/CMD	41.488	41.800	0.312	0.75%
FREE TOTL/CMD	75.215	52.645	-22.570	-30.01%
VTAM Machines				
WKSET (V)	504	504	0	0.00%
TOT CPU/CMD (V)	4.1610	4.0180	-0.1430	-3.44%
CP CPU/CMD (V)	1.6312	1.5116	-0.1196	-7.33%
VIRT CPU/CMD (V)	2.5298	2.5063	-0.0235	-0.93%
DIAG 98/CMD (V)	1.437	1.315	-0.122	-8.50%
SFS Servers				
WKSET (V)	3194	3164	-30	-0.94%
TOT CPU/CMD (V)	3.3405	3.1597	-0.1808	-5.41%
CP CPU/CMD (V)	1.5824	1.4141	-0.1683	-10.64%
VIRT CPU/CMD (V)	1.7582	1.7457	-0.0125	-0.71%
FP REQ/CMD(Q)	1.136	1.117	-0.019	-1.67%
IO/CMD (Q)	1.681	1.577	-0.104	-6.19%
IO TIME/CMD (Q)	0.020	0.021	0.001	5.00%
SFS TIME/CMD (Q)	0.026	0.026	0.000	0.00%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Query Filepool Counters, Unmarked=RTM

The SFS counts and timings in the following two tables are provided to supplement the information provided above. 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. A description of the QUERY FILEPOOL STATUS output can be found in *SFS and CRR Planning, Administration, and Operation*.

Table 5 consists of counts and timings that are 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 measurement interval. Counts and timings that 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.

Table 5 (Page 1 of 2). SFS CMS-intensive migration from VM/ESA Release 2.1 on the 9121-480

Release **Run ID**	**VM/ESA 2.1** **L26S1621**	**VM/ESA 2.2** **L27S1622**	**Difference**	**%Difference**
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Close File Requests	0.3614	0.3624	0.0010	0.28%
Commit Requests	0.0163	0.0163	0.0000	0.00%
Connect Requests	0.0078	0.0078	0.0000	0.00%
Delete File Requests	0.0736	0.0736	0.0000	0.00%
Lock Requests	0.0248	0.0247	-0.0001	-0.40%
Open File New Requests	0.0033	0.0033	0.0000	0.00%
Open File Read Requests	0.2159	0.2168	0.0009	0.42%
Open File Replace Requests	0.1210	0.1211	0.0001	0.08%
Open File Write Requests	0.0212	0.0212	0.0000	0.00%
Query File Pool Requests	0.0000	0.0000	0.0000	na
Query User Space Requests	0.0213	0.0213	0.0000	0.00%
Read File Requests	0.1607	0.1457	-0.0150	-9.33%
Refresh Directory Requests	0.0228	0.0228	0.0000	0.00%
Rename Requests	0.0049	0.0049	0.0000	0.00%
Unlock Requests	0.0247	0.0246	-0.0001	-0.40%
Write File Requests	0.0557	0.0508	-0.0049	-8.80%
Total File Pool Requests	1.1356	1.1172	-0.0184	-1.62%
File Pool Request Service Time	26.2170	26.1436	-0.0734	-0.28%
Local File Pool Requests	1.1356	1.1172	-0.0184	-1.62%
Begin LUWs	0.4417	0.4421	0.0004	0.09%
Agent Holding Time (msec)	96.8461	85.1939	-11.6522	-12.03%
SAC Calls	5.4532	5.4368	-0.0164	-0.30%
Catalog Lock Conflicts	0.0007	0.0003	-0.0004	-57.14%
Total Lock Conflicts	0.0007	0.0003	-0.0004	-57.14%
Lock Wait Time (msec)	0.0082	0.0037	-0.0045	-54.88%

Migration: CMS-Intensive

Table 5 (Page 2 of 2). SFS CMS-intensive migration from VM/ESA Release 2.1 on the 9121-480

Release Run ID	VM/ESA 2.1 L26S1621	VM/ESA 2.2 L27S1622	Difference	%Difference
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
File Blocks Read	0.9708	0.9048	-0.0660	-6.80%
File Blocks Written	0.5238	0.5028	-0.0210	-4.01%
Catalog Blocks Read	0.4962	0.4954	-0.0008	-0.16%
Catalog Blocks Written	0.2565	0.2559	-0.0006	-0.23%
Control Minidisk Blocks Written	0.0497	0.0497	0.0000	0.00%
Log Blocks Written	0.4584	0.4590	0.0006	0.13%
Total DASD Block Transfers	2.7554	2.6676	-0.0878	-3.19%
BIO Requests to Read File Blocks	0.4083	0.3893	-0.0190	-4.65%
BIO Requests to Write File Blocks	0.1864	0.1793	-0.0071	-3.81%
BIO Requests to Read Catalog Blks	0.4962	0.4954	-0.0008	-0.16%
BIO Requests to Write Catalog Blks	0.2077	0.2068	-0.0009	-0.43%
BIO Requests to Write Ctl Mdsk Blks	0.0020	0.0020	0.0000	0.00%
BIO Requests to Write Log Blocks	0.4000	0.4000	0.0000	0.00%
Total BIO Requests	1.7006	1.6730	-0.0276	-1.62%
Total BIO Request Time (msec)	20.2410	21.0825	0.8415	4.16%
I/O Requests to Read File Blocks	0.3608	0.2664	-0.0944	-26.16%
I/O Requests to Write File Blocks	0.2052	0.1968	-0.0084	-4.09%
I/O Requests to Read Catalog Blocks	0.4962	0.4954	-0.0008	-0.16%
I/O Requests to Write Catalog Blocks	0.2149	0.2138	-0.0011	-0.51%
I/O Requests to Write Ctl Mdsk Blks	0.0039	0.0038	-0.0001	-2.56%
I/O Requests to Write Log Blocks	0.4004	0.4003	-0.0001	-0.02%
Total I/O Requests	1.6813	1.5766	-0.1047	-6.23%
Get Logname Requests	0.0033	0.0033	0.0000	0.00%
Get LUWID Requests	0.0033	0.0033	0.0000	0.00%
Total CRR Requests	0.0065	0.0065	0.0000	0.00%
CRR Request Service Time (msec)	0.0772	0.0752	-0.0020	-2.59%

Note: Query Filepool Counters — normalized by command

The 26% decrease in *I/O Requests to Read File Blocks* results from a change in how CP counts I/O requests produced by multiblock block I/O read requests when minidisk caching is present. It does not reflect an actual change in the number of real I/Os required. The actual decrease was similar to the 4.65% decrease shown for *BIO Requests to Read File Blocks*.

Table 6 consists of derived relationships that were calculated from a combination of two or more individual counts or timings. See the glossary for definitions of these derived values.

Table 6. SFS CMS-intensive migration from VM/ESA Release 2.1 on the 9121-480

Release **Run ID**	**VM/ESA 2.1** **L26S1621**	**VM/ESA 2.2** **L27S1622**	**Difference**	**%Difference**
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Agents Held	5.5	4.9	-0.7	-11.89%
Agents In-call	1.5	1.5	0.0	-0.12%
Avg LUW Time (msec)	219.3	192.7	-26.6	-12.11%
Avg File Pool Request Time (msec)	23.1	23.4	0.3	1.36%
Avg Lock Wait Time (msec)	11.7	12.3	0.6	5.28%
SAC Calls / FP Request	4.80	4.87	0.06	1.34%
Deadlocks (delta)	0	0	0	na
Rollbacks Due to Deadlock (delta)	0	0	0	na
Rollback Requests (delta)	0	0	0	na
LUW Rollbacks (delta)	0	0	0	na
Checkpoints Taken (delta)	32	32	0	0.00%
Checkpoint Duration (sec)	1.7	1.6	-0.1	-7.09%
Seconds Between Checkpoints	60.2	60.2	0.0	0.00%
Checkpoint Utilization	2.9	2.7	-0.2	-7.09%
BIO Request Time (msec)	11.90	12.60	0.70	5.88%
Blocking Factor (Blocks/BIO)	1.62	1.59	-0.03	-1.59%
Chaining Factor (Blocks/IO)	1.64	1.69	0.05	3.24%

Note: Query Filepool Counters — derived results

9221-170 / Minidisk

Workload: FS8F0R

Hardware Configuration

Processor model: 9221-170
Processors used: 1
Storage, VM/ESA 2.1:
  Real: 48MB
  Expanded: 16MB (all for MDC)
Storage, VM/ESA 2.2:
  Real: 64MB (default MDC)
  Expanded: 0MB
Tape: 3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	1	16	6	6			
3390-2	3990-3	1		2	2	8 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	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	350	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	300	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	300	Users	3MB/XC	100		

Measurement Discussion: The following table shows that VM/ESA Release 2.2 compared to VM/ESA Release 2.1 has improved its performance characteristics. The external response time (AVG LAST(T)) increased by 0.7% and the internal throughput rate (ITR(H)) improved by 2.6%.

Also worth noting is the difference in internal trivial response time (TRIV INT). As a side effect of the scheduler enhancements put into VM/ESA Release 2.2, the default favoring of short requests was reduced. This small processor

environment is more sensitive to this change because of the fact that a larger percentage of command response time is due to processor usage and processor contention. The amount by which the scheduler favors short requests can be tuned through the use of the SET SRM IABIAS command. Refer to "Effect of IABIAS on Response Times" on page 192 for further discussion on this subject. For these measurements, IABIAS was defaulted.

As in the previous section, storage used by the enhanced minidisk caching (MDC) for the VM/ESA Release 2.2 measurement was taken directly from real storage.

Table 7 (Page 1 of 3). Minidisk-only CMS-intensive migration from VM/ESA Release 2.1 on a 9221-170

Release **Run ID**	**VM/ESA 2.1** **H16E0302**	**VM/ESA 2.2** **H17E0303**	**Difference**	**%Difference**
Environment				
Real Storage	**48MB**	**64MB**		
Exp. Storage	**16MB**	**0MB**		
Users	**300**	**300**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**1**	**1**		
Response Time				
TRIV INT	0.149	0.229	0.080	53.69%
NONTRIV INT	0.970	0.985	0.015	1.55%
TOT INT	0.728	0.755	0.027	3.71%
TOT INT ADJ	0.626	0.644	0.018	2.80%
AVG FIRST (T)	0.307	0.391	0.084	27.36%
AVG LAST (T)	0.679	0.684	0.005	0.74%
Throughput				
AVG THINK (T)	28.11	28.06	-0.05	-0.18%
ETR	8.96	8.89	-0.07	-0.78%
ETR (T)	10.42	10.43	0.01	0.10%
ETR RATIO	0.860	0.853	-0.008	-0.88%
ITR (H)	11.46	11.76	0.30	2.62%
ITR	9.86	10.03	0.17	1.74%
EMUL ITR	15.19	15.05	-0.14	-0.89%
ITRR (H)	1.000	1.026	0.026	2.62%
ITRR	1.000	1.017	0.017	1.74%
Proc. Usage				
PBT/CMD (H)	87.264	85.033	-2.231	-2.56%
PBT/CMD	87.372	85.366	-2.006	-2.30%
CP/CMD (H)	36.158	33.958	-2.199	-6.08%
CP/CMD	30.724	28.775	-1.949	-6.34%
EMUL/CMD (H)	51.106	51.075	-0.031	-0.06%
EMUL/CMD	56.647	56.591	-0.057	-0.10%
Processor Util.				
TOTAL (H)	90.89	88.65	-2.23	-2.46%
TOTAL	91.00	89.00	-2.00	-2.20%
TOTAL EMUL (H)	53.23	53.25	0.02	0.04%
TOTAL EMUL	59.00	59.00	0.00	0.00%
TVR(H)	1.71	1.66	-0.04	-2.50%
TVR	1.54	1.51	-0.03	-2.20%

Migration: CMS-Intensive

Table 7 (Page 2 of 3). Minidisk-only CMS-intensive migration from VM/ESA Release 2.1 on a 9221-170

Release Run ID	VM/ESA 2.1 H16E0302	VM/ESA 2.2 H17E0303	Difference	%Difference
Environment				
Real Storage	**48MB**	**64MB**		
Exp. Storage	**16MB**	**0MB**		
Users	**300**	**300**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**1**	**1**		
Storage				
NUCLEUS SIZE (V)	2368KB	2556KB	188KB	7.94%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
WKSET (V)	85	86	1	1.18%
PGBLPGS	9482	13367	3885	40.97%
PGBLPGS/USER	31.6	44.6	13.0	40.97%
FREEPGS	917	936	19	2.07%
FREE UTIL	0.89	0.90	0.01	0.78%
SHRPGS	1074	1095	21	1.96%
Paging				
READS/SEC	83	84	1	1.20%
WRITES/SEC	69	69	0	0.00%
PAGE/CMD	14.594	14.675	0.081	0.56%
PAGE IO RATE (V)	26.000	26.300	0.300	1.15%
PAGE IO/CMD (V)	2.496	2.523	0.026	1.05%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	6.625	8.537	1.912	28.86%
Queues				
DISPATCH LIST	10.88	11.86	0.98	9.03%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	118	117	-1	-0.85%
VIO/CMD	11.329	11.222	-0.107	-0.95%
RIO RATE (V)	69	70	1	1.45%
RIO/CMD (V)	6.625	6.714	0.089	1.35%
NONPAGE RIO/CMD (V)	4.129	4.192	0.063	1.53%
DASD RESP TIME (V)	23.000	23.800	0.800	3.48%
MDC REAL SIZE (MB)	na	11.1		
MDC XSTOR SIZE (MB)	15.9	na		
MDC READS (blks)	58	na		
MDC READS (I/Os)	na	28		
MDC WRITES (blks)	22	na		
MDC WRITES (I/Os)	na	1.40		
MDC MODS	16	na		
MDC AVOID	27	26	-1	-3.70%
MDC HIT RATIO	0.91	0.88		

Table 7 (Page 3 of 3). Minidisk-only CMS-intensive migration from VM/ESA Release 2.1 on a 9221-170

Release Run ID	VM/ESA 2.1 H16E0302	VM/ESA 2.2 H17E0303	Difference	%Difference
Environment				
Real Storage	**48MB**	**64MB**		
Exp. Storage	**16MB**	**0MB**		
Users	**300**	**300**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**1**	**1**		
PRIVOPs				
PRIVOP/CMD	14.571	14.521	-0.050	-0.35%
DIAG/CMD	32.686	32.475	-0.210	-0.64%
DIAG 04/CMD	6.226	6.228	0.001	0.02%
DIAG 08/CMD	0.735	0.735	0.000	0.02%
DIAG 0C/CMD	1.323	1.128	-0.195	-14.73%
DIAG 14/CMD	0.024	0.024	0.000	-0.13%
DIAG 58/CMD	1.252	1.251	-0.002	-0.13%
DIAG 98/CMD	2.301	2.283	-0.018	-0.78%
DIAG A4/CMD	3.647	3.582	-0.065	-1.78%
DIAG A8/CMD	2.832	2.840	0.008	0.28%
DIAG 214/CMD	13.195	13.254	0.059	0.45%
SIE/CMD	61.160	63.209	2.049	3.35%
SIE INTCPT/CMD	44.647	44.878	0.232	0.52%
FREE TOTL/CMD	76.426	58.509	-17.917	-23.44%
VTAM Machines				
WKSET (V)	282	289	7	2.48%
TOT CPU/CMD (V)	18.0024	17.7446	-0.2578	-1.43%
CP CPU/CMD (V)	7.3076	6.9806	-0.3270	-4.47%
VIRT CPU/CMD (V)	10.6947	10.7640	0.0693	0.65%
DIAG 98/CMD (V)	2.301	2.283	-0.018	-0.77%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

VSE/ESA Guest

In the following two sections, VSE/ESA* 1.3.2 guest performance measurement results are presented and discussed for the PACEX8 batch workload and the VSECICS transaction processing workload. These sections compare VSE/ESA guest performance on VM/ESA Release 2.2 to VM/ESA Release 2.1 and to VSE/ESA native.

9121-320 / PACEX8

This section examines VSE/ESA 1.3.2 guest performance of VM/ESA Release 2.2 compared to VM/ESA Release 2.1 running the PACEX8 workload. PACEX8 is batch-only and is characterized by heavy I/O. See Appendix A, “Workloads” on page 208 for a detailed description of the workload.

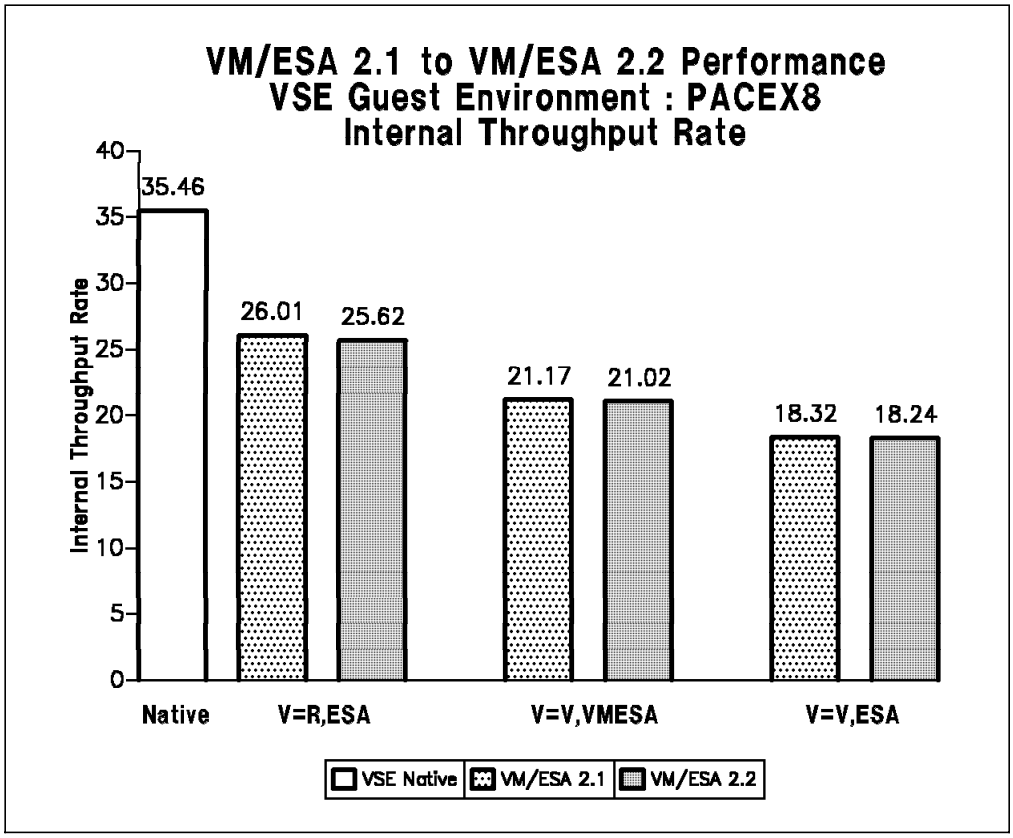

Figure 3. VSE guest migration from VM/ESA Release 2.1. PACEX8 workload on a single VSE/ESA guest of VM/ESA Release 2.1 and VM/ESA Release 2.2 on the 9121-320 processor.

The VSE PACEX8 workload was run as a guest of VM/ESA Release 2.1 and VM/ESA Release 2.2 in three modes. The internal throughput rate (ITR(H)) decreased slightly relative to VM/ESA 2.1. The amount of decrease ranged from 0.4% to 1.5%.

When comparing guest performance to VSE native, the V=R guest, with this workload, achieves about 72% of native. The V=V guest with a VMESA mode supervisor achieves about 59% of native performance, and the V=V guest with

an ESA mode supervisor is approximately 51% of native. These rates are lower than average because the PACEX8 workload is very I/O intensive.

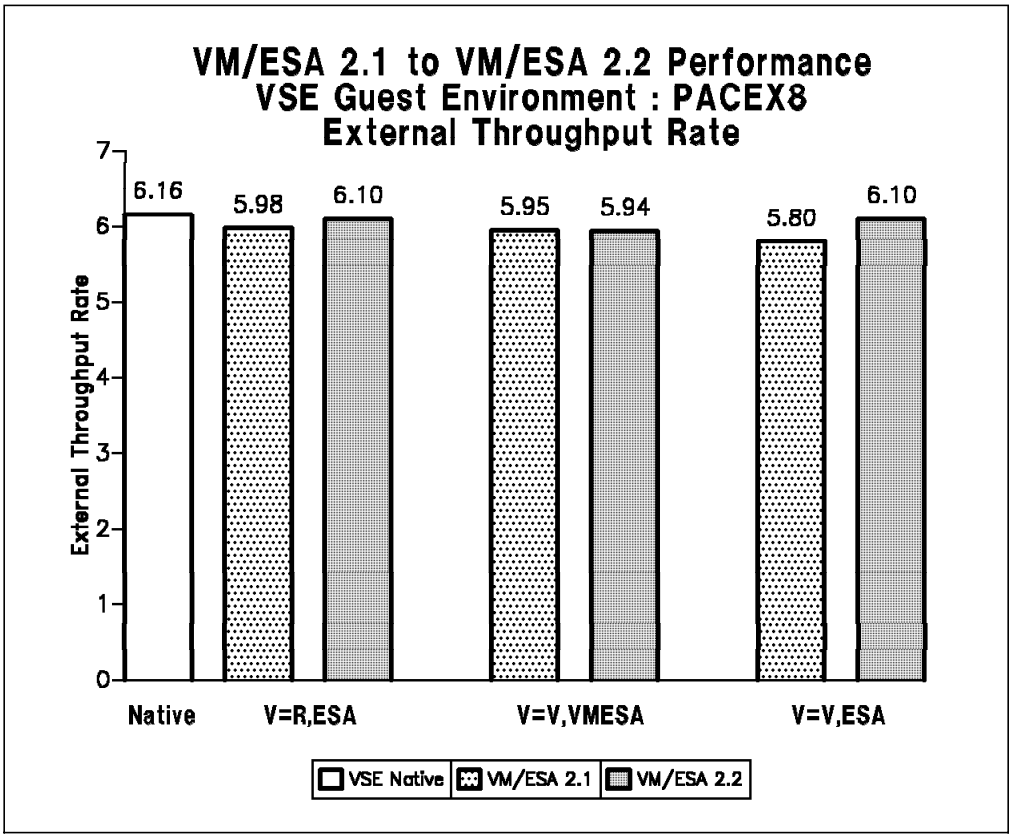

Figure 4. VSE guest migration from VM/ESA Release 2.1. PACEX8 workload on a single VSE/ESA guest of VM/ESA Release 2.1 and VM/ESA Release 2.2 on the 9121-320 processor.

Figure 4 reflects elapsed time comparisons of the various guest modes running under VM/ESA Release 2.1 and VM/ESA Release 2.2. The elapsed time duration of the batch jobs remains unchanged (within run variability) for all environments. When comparing external throughput rates, all guest modes are nearly identical to VSE native.

Migration: VSE/ESA Guest

Workload: PACEX8

Hardware Configuration

Processor models: 9121-320[3]
Storage
  Real: 96MB
  Expanded: 0MB
DASD:

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

Software Configuration

VSE version: 1.3.2

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
VSEVR	1	VSE V=R	16MB/ESA	100		IOASSIST ON CCWTRANS OFF
or						
VSEVV	1	VSE V=V	16MB/ESA	100		IOASSIST OFF
SMART	1	RTM	16MB/370	100		
WRITER	1	CP monitor	2MB/XA	100		

Additional Information: Minidisk caching (MDC) for non-CMS guests is a new function that is not available on VM/ESA Release 2.1. Therefore, all VM/ESA Release 2.2 guest measurements in this section were run with MDC inactive. See “VSE Guest Environment” on page 90 for VM/ESA Release 2.2 guest measurements with minidisk caching.

The PACEX8 workload consists of the base PACEX1 workload (7 batch jobs), running eight copies in eight partitions. All partitions are balanced with equal priorities.

For all guest measurements in this section, VSE/ESA was run in an ESA virtual machine and the VSE supervisor was defined as MODE=ESA. For all comparisons with PACEX8, the guest was run in three modes:

- V=R, mode=ESA
- V=V, mode=ESA
- V=V, mode=VMESA

All DASD are dedicated to the VSE guest for these measurements (except for the VM system DASD volumes). The VM system used for these guest measurements has a 32MB V=R area defined. For measurements with V=V

[3] See “Hardware Used” on page 17 for an explanation of how this processor model was defined.

guests, the V=R area is configured, but not used. Therefore, if the real storage configuration on the processor is 96MB, then 64MB of effective useable storage is available for the VM system and V=V guest. For the V=V measurements, it is this effective real storage size that is shown in this section's measurement results tables.

Measurement Results: The VSE guest measurement results for the PACEX8 workload are provided in the following three tables. The VSE native results are provided in Table 13 on page 63.

Migration: VSE/ESA Guest

Table 8 (Page 1 of 2). VSE/ESA V=R guest migration from VM/ESA 2.1 on the 9121-320, PACEX8 workload

Release Run ID	VM/ESA 2.1 LB68REX1	VM/ESA 2.2 LB78REX8	Difference	%Difference
Environment				
Real Storage	**96MB**	**96MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**16MB**	**16MB**		
Guest Setting	**V=R**	**V=R**		
Guest DASD	**Dedicated**	**Dedicated**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (C)	562.0	551.0	-11.0	-1.95%
ETR (C)	5.98	6.10	0.09	2.00%
ITR (H)	26.01	25.62	-0.39	-1.50%
ITR	25.99	25.27	-0.72	-2.78%
ITRR (H)	1.000	0.985	-0.015	-1.50%
ITRR	1.000	0.972	-0.028	-2.78%
Proc. Usage (Sec)				
PBT/CMD (H)	2.304	2.343	0.040	1.72%
PBT/CMD	2.308	2.374	0.066	2.86%
CP/CMD (H)	0.352	0.380	0.028	8.01%
CP/CMD	0.301	0.396	0.095	31.44%
EMUL/CMD (H)	1.951	1.963	0.011	0.59%
EMUL/CMD	2.007	1.979	-0.029	-1.42%
Processor Util.				
TOTAL (H)	22.95	23.69	0.73	3.19%
TOTAL	23.00	24.00	1.00	4.35%
TOTAL EMUL (H)	19.44	19.84	0.40	2.04%
TOTAL EMUL	20.00	20.00	0.00	0.00%
TVR(H)	1.18	1.19	0.01	1.13%
TVR	1.15	1.20	0.05	4.35%
Storage				
NUCLEUS SIZE (V)	2360KB	2536KB	176KB	7.46%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	14571	14525	-46	-0.32%
FREEPGS	70	69	-1	-1.43%
FREE UTIL	0.39	0.42	0.03	6.98%
SHRPGS	1053	1053	0	0.00%
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	0.000	na
I/O				
VIO RATE	5.000	5.000	0.000	0.00%
VIO/CMD	50.179	49.464	-0.714	-1.42%
RIO RATE (V)	1.000	1.000	0.000	0.00%
RIO/CMD (V)	10.036	9.893	-0.143	-1.42%
DASD IO TOTAL (V)	173364	174783	1419	0.82%
DASD IO RATE (V)	321.04	323.67	2.63	0.82%
DASD IO/CMD (V)	3221.91	3202.04	-19.87	-0.62%
MDC HIT RATIO	na	0		

Table 8 (Page 2 of 2). VSE/ESA V=R guest migration from VM/ESA 2.1 on the 9121-320, PACEX8 workload

Release Run ID	VM/ESA 2.1 LB68REX1	VM/ESA 2.2 LB78REX8	Difference	%Difference
Environment				
Real Storage	**96MB**	**96MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**16MB**	**16MB**		
Guest Setting	**V=R**	**V=R**		
Guest DASD	**Dedicated**	**Dedicated**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
PRIVOPs				
PRIVOP/CMD (R)	55.090	54.509	-0.581	-1.05%
DIAG/CMD (R)	620.648	628.492	7.843	1.26%
SIE/CMD	3462.321	3462.500	0.179	0.01%
SIE INTCPT/CMD	3081.466	3081.625	0.159	0.01%
FREE TOTL/CMD	953.393	949.714	-3.679	-0.39%

Note: V=VMPRF, H=Hardware Monitor, C=VSE Console Timestamps, Unmarked=RTM

Migration: VSE/ESA Guest

Table 9 (Page 1 of 2). VSE/ESA V=V mode=VMESA guest migration from VM/ESA 2.1 on the 9121-320, PACEX8 workload

Release Run ID	VM/ESA 2.1 LB68VEM1	VM/ESA 2.2 LB78VEMG	Difference	%Difference
Environment				
Real Storage	**64MB**	**64MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**16MB**	**16MB**		
Guest Setting	**V=V**	**V=V**		
Guest DASD	**Dedicated**	**Dedicated**		
VSE Supervisor	**VMESA**	**VMESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (C)	565.0	566.0	1.0	0.17%
ETR (C)	5.95	5.94	-0.01	-0.17%
ITR (H)	21.17	21.02	-0.15	-0.73%
ITR	21.20	21.13	-0.07	-0.35%
ITRR (H)	1.000	0.993	-0.007	-0.73%
ITRR	1.000	0.996	-0.004	-0.35%
Proc. Usage (Sec)				
PBT/CMD (H)	2.834	2.855	0.021	0.74%
PBT/CMD	2.830	2.840	0.010	0.35%
CP/CMD (H)	1.117	1.141	0.025	2.21%
CP/CMD	1.011	1.014	0.004	0.35%
EMUL/CMD (H)	1.717	1.713	-0.004	-0.22%
EMUL/CMD	1.819	1.826	0.006	0.35%
Processor Util.				
TOTAL (H)	28.04	28.15	0.11	0.38%
TOTAL	28.00	28.00	0.00	0.00%
TOTAL EMUL (H)	16.99	16.89	-0.10	-0.57%
TOTAL EMUL	18.00	18.00	0.00	0.00%
TVR(H)	1.65	1.67	0.02	0.96%
TVR	1.56	1.56	0.00	0.00%
Storage				
NUCLEUS SIZE (V)	2360KB	2536KB	176KB	7.46%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	14530	14504	-26	-0.18%
FREEPGS	76	76	0	0.00%
FREE UTIL	0.43	0.53	0.10	21.97%
SHRPGS	1053	1053	0	0.00%
Paging				
PAGE/CMD	0.000	0.000	0.000	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	80.857	81.143	0.286	0.35%
I/O				
VIO RATE	313.000	312.000	-1.000	-0.32%
VIO/CMD	3163.536	3164.571	1.036	0.03%
RIO RATE (V)	298.000	327.000	29.000	9.73%
RIO/CMD (V)	3011.929	3316.714	304.786	10.12%
DASD IO TOTAL (V)	178548	176178	-2370	-1.33%
DASD IO RATE (V)	297.58	326.26	28.68	9.64%
DASD IO/CMD (V)	3007.68	3309.16	301.48	10.02%
MDC HIT RATIO	na	0		

Table 9 (Page 2 of 2). VSE/ESA V=V mode=VMESA guest migration from VM/ESA 2.1 on the 9121-320, PACEX8 workload

Release **Run ID**	**VM/ESA 2.1** **LB68VEM1**	**VM/ESA 2.2** **LB78VEMG**	**Difference**	**%Difference**
Environment				
Real Storage	**64MB**	**64MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**16MB**	**16MB**		
Guest Setting	**V=V**	**V=V**		
Guest DASD	**Dedicated**	**Dedicated**		
VSE Supervisor	**VMESA**	**VMESA**		
Processors	**1**	**1**		
PRIVOPs				
PRIVOP/CMD	3164.284	3167.210	2.925	0.09%
DIAG/CMD	460.419	464.718	4.299	0.93%
SIE/CMD	13391.964	13358.143	-33.821	-0.25%
SIE INTCPT/CMD	12454.527	12556.654	102.127	0.82%
FREE TOTL/CMD	4083.286	4118.000	34.714	0.85%

Note: V=VMPRF, H=Hardware Monitor, C=VSE Console Timestamps, Unmarked=RTM

Migration: VSE/ESA Guest

Table 10 (Page 1 of 2). VSE/ESA V=V guest migration from VM/ESA 2.1 on the 9121-320, PACEX8 workload

Release Run ID	VM/ESA 2.1 LB68VEXC	VM/ESA 2.2 LB78VEXK	Difference	%Difference
Environment				
Real Storage	**64MB**	**64MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**16MB**	**16MB**		
Guest Setting	**V=V**	**V=V**		
Guest DASD	**Dedicated**	**Dedicated**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (C)	573.0	548.0	-25.0	-4.36%
ETR (C)	5.86	6.13	0.27	4.61%
ITR (H)	18.32	18.24	-0.08	-0.43%
ITR	18.72	19.06	0.34	1.80%
ITRR (H)	1.000	0.996	-0.004	-0.43%
ITRR	1.000	1.018	0.018	1.80%
Proc. Usage (Sec)				
PBT/CMD (H)	3.276	3.290	0.014	0.43%
PBT/CMD	3.205	3.149	-0.057	-1.77%
CP/CMD (H)	1.197	1.206	0.009	0.77%
CP/CMD	1.034	1.082	0.048	4.68%
EMUL/CMD (H)	2.079	2.084	0.005	0.23%
EMUL/CMD	2.171	2.066	-0.105	-4.84%
Processor Util.				
TOTAL (H)	31.68	33.44	1.75	5.53%
TOTAL	31.00	32.00	1.00	3.23%
TOTAL EMUL (H)	20.11	21.18	1.07	5.33%
TOTAL EMUL	21.00	21.00	0.00	0.00%
TVR(H)	1.58	1.58	0.00	0.20%
TVR	1.48	1.52	0.05	3.23%
Storage				
NUCLEUS SIZE (V)	2368KB	2536KB	168KB	7.09%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	14529	14506	-23	-0.16%
FREEPGS	77	77	0	0.00%
FREE UTIL	0.44	0.52	0.08	17.52%
SHRPGS	538	1053	515	95.72%
Paging				
PAGE/CMD	0.000	0.000	0.000	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	82.714	78.714	-4.000	-4.84%
I/O				
VIO RATE	300.000	303.000	3.000	1.00%
VIO/CMD	3101.786	2981.304	-120.482	-3.88%
RIO RATE (V)	320.000	325.000	5.000	1.56%
RIO/CMD (V)	3308.571	3197.768	-110.804	-3.35%
DASD IO TOTAL (V)	172554	175290	2736	1.59%
DASD IO RATE (V)	319.54	324.61	5.07	1.59%
DASD IO/CMD (V)	3303.86	3193.94	-109.92	-3.33%
MDC HIT RATIO	na	0		

Table 10 (Page 2 of 2). VSE/ESA V=V guest migration from VM/ESA 2.1 on the 9121-320, PACEX8 workload

Release Run ID	VM/ESA 2.1 LB68VEXC	VM/ESA 2.2 LB78VEXK	Difference	%Difference
Environment				
Real Storage	**64MB**	**64MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**16MB**	**16MB**		
Guest Setting	**V=V**	**V=V**		
Guest DASD	**Dedicated**	**Dedicated**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
PRIVOPs				
PRIVOP/CMD	3108.555	2979.746	-128.809	-4.14%
DIAG/CMD	656.112	623.973	-32.138	-4.90%
SIE/CMD	13513.446	12938.661	-574.786	-4.25%
SIE INTCPT/CMD	12432.371	11903.568	-528.803	-4.25%
FREE TOTL/CMD	3959.946	3896.357	-63.589	-1.61%

Note: V=VMPRF, H=Hardware Monitor, C=VSE Console Timestamps, Unmarked=RTM

9121-320 / VSECICS

This section examines VSE/ESA 1.3.2 guest performance running under VM/ESA Release 2.2 compared to VM/ESA Release 2.1 guest performance. The VSECICS workload is used for these measurements. VSECICS is an online transaction processing workload and is characterized by light I/O. See Appendix A, "Workloads" on page 208 for a detailed description of the workload. All DASD are dedicated to the VSE guest for these measurements (except for the VM system DASD volumes).

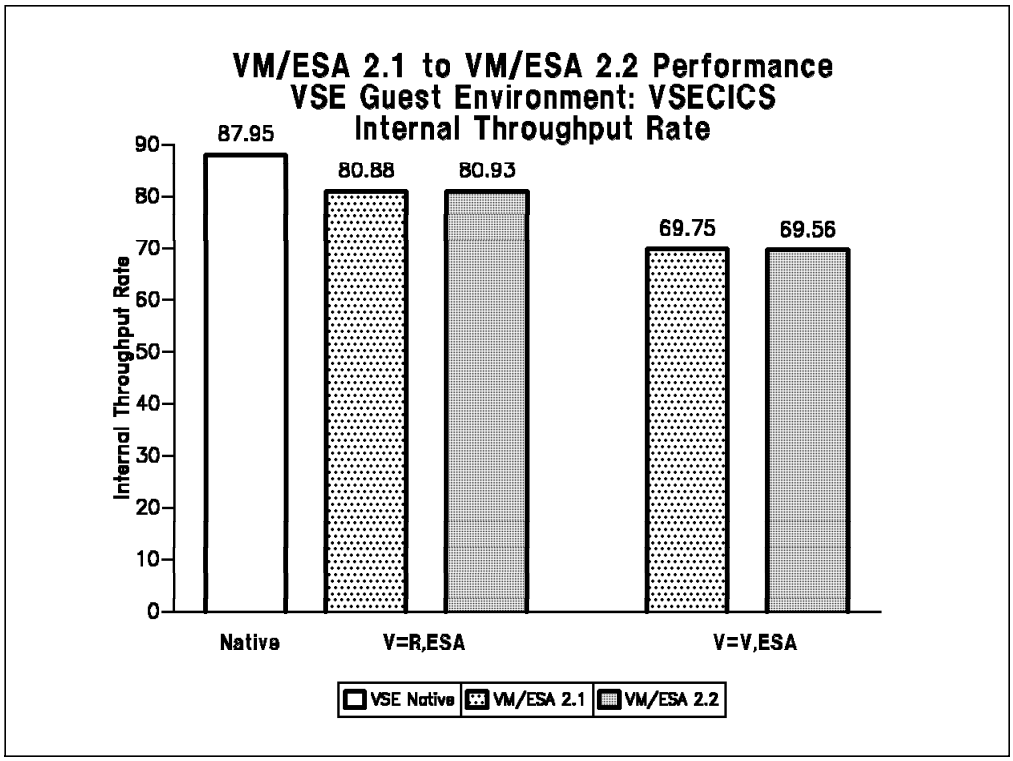

Figure 5. VSE guest migration from VM/ESA Release 2.1. VSECICS workload on a single VSE/ESA guest of VM/ESA Release 2.1 and VM/ESA Release 2.2 on the 9121-320 processor.

Comparing VM/ESA Release 2.2 to VM/ESA Release 2.1, internal throughput rates were equivalent within measurement variability.

When comparing guest performance to VSE running native, the V=R guest, with this workload, achieved about 92% of native. The V=V guest achieved approximately 80% of native mode performance.

Workload: VSECICS

Hardware Configuration

Processor models: 9121-320[4]
Storage
Real: 128MB
Expanded: 0MB

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*VSAM*	*VSE Sys.*	*VM Sys.*
3380-A	3880-03	2				18	2	2

Software Configuration

VSE version: 1.3.2

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
VSEVR	1	VSE V=R	32MB/ESA	100		IOASSIST ON CCWTRANS OFF
or						
VSEVV	1	VSE V=V	32MB/ESA	100		IOASSIST OFF
SMART	1	RTM	16MB/370	100		
WRITER	1	CP monitor	2MB/XA	100		

Measurement Discussion: The VSECICS workload was used to compare guest environments of VM/ESA Release 2.1 and VM/ESA Release 2.2 as well as VSE/ESA running native. For these measurement comparisons, the number of terminals was adjusted so that when running as a guest of VM/ESA Release 2.1, the CPU utilization was near 70%. Then, the same number of terminals was run in the same guest mode under VM/ESA Release 2.2.

The VSE guest measurement results for the VSECICS workload are provided in the following two tables. The VSE native results are provided in Table 13 on page 63.

[4] See "Hardware Used" on page 17 for an explanation of how this processor model was defined.

Migration: VSE/ESA Guest

Table 11 (Page 1 of 2). VSE/ESA V=R guest migration from VM/ESA 2.1 on the 9121-320, VSECICS workload

RELEASE RUN ID	VM/ESA 2.1 VR13207B	VM/ESA 2.2 VR13207C	Difference	%Difference
Environment				
RUN MODE	**ESA**	**ESA**		
REAL STORAGE	**128M**	**128M**		
EXP. STORAGE	**0M**	**0M**		
USERS	**1140**	**1140**		
VM MODE	**ESA**	**ESA**		
VM SIZE	**32MB**	**32MB**		
GUEST SETTING	**V=R**	**V=R**		
GUEST DASD	**Dedicated**	**Dedicated**		
VSE MODE	**ESA**	**ESA**		
PROCESSORS	**1**	**1**		
Response Time				
AVG RESP (C)	0.212	0.241	0.029	13.68%
Throughput				
ETR (C)	57.68	57.64	-0.04	-0.08%
ITR (H)	80.88	80.93	0.05	0.07%
ITRR (H)	1.000	1.001	0.001	0.07%
Proc. Usage				
PBT/CMD (H)	12.364	12.356	-0.008	-0.07%
PBT/CMD	12.310	12.319	0.009	0.08%
CP/CMD (H)	0.625	0.648	0.023	3.69%
CP/CMD	0.520	0.521	0.000	0.08%
EMUL/CMD (H)	11.739	11.708	-0.031	-0.27%
EMUL/CMD	11.789	11.798	0.009	0.08%
Processor Util.				
TOTAL (H)	71.31	71.21	-0.10	-0.14%
TOTAL	71.00	71.00	0.00	0.00%
TOTAL EMUL (H)	67.71	67.48	-0.23	-0.34%
TOTAL EMUL	68.00	68.00	0.00	0.00%
TVR(H)	1.05	1.06	0.00	0.20%
TVR	1.04	1.04	0.00	0.00%
Storage				
NUCLEUS SIZE (V)	2360K	2576K	216K	9.15%
TRACE TABLE (V)	200K	200K	0K	0.00%
PGBLPGS	22686	23237	551	2.43%
FREEPGS	65	74	9	13.85%
FREE UTIL	0.43	0.52	0.09	21.56%
SHRPGS	1053	1079	26	2.47%
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	0.000	na
I/O				
VIO RATE	0.000	0.000	0.000	na
VIO/CMD	0.000	0.000	0.000	na
RIO RATE (V)	0.000	1.000	1.000	na
RIO/CMD (V)	0.000	0.017	0.017	na
MDC HIT RATIO	na	0		

Table 11 (Page 2 of 2). VSE/ESA V=R guest migration from VM/ESA 2.1 on the 9121-320, VSECICS workload

RELEASE RUN ID	VM/ESA 2.1 VR13207B	VM/ESA 2.2 VR13207C	Difference	%Difference
Environment				
RUN MODE	**ESA**	**ESA**		
REAL STORAGE	**128M**	**128M**		
EXP. STORAGE	**0M**	**0M**		
USERS	**1140**	**1140**		
VM MODE	**ESA**	**ESA**		
VM SIZE	**32MB**	**32MB**		
GUEST SETTING	**V=R**	**V=R**		
GUEST DASD	**Dedicated**	**Dedicated**		
VSE MODE	**ESA**	**ESA**		
PROCESSORS	**1**	**1**		
PRIVOPs				
PRIVOP/CMD (R)	0.004	0.004	0.000	4.82%
DIAG/CMD (R)	3.835	4.030	0.195	5.08%
SIE/CMD	7.889	7.808	-0.081	-1.02%
SIE INTCPT/CMD	5.759	5.934	0.175	3.04%
FREE TOTL/CMD	0.815	1.128	0.313	38.40%

Note: V=VMPRF, H=Hardware Monitor, C=CICSPARS, Unmarked=RTM. Per command (/CMD), for this workload, means per CICS transaction.

Migration: VSE/ESA Guest

Table 12 (Page 1 of 2). VSE/ESA V=V guest migration from VM/ESA 2.1 on the 9121-320, VSECICS workload

RELEASE RUN ID	VM/ESA 2.1 VV13207H	VM/ESA 2.2 VV13207K	Difference	%Difference
Environment				
RUN MODE	**ESA**	**ESA**		
REAL STORAGE	**128M**	**128M**		
EXP. STORAGE	**0M**	**0M**		
USERS	**976**	**976**		
VM MODE	**ESA**	**ESA**		
VM SIZE	**32MB**	**32MB**		
GUEST SETTING	**V=V**	**V=V**		
GUEST DASD	**Dedicated**	**Dedicated**		
VSE MODE	**ESA**	**ESA**		
PROCESSORS	**1**	**1**		
Response Time				
AVG RESP (C)	0.216	0.202	-0.014	-6.25%
Throughput				
ETR (C)	49.53	49.51	-0.02	-0.04%
ITR (H)	69.75	69.56	-0.19	-0.27%
ITRR (H)	1.000	0.997	-0.003	-0.27%
Proc. Usage				
PBT/CMD (H)	14.337	14.377	0.039	0.27%
PBT/CMD	14.333	14.340	0.006	0.04%
CP/CMD (H)	2.000	2.053	0.053	2.65%
CP/CMD	1.817	1.818	0.001	0.04%
EMUL/CMD (H)	12.337	12.323	-0.014	-0.11%
EMUL/CMD	12.517	12.522	0.005	0.04%
Processor Util.				
TOTAL (H)	71.02	71.18	0.16	0.23%
TOTAL	71.00	71.00	0.00	0.00%
TOTAL EMUL (H)	61.11	61.02	-0.09	-0.15%
TOTAL EMUL	62.00	62.00	0.00	0.00%
TVR(H)	1.16	1.17	0.00	0.39%
TVR	1.15	1.15	0.00	0.00%
Storage				
NUCLEUS SIZE (V)	2360K	2576K	216K	9.15%
TRACE TABLE (V)	200K	200K	0K	0.00%
PGBLPGS	22675	23239	564	2.49%
FREEPGS	71	70	-1	-1.41%
FREE UTIL	0.44	0.47	0.03	6.30%
SHRPGS	1080	491	-589	-54.54%
Paging				
PAGE/CMD	0.000	0.000	0.000	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	3.129	3.151	0.022	0.69%
I/O				
VIO RATE	157.000	158.000	1.000	0.64%
VIO/CMD	3.170	3.191	0.022	0.68%
RIO RATE (V)	157.000	158.000	1.000	0.64%
RIO/CMD (V)	3.170	3.191	0.022	0.68%
MDC HIT RATIO	na	0		

Table 12 (Page 2 of 2). VSE/ESA V=V guest migration from VM/ESA 2.1 on the 9121-320, VSECICS workload

RELEASE RUN ID	VM/ESA 2.1 VV13207H	VM/ESA 2.2 VV13207K	Difference	%Difference
Environment				
RUN MODE	**ESA**	**ESA**		
REAL STORAGE	**128M**	**128M**		
EXP. STORAGE	**0M**	**0M**		
USERS	**976**	**976**		
VM MODE	**ESA**	**ESA**		
VM SIZE	**32MB**	**32MB**		
GUEST SETTING	**V=V**	**V=V**		
GUEST DASD	**Dedicated**	**Dedicated**		
VSE MODE	**ESA**	**ESA**		
PROCESSORS	**1**	**1**		
PRIVOPs				
PRIVOP/CMD (R)	3.174	3.189	0.015	0.47%
DIAG/CMD (R)	4.004	4.115	0.111	2.76%
SIE/CMD	23.378	23.266	-0.111	-0.48%
SIE INTCPT/CMD	15.897	16.054	0.157	0.99%
FREE TOTL/CMD	4.159	4.282	0.123	2.96%

Note: V=VMPRF, H=Hardware Monitor, C=CICSPARS, Unmarked=RTM. Per command (/CMD), for this workload, means per CICS transaction.

Table 13. VSE/ESA 1.3.2 native on the 9121-320, PACEX8 and VSECICS workloads

Release Workload Run ID	VSE/ESA 1.3.2 PACEX8 LBNAT132	VSE/ESA 1.3.2 VSECICS NAT3207B
Environment		
IML Mode	**ESA**	**ESA**
Real Storage	**64MB**	**64MB**
VSE Supervisor	**ESA**	**ESA**
Processors	**1**	**1**
Throughput		
ITR (H)	35.46	87.95
ETR(C)	6.16	60.81
Proc. Usage (Sec)		
PBT/CMD (H)	1.69	0.0164
Processor Util.		
TOTAL (H)	17.39	69.14

Note: H=Hardware Monitor, C=VSE Console Timestamps (PACEX8), C=CICSPARS (VSECICS)

VMSES/E

A number of requirements were addressed by VM Service Enhancements Staged/Extended (VMSES/E) in VM/ESA Release 2.2. Some of these changes improve execution performance:

- Performance improvements for VMFBLD and Test Build.
- VMFINS Requisite Support Enhancements provides a requisite checking enhancement that improves processing performance.
- VMFSETUP enhancements identify as many errors as possible before exiting to help identify all of the problems in one pass. VMFSETUP with the LINK option minimizes the number of links performed.

Others reduce the number of manual steps through automation:

- The PSU/RSU Planning and Local Mod Identification enhances the VMSES/E automation of the RSU process.
- The CP Loadlist Build (GENCPBLS EXEC) requirement automates the modifying of the CP LOADLIST with a local modification.
- VMSES/E Execupdt Support (VMFEXUPD EXEC) provides automation to the EXECUPDT command to facilitate the application of local modifications.
- VMFPPF Compile of all components in a source $PPF automates manual customer steps and tasks.
- VMFREC part handler for upper casing files eliminates user install tasks.

Three primary VMSES/E tools that help with the servicing of products were measured to quantify the effects of the new function and the performance enhancements:

- 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.

The biggest performance impact of all the new enhancements found in VMSES/E in VM/ESA Release 2.2 came from the VMFBLD improvements. The regression measurements showed significant savings in build response time. Most of these improvements came from reducing storage requirements in the build data structure and from the elimination of extraneous code paths for parts that were serviced in the build lists.

Overall, for the dedicated, single-user measurements reported here, the process of receiving and applying CMS service, and building CMS on VMSES/E in VM/ESA Release 2.2 improved total elapsed time (Total Time (Q)) when compared to VMSES/E in VM/ESA Release 2.1:

- 11% improvement on the 9121-742 configuration
- 15% improvement on the 9121-480 configuration

- 22% improvement on the 9221-120 configuration

An additional measurement comparison was made in a multi-user environment on the 9121-480 configuration. There were 1860 users executing CMS commands (processor busy at 90%) and one user executing the VMFBLD command once. The results showed a 57% total elapsed time improvement for VMSES/E in VM/ESA Release 2.2 when compared to the same measurement on VMSES/E in VM/ESA Release 2.1.

The measurements described in the following sections are provided to demonstrate the performance of these changes on VMSES/E in VM/ESA Release 2.2.

9121-742

Hardware Configuration

Processor model: 9121-742
Processors used: 4
Storage
  Real: 1024MB
  Expanded: 1024MB
Tape: 3480 (one service tape for the receive command)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	4	6	7	7	32 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.

Software Configuration

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
MAINT	1	MAINT	30MB/XC			

Measurement Discussion: All measurements were performed on a dedicated, first-level system with only one active user logged on (MAINT user ID). The objective of these measurements was to show that the new functional enhancements to VMSES/E in VM/ESA Release 2.2 did not degrade performance when compared to VMSES/E in VM/ESA Release 2.1 in an established service environment — where all Software Inventory Management (SIM) tables had been previously initialized. The SIM tables were initialized by 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 SIM was 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 2.1 and VMSES/E in VM/ESA Release 2.2; 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, 15 build lists were processed after running the VMFBLD command with the SERVICED option.

The methodology described in this section applies to both VMSES/E in VM/ESA Release 2.2 and VMSES/E in VM/ESA Release 2.1. 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. 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 running 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 running the command.

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

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.) Note that the above measurement discussion applies to all measurement configurations described in the VMSES/E sections.

Workload: Receive

Command: VMFREC PPF ESA CMS
Scenario Details: 1728 parts received from 7 tape files.

Table 14. VMFREC measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9121-742

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	450	451	1	0%
Total CPU (Q)	128	126	-2	-2%
Virtual CPU (Q)	117	116	-1	-1%

Note: Q=CP QUERY TIME.

Migration: VMSES/E

Workload: Apply

Command: VMFAPPLY PPF ESA CMS
Scenario Details: 206 PTFs after receiving parts from COR tape.

Table 15. VMFAPPLY measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9121-742

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	251	254	3	1%
Total CPU (Q)	196	195	-1	-1%
Virtual CPU (Q)	191	192	1	1%

Note: Q=CP QUERY TIME.

Workload: Build with STATUS Option

Command: VMFBLD PPF ESA CMS (STATUS
Scenario Details: 149 build requirements identified.

Table 16. VMFBLD STATUS measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9121-742

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	101	93	-8	-8%
Total CPU (Q)	91	81	-10	-11%
Virtual CPU (Q)	91	81	-10	-11%

Note: Q=CP QUERY TIME.

Workload: Build with SERVICED Option

Command: VMFBLD PPF ESA CMS (SERVICED
Scenario Details: 16 build lists processed; 149 objects built.

Table 17. VMFBLD SERVICED measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9121-742

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	707	546	-161	-23%
Total CPU (Q)	374	210	-164	-44%
Virtual CPU (Q)	364	202	-162	-45%

Note: Q=CP QUERY TIME.

9121-480

Hardware Configuration

Processor model: 9121-480
Processors used: 2
Storage
  Real: 224MB
  Expanded: 32MB
Tape: 3480 (one service tape for the receive command)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*TDSK*	*- Number of Volumes - User*	*Server*	*System*
3390-2	3990-2	4	16	5	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2	16 R		

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

Software Configuration

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
MAINT	1	MAINT	30MB/XC			

Measurement Discussion: See “9121-742” on page 66.

Workload: Receive

Command: VMFREC PPF ESA CMS
Scenario Details: 1728 parts received from 7 tape files.

Table 18. VMFREC measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9121-480

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	575	578	3	1%
Total CPU (Q)	187	188	1	1%
Virtual CPU (Q)	170	171	1	1%

Note: Q=CP QUERY TIME.

Migration: VMSES/E

Workload: Apply

Command: VMFAPPLY PPF ESA CMS
Scenario Details: 206 PTFs after receiving parts from COR tape.

Table 19. VMFAPPLY measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9121-480

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	356	362	6	2%
Total CPU (Q)	283	288	5	2%
Virtual CPU (Q)	275	280	5	2%

Note: Q=CP QUERY TIME.

Workload: Build with STATUS Option

Command: VMFBLD PPF ESA CMS (STATUS
Scenario Details: 149 build requirements identified.

Table 20. VMFBLD STATUS measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9121-480

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	148	133	-15	-10%
Total CPU (Q)	133	118	-15	-11%
Virtual CPU (Q)	132	117	-15	-11%

Note: Q=CP QUERY TIME.

Workload: Build with SERVICED Option

Command: VMFBLD PPF ESA CMS (SERVICED
Scenario Details: 16 build lists processed; 149 objects built.

Table 21. VMFBLD SERVICED measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9121-480

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	913	626	-287	-31%
Total CPU (Q)	545	280	-263	-50%
Virtual CPU (Q)	529	266	-265	-49%

Note: Q=CP QUERY TIME.

9221-120

Hardware Configuration

Processor model: 9221-120
Processors used: 1
Storage
  Real: 48MB
  Expanded: 208MB
Tape: 3480 (one service tape for the receive command)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*TDSK*	*- Number of Volumes - User*	*Server*	*System*
3390-2	3990-2	1	16	6	6			
3390-2	3990-3	1		2	2	8 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.

Software Configuration

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
MAINT	1	MAINT	30MB/XC			

Measurement Discussion: See “9121-742” on page 66.

Workload: Receive

Command: VMFREC PPF ESA CMS
Scenario Details: 1728 parts received from 7 tape files.

Table 22. VMFREC measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9221-120

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	2006	2069	63	3%
Total CPU (Q)	1604	1656	55	4%
Virtual CPU (Q)	1435	1490	52	3%

Note: Q=CP QUERY TIME.

Migration: VMSES/E

Workload: Apply

Command: VMFAPPLY PPF ESA CMS
Scenario Details: 206 PTFs after receiving parts from COR tape.

Table 23. VMFAPPLY measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9221-120

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	2598	2627	29	1%
Total CPU (Q)	2485	2526	43	2%
Virtual CPU (Q)	2393	2436	41	2%

Note: Q=CP QUERY TIME.

Workload: Build with STATUS Option

Command: VMFBLD PPF ESA CMS (STATUS
Scenario Details: 149 build requirements identified.

Table 24. VMFBLD STATUS measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9221-120

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	1217	1081	-136	-11%
Total CPU (Q)	1186	1057	-129	-11%
Virtual CPU (Q)	1173	1046	-127	-11%

Note: Q=CP QUERY TIME.

Workload: Build with SERVICED Option

Command: VMFBLD PPF ESA CMS (SERVICED
Scenario Details: 16 build lists processed; 149 objects built.

Table 25. VMFBLD SERVICED measurement data: VMSES/E in VM/ESA Release 2.1 and VMSES/E in VM/ESA Release 2.2 on the 9221-120

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	5686	3210	-2476	-44%
Total CPU (Q)	5235	2801	-2434	-48%
Virtual CPU (Q)	5054	2635	-2419	-47%

Note: Q=CP QUERY TIME.

9121-480 / 90% Processor Busy plus VMFBLD

One environment that is important to VMSES/E is when other, non-VMSES/E activity is running on the system and the processor is at a high utilization. The build function of VMSES/E is one command that is periodically used in such an environment. This section contains performance results for this environment with VM/ESA Release 2.1 and VM/ESA Release 2.2.

Workload: FS8F0R plus Build with SERVICED Option

VMSES/E Command: VMFBLD PPF ESA CMS (SERVICED
VMSES/E Scenario Details: Four build lists processed.

Hardware Configuration

Processor model:	9121-480
Processors used:	2
Storage	
Real:	224MB
Expanded:	32MB (all for MDC)
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2	16 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*
3088-08	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*	*SHARE*	*RESERVED*	*Other Options*
MAINT	1	VMFBLD	30MB/XC			
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XC	10000	512	QUICKDSP ON
WRITER	1	CP Monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1860	Users	3MB/XC	100		

Measurement Discussion: These VMSES/E plus CMS-intensive measurements at 90% processor utilization were performed on a dedicated, first-level system. The 1860 CMS users were set up in a similar way (except for expanded storage usage) as set up in section "9121-480 / Minidisk" on page 33. For the VMSES/E command, the MAINT user was logged on and set up to run a subset of the build command. One minute after the 30 minute FS8F measurement period was started, the MAINT user started the command VMFBLD PPF ESA CMS (SERVICED. The reason for the subset build was to be able to complete within the 30 minute measurement interval. The VMFBLD for all of CMS described in the previous VMSES/E measurement sections would not have finished within the 30 minutes.

The changes that went into the build function for VMSES/E 2.2 show a significant performance improvement when compared to the same VMSES/E 2.1 function. Elapsed time for the VMFBLD command improved 57%. For the CMS users, overall external response time (AVG LAST(T)) improved 20% and the internal throughput rate (ITR(H)) improved 5.1%.

Table 26 contains performance data for the build command from the MAINT user (with CP QUERY TIME). Table 27 on page 75 contains system performance data collected by various performance tools for the complete environment.

Table 26. VMFBLD SERVICED at 90% processor utilization on the 9121-480.

Release	VMSES/E 2.1	VMSES/E 2.2	Difference	%Difference
Total Time (Q)	1604	683	-921	-57%
Total CPU (Q)	184	106	-78	-42%
Virtual CPU (Q)	192	111	-81	-42%

Note: Q=CP QUERY TIME.

Table 27 (Page 1 of 2). VMFBLD SERVICED at 90% processor utilization on the 9121-480.

Release Run ID	VM/ESA 2.1 L26E18FF	VM/ESA 2.2 L27E18FF	Difference	%Difference
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.131	0.134	0.003	2.29%
NONTRIV INT	0.528	0.441	-0.087	-16.48%
TOT INT	0.392	0.338	-0.054	-13.78%
TOT INT ADJ	0.354	0.298	-0.056	-15.91%
AVG FIRST (T)	0.291	0.254	-0.036	-12.39%
AVG LAST (T)	0.450	0.360	-0.090	-20.00%
Throughput				
AVG THINK (T)	26.23	26.19	-0.04	-0.15%
ETR	58.78	57.49	-1.29	-2.19%
ETR (T)	65.09	65.28	0.19	0.29%
ETR RATIO	0.903	0.881	-0.022	-2.48%
ITR (H)	68.93	72.46	3.53	5.12%
ITR	31.13	31.92	0.79	2.53%
EMUL ITR	45.24	45.86	0.63	1.39%
ITRR (H)	1.000	1.051	0.051	5.12%
ITRR	1.000	1.025	0.025	2.53%
Proc. Usage				
PBT/CMD (H)	29.015	27.603	-1.412	-4.87%
PBT/CMD	29.036	27.572	-1.464	-5.04%
CP/CMD (H)	9.631	8.977	-0.655	-6.80%
CP/CMD	9.064	8.425	-0.639	-7.05%
EMUL/CMD (H)	19.384	18.627	-0.757	-3.91%
EMUL/CMD	19.972	19.147	-0.824	-4.13%
Processor Util.				
TOTAL (H)	188.87	180.20	-8.66	-4.59%
TOTAL	189.00	180.00	-9.00	-4.76%
UTIL/PROC (H)	94.43	90.10	-4.33	-4.59%
UTIL/PROC	94.50	90.00	-4.50	-4.76%
TOTAL EMUL (H)	126.18	121.60	-4.57	-3.63%
TOTAL EMUL	130.00	125.00	-5.00	-3.85%
MASTER TOTAL (H)	94.10	89.78	-4.32	-4.60%
MASTER TOTAL	94.00	90.00	-4.00	-4.26%
MASTER EMUL (H)	56.26	53.84	-2.42	-4.30%
MASTER EMUL	58.00	56.00	-2.00	-3.45%
TVR(H)	1.50	1.48	-0.01	-1.00%
TVR	1.45	1.44	-0.01	-0.95%
Storage				
NUCLEUS SIZE (V)	2364KB	2536KB	172KB	7.28%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	82	84	2	2.44%
PGBLPGS	47474	47348	-126	-0.27%
PGBLPGS/USER	25.5	25.5	-0.1	-0.27%
FREEPGS	5123	5252	129	2.52%
FREE UTIL	0.95	0.93	-0.02	-2.46%
SHRPGS	1338	1275	-63	-4.71%

Table 27 (Page 2 of 2). VMFBLD SERVICED at 90% processor utilization on the 9121-480.

Release Run ID	VM/ESA 2.1 L26E18FF	VM/ESA 2.2 L27E18FF	Difference	%Difference
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	627	633	6	0.96%
WRITES/SEC	460	443	-17	-3.70%
PAGE/CMD	16.699	16.482	-0.217	-1.30%
PAGE IO RATE (V)	187.000	181.300	-5.700	-3.05%
PAGE IO/CMD (V)	2.873	2.777	-0.096	-3.33%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	6.852	8.854	2.002	29.22%
Queues				
DISPATCH LIST	41.13	39.31	-1.81	-4.40%
ELIGIBLE LIST	0.00	0.02	0.02	na
I/O				
VIO RATE	666	664	-2	-0.30%
VIO/CMD	10.232	10.171	-0.061	-0.59%
RIO RATE (V)	379	375	-4	-1.06%
RIO/CMD (V)	5.822	5.744	-0.078	-1.35%
NONPAGE RIO/CMD (V)	2.950	2.967	0.017	0.59%
DASD RESP TIME (V)	18.900	19.000	0.100	0.53%
MDC READS (blks)	365	na		
MDC READS (I/Os)	na	189		
MDC WRITES (blks)	142	na		
MDC WRITES (I/Os)	na	9.20		
MDC MODS	105	na		
MDC AVOID	174	174	0	0.00%
MDC HIT RATIO	0.90	0.90		
PRIVOPs				
PRIVOP/CMD	14.040	14.000	-0.040	-0.29%
DIAG/CMD	28.592	28.042	-0.549	-1.92%
DIAG 04/CMD	2.526	2.430	-0.096	-3.82%
DIAG 08/CMD	0.754	0.753	0.000	-0.03%
DIAG 0C/CMD	1.324	1.126	-0.198	-14.98%
DIAG 14/CMD	0.025	0.025	0.000	0.09%
DIAG 58/CMD	1.249	1.248	-0.001	-0.09%
DIAG 98/CMD	1.115	1.105	-0.010	-0.90%
DIAG A4/CMD	3.705	3.665	-0.040	-1.08%
DIAG A8/CMD	2.836	2.840	0.004	0.14%
DIAG 214/CMD	13.912	13.701	-0.211	-1.52%
SIE/CMD	52.556	54.087	1.531	2.91%
SIE INTCPT/CMD	34.687	34.616	-0.071	-0.21%
FREE TOTL/CMD	67.443	50.135	-17.307	-25.66%
VTAM Machines				
WKSET (V)	523	512	-11	-2.10%
TOT CPU/CMD (V)	3.9090	3.8124	-0.0966	-2.47%
CP CPU/CMD (V)	1.5363	1.4467	-0.0896	-5.83%
VIRT CPU/CMD (V)	2.3727	2.3657	-0.0070	-0.30%
DIAG 98/CMD (V)	1.114	1.105	-0.009	-0.81%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Migration from Other VM Releases

The performance results provided in this report apply to migration from VM/ESA Release 2.1. This section discusses how to use the information in this report along with similar information from earlier reports to get an understanding of the performance of migrating from earlier VM releases.

Migration Performance Measurements Matrix

The matrix on the following page is provided as an index to all the performance measurements pertaining to VM migration that are available in the VM/ESA performance reports. The numbers that appear in the matrix indicate which report includes migration results for that case:

10 *VM/ESA Release 1.0 Performance Report*

11 *VM/ESA Release 1.1 Performance Report*

20 *VM/ESA Release 2.0 Performance Report*

21 *VM/ESA Release 2.1 Performance Report*

22 *VM/ESA Release 2.2 Performance Report* (this document)

See “Referenced Publications” on page ix for more information on these reports.

Many of the comparisons listed in the matrix are for two consecutive VM releases. For migrations that skip one or more VM releases, you can get a general idea how the migration will affect performance by studying the applicable results for those two or more comparisons that, in combination, span those VM releases. For example, to get a general understanding of how migrating from VM/ESA 2.0 to VM/ESA 2.2 will tend to affect VSE guest performance, look at the VM/ESA 2.0 to VM/ESA 2.1 comparison measurements and the VM/ESA 2.1 to VM/ESA 2.2 comparison measurements. In each case, use the measurements from the system configuration that best approximates your VM system. For more discussion on the use of multiple comparisons, see page 82.

The comparisons listed for the CMS-intensive environment primarily consist of minidisk-only measurements but there are some SFS comparisons as well.

Internal throughput rate ratio (ITRR) information for the minidisk-only CMS-intensive environment has been extracted from the CMS comparisons listed in the matrix and is summarized in “Migration Summary: CMS-Intensive Environment” on page 79.

Migration from Other VM Releases

Table 28. Sources of VM migration performance measurement results

Source	Target	Processor	Report Number			
			CMS	OV/VM	VSE Guest	MVS Guest
VM/SP 5	VM/ESA 1.0 (370)	4381-13	10			
	VM/ESA 1.0 (370)	9221-170			20	
	VM/ESA 1.0 (370)	9221-120	20		20	
	VM/ESA 2.0	9221-170			20	
	VM/ESA 2.0	9221-120	20		20	
VM/SP 6	VM/ESA 1.0 (370)	4381-13	10			
		9370-80	10			
		9370-30	10			
VM/SP HPO5	VM/ESA 1.0 (ESA)	3090*-200J	10			
	VM/ESA 2.0	9121-480	20			
	VM/ESA 2.0	9121-320	20			
VM/ESA 1.0 (370)	VM/ESA 1.5 (370)	9221-120	22			
	VM/ESA 1.1	9221-170	11			
	VM/ESA 2.0	9221-170	20		20	
	VM/ESA 2.0	9221-120	20		20	
VM/XA 2.0	VM/ESA 1.0 (ESA)	3090-600J	10			
VM/XA 2.1	VM/ESA 1.0 (ESA)	3090-600J	10			10
	VM/ESA 1.0 (ESA)	3090-200J	10			
	VM/ESA 1.0 (ESA)	9021-720		11		
	VM/ESA 1.0 (ESA)	9121-320			11	
	VM/ESA 1.1	9021-720		11		
	VM/ESA 1.1	9121-320			11	
VM/ESA 1.0 (ESA)	VM/ESA 1.1	3090-600J				11
		9021-720	11	11		
		9021-580	11			
		9121-480	11			
		9121-320	11		11	
		9221-170	11			
VM/ESA 1.1	VM/ESA 2.0	9021-900	20			20
		9021-720		20		
		9121-480	20	20		
		9121-320			20	
		9221-170	20			
VM/ESA 2.0	VM/ESA 2.1	9121-742	21	21		
		9121-480	21	21		
		9121-320			21	
		9221-170	21			
VM/ESA 2.1	VM/ESA 2.2	9121-742	22			
		9121-480	22			
		9121-320			22	
		9221-170	22			

Migration Summary: CMS-Intensive Environment

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 minidisk-only 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 77.

Table 29 summarizes the ITR relationships that were observed for the CMS-intensive environment for a number of VM release-to-release transitions:

Table 29. Approximate VM relative capacity: CMS-intensive environment

Source	Target	Case	ITRR	ITRR Derivation	Notes
VM/SP 5	VM/ESA 1.5 (370) VM/ESA 2.2	 9221-120	0.94 0.87	R5*R13c R5*R13a*R2*R21	1,5,7 1,2,4,6,7
VM/SP 6	VM/ESA 1.5 (370) VM/ESA 2.2	 9221-120	1.09 1.00	R6*R13c R6*R13a*R2*R21	5 2,4,6,7
VM/ESA 1.0 (370)	VM/ESA 1.5 (370) VM/ESA 2.2	 9221-120 9221-170	1.02 0.94 1.00	R13c R13a*R2*R21 R13b*R11*R2*R21	 2,6,7 4,5,6,7
VM/ESA 1.5 (370)	VM/ESA 2.2	9221-120 9221-170	0.92 0.98	(1/R13c)*R13a*R2*R21 (1/R13c)*R13b*R11* R2*R21	2,6,7 4,5,6,7
VM/SP HPO 5	VM/ESA 2.2	UP, ¬4381 MP, ¬4381	0.95 1.06	RHa*R2*R21 RHb*R1E*R11* R2*R21	4,5,7 3,4,5,7
VM/XA 2.0	VM/ESA 2.2		1.15	RX20*RX21*R1E* R11*R2*R21	
VM/XA 2.1	VM/ESA 2.2		1.13	RX21*R1E*R11* R2*R21	
VM/ESA 1.0 ESA	VM/ESA 2.2		1.09	R1E*R11*R2*R21	
VM/ESA 1.1	VM/ESA 2.2		1.05	R11*R2*R21	
VM/ESA 2	VM/ESA 2.2		1.04	R2*R21	
VM/ESA 2.1	VM/ESA 2.2		1.03	R21	

Migration from Other VM Releases

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, there is still some variability. These ITRR variations are shown in "Derivation and Supporting Data" on page 82.

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 82 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 a 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. These estimates do not apply to 4381 processors. 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 than shown here when migrating on a 4381 processor from VM/SP, VM/SP HPO, or VM/ESA (370) to VM/ESA Release 2.2.
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 1.5 (370 Feature) ITRR is based (in part) on a comparison of VM/SP 5 to VM/ESA Release 1.0 (370 Feature), which showed an ITRR of 0.92. This comparison was done 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" in the *VM/ESA Release 2 Performance Report* for details.

   A VM/SP HPO example: The stated VM/SP HPO 5 to VM/ESA 2.2 ITRR for uniprocessors is based (in part) on a VM/SP HPO 5 to VM/ESA 2 comparison, which showed an ITRR of 0.91. Those measurements were 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 VM/SP HPO Release 5.0 (64MB/192MB) was compared to VM/ESA 2 (192MB/64MB), an ITRR of 0.95 was observed. See "CMS-Intensive Migration from VM/SP HPO Release 5" in the *VM/ESA Release 2 Performance Report* 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. Otherwise, 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 section "CMS-Intensive Migration from VM/ESA 1.1," subsection "9221-170 / Minidisk" in the *VM/ESA Release 2 Performance Report* 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*.
7. There has been growth in CMS real storage requirements on a per user basis. This growth is reflected in the ITR ratios to only a limited extent and should therefore be taken into consideration separately. The most significant growth took place in VM/SP 6 and in VM/ESA 2.0. The VM/SP 6 increase can affect the performance of migrations from VM/SP 5 and VM/SP HPO 5. The VM/ESA 2.0 growth can affect the performance of migrations from VM releases prior to VM/ESA 2.0. Storage constrained environments with large numbers of CMS users will be the most affected.

Table 29 on page 79 only shows 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 a positive effect on response time because they reduce I/O wait time, which is often the largest component of system response time.

Migration from Other VM Releases

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 2.2. From Table 29 on page 79, the estimated ITRR for migrating from VM/XA 2.1 to VM/ESA 2.2 is 1.13. Therefore, the estimated overall increase in system capacity is 1.30*1.13 = 1.47.

Table 29 on page 79 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 77.

Derivation and Supporting Data

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

The derivation column in Table 29 on page 79 shows how the stated ITR ratio was calculated. For example, the ITRR of 1.04 for migrating from VM/ESA 2.0 to VM/ESA 2.2 was calculated by multiplying the average ITRR for migrating from VM/ESA 2.0 to VM/ESA 2.1 (R2) by the average ITRR for migrating from VM/ESA 2.1 to VM/ESA 2.2 (R21): 1.01*1.03 = 1.04. R2 was calculated by averaging the ITRRs for VM measurement pairs 21 through 23 (see Table 30 on page 83). Likewise, R21 was calculated by averaging the ITRRs for VM measurement pairs 24 through 27.

For the case where the source system level is VM/ESA 1.5 (370), the term "1/R13c" resolves to "1/1.02." This takes into account the fact that VM/ESA 1.5 (370) has a somewhat higher ITR than VM/ESA 1.0 (370). This makes the ITRR smaller when migrating to VM/ESA 2.2 from VM/ESA 1.5 (370) as compared to migrating from VM/ESA 1.0 (370).

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, Table 29 on page 79 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 29 on page 79 are based on the following pairs of measurements:

Table 30 (Page 1 of 2). Derivation and supporting data: VM measurement pairs

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): FS7B0R Workload; Report 20								
1	H1SR0091	H17R0090	9221-120	16MB	80	9	0.92	(R5)
VM/SP 6 to VM/ESA 1.0 (370 Feature): FS7B0; Report 10								
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: FS7B0R; Report 20								
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: FS7B0R; Report 11								
5	H17R0281	H14R0287	9221-170	64MB	80	7	0.95	(R13b)
VM/ESA 1.0 (370 Feature) to VM/ESA 1.5 (370 Feature: FS7F0; Report 22								
6	H17E0106	H17E0113		16MB	90	10	0.985	
7	H17E0108	H17E0113		16MB	90	10	1.032	
avg							1.02	(R13c)
VM/SP HPO 5 to VM/ESA 2: FS7B0R; Report 20								
8	L1HR1033	L15R0951	9121-320	64MB/192MB	90	17	0.91	(RHa)
VM/SP HPO 5 to VM/ESA 1.0 (ESA Feature): FS7B0R; Report 10								
9	Y25R1141	Y23R1143	3090-200J	64MB/512MB	90	22	0.97	(RHb)
VM/XA 2.0 to VM/XA 2.1: FS7B0R; Report 10								
10	Y62R5401	Y6$R5401	3090-600J	512MB/2GB	90	15	1.02	(RX20)
VM/XA 2.1 to VM/ESA 1.0 (ESA Feature): FS7B0R; Report 10								
11	Y2$R2001	Y23R2001	3090-200J	256MB/2GB	90	11	1.064	
12	Y6$R5401	Y63R5405	3090-600J	512MB/2GB	90	12	1.029	
avg							1.04	(RX21)
VM/ESA 1.0 (ESA Feature) to VM/ESA 1.1: FS7B0R; Report 11								
13	Y63R5866	Y64R5865	9021-720	512MB/2GB	90	13	1.059	
14	L23R1770	L24R1770	9121-480	192MB/64MB	90	13	1.032	
15	L13R0911	L14R0910	9121-320	192MB/64MB	90	12	1.045	
16	H13R0280	H14R0287	9221-170	48M/16MB	80	11	1.043	
avg							1.04	(R1E)

Table 30 (Page 2 of 2). Derivation and supporting data: VM measurement pairs

Pair Number	Source Run ID	Target Run ID	Processor	Memory	Proc. Util.	Base Pg/cmd	ITR Ratio	Symbol
VM/ESA 1.1 to VM/ESA 2: FS7B0R; Report 20								
17	264RB424	265RB426	9021-900	1GB/4GB	90	16	1.018	
18	L24R1876	L25R187F	9121-480	192MB/64MB	90	14	1.005	
19	L24R1821	L25R1823	9121-480	128MB/0MB	90	15	1.009	
20	H14R0292	H15R0294	9221-170	48MB/16MB	90	12	1.009	(R11)
avg							1.01	
VM/ESA 2 to VM/ESA 2.1: FS7F0R; Report 21								
21	S45E5400	S46E5400	9121-742	1GB/1GB	90	17	1.012	
22	S45E5201	S46E5200	9121-742	320MB/64MB	90	19	1.011	
23	H15E0290	H16E0290	9221-170	48MB/16MB	90	15	1.016	
avg							1.01	(R2)
VM/ESA 2.1 to VM/ESA 2.2: FS8F0R; Report 22								
24	S46E5505	S47E550A	9121-742	1GB/1GB	90	5	1.026	
25	S46E5202	S47E5201	9121-742	320MB/64MB[5]	90	20	1.037	
26	L26E186I	L27E186J	9121-480	224MB/32MB[5]	90	16	1.026	
27	H16E0302	H17E0303	9221-170	48MB/16MB[5]	90	15	1.026	
avg							1.03	(R21)

Note: The report numbers refer to the list of VM performance reports on page 77.

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 29 on page 79.

The FS7B0R, FS7F0R, or FS8F0R workload (CMS-intensive, minidisks, remote users simulated by TPNS) was used for all comparisons except those involving

[5] These are the storage sizes used for the VM/ESA Release 2.1 measurements. For VM/ESA Release 2.2, the total storage size was the same but all of the expanded storage was reconfigured as real storage. This conforms to the usage guidelines for enhanced minidisk caching.

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.

New Functions

A number of the functional enhancements in VM/ESA Release 2.2 have performance implications. This section contains performance evaluation results for the most significant changes:

Enhanced Minidisk Cache

Scheduler Share Capping and Proportional Distribution

SPXTAPE

ISFC Changes

Improved CMS Block Allocation for Virtual Disk in Storage

Enhanced Minidisk Cache

Minidisk caching (MDC) was substantially enhanced in VM/ESA Release 2.2. These changes broaden MDC's scope of applicability, improve MDC performance, and provide better controls over the caching process. See "Enhanced Minidisk Cache" on page 5 for an overview of minidisk caching in VM/ESA Release 2.2 and how it compares to minidisk caching in prior releases.

Perhaps the most significant change is that minidisk caching has been extended to support non-CMS environments. This can result in significant performance benefits for guest operating systems. These benefits are illustrated by the performance measurement results provided in "VSE Guest Environment" on page 90.

Other improvements result in improved performance for CMS environments that already use minidisk caching. This is reflected in the CMS-intensive migration results provided in "CMS-Intensive" on page 22. Further performance gains can be realized in situations where more data becomes eligible for minidisk caching due to removed restrictions. Minidisk caching can now handle all CMS blocksizes and can be used on systems that do not have expanded storage.

Since minidisk caching has changed so much, this raises the question as to what system setup and tuning changes should be considered to ensure that enhanced minidisk caching is used effectively after migrating to VM/ESA Release 2.2. This question is addressed by the following usage guidelines and the measurement results provided in "CMS Environment" on page 111.

Guidelines: With the enhancements to the CP minidisk cache feature in VM/ESA Release 2.2, the following are potential items to consider when migrating from previous releases of VM/ESA or VM/XA.[6]

- Reconfigure expanded storage as main storage if it is specifically for minidisk caching.

  Some processors allow a portion of main storage to be configured as expanded storage. Installations that configure storage as expanded storage in order to do minidisk caching should consider reconfiguring this expanded storage as main storage.

- Review storage allocation for minidisk cache.

  Prior to VM/ESA Release 2.2, the amount of expanded storage allocated for minidisk cache could be fixed or bounded (minimum and maximum) via the RETAIN XSTORE MDC command. Now the amount of storage, either real or expanded, can be fixed or bounded via the SET MDCACHE command. The RETAIN XSTORE MDC command is maintained for compatibility purposes. When the amount of storage allocated for minidisk cache is not a fixed amount, the arbiter function of minidisk cache determines how much storage to allocate.

[6] Most of this information also appears in *VM/ESA Planning and Administration* in the chapter on "Expanded Storage Planning and Administration."

The arbiter was completely changed in VM/ESA Release 2.2. This eliminated some problems seen with the previous arbiter. If the arbiter was ineffective for your system in the past, it may now be acceptable and should be tried. If the size of the minidisk cache was previously fixed or bounded, you will want to either allow the arbiter to decide or you will want to alter the values in consideration of both real and expanded storage usage.

A new method of influencing the arbiter is setting a bias with the SET MDCACHE STORAGE BIAS or SET MDCACHE XSTORE BIAS commands. This can be preferable to setting a fixed minimum or maximum for systems where paging and/or minidisk cache size requirements vary significantly over time.

- Try letting the minidisk cache size default.

  For both MDC in real storage and MDC in expanded storage, the default is no cache size restrictions and no biasing (that is, a bias setting of 1.00). Under these conditions, measurement results to date show that the real storage arbiter usually strikes a good balance between using real storage for MDC and using it for demand paging. Likewise, the expanded storage arbiter usually makes a good trade-off between using expanded storage for MDC and using it for paging. The default settings should be good for most systems and are at least a good starting point.

  The one exception we have found is when storage is very unconstrained. In that case, we have observed that, under default conditions, both the real storage arbiter and the expanded storage arbiter tend to use too much storage for MDC. In the environments we looked at, bias settings in the range of 0.1 to 0.2 gave better results.

- When using expanded storage for minidisk caching, also use some real storage for minidisk caching.

  Measurement results indicate that in situations where expanded storage is to be used for most minidisk caching, it is better to allow some real storage to also be used for minidisk caching rather than setting the real MDC size to 0 or setting it off.

- Use SET MDCACHE or SET RDEVICE commands instead of SET SHARED to enable minidisk cache on volumes.

  Prior to VM/ESA Release 2.2, one could define all DASD volumes as shared volumes in the system gen or system configuration file and then enable particular volumes for minidisk cache via SET SHARE OFF command. With VM/ESA Release 2.2, one must explicitly use the SET MDCACHE RDEV command or the SET RDEVICE command with MDC option to enable minidisk cache.

- Enable caching for minidisks that were poor candidates in the past.

  There may be minidisks that were disabled for caching because most of the I/O to them were writes. It may now be reasonable to enable these minidisks because the writes will not cause inserts in the cache and the reads might benefit from caching. Some overhead to check for data in the cache when performing a write I/O (to maintain integrity) does exist. Therefore, disabling minidisk cache for write only minidisks would save some system resources.

- Disable caching for minidisks that are poor candidates.

  You may want to re-visit and change the devices for which you disabled or enabled minidisk caching. There may be some minidisks that are poor candidates for minidisk caching, but did not matter in the past since the type of I/O or format made them ineligible for minidisk caching. With several restrictions being lifted, it may be worthwhile to re-visit these minidisks and make sure they have minidisk caching disabled. VSE paging volumes may be candidates.

- Disable the minidisk cache fair share limit for key users.

  The minidisk cache feature uses a fair share concept which prevents any one virtual machine from flooding the cache storage by inserting more than its fair share. Prior to VM/ESA Release 2.2, it was recommended to use the directory option NOMDCFS (no minidisk cache fair share) to turn off the minidisk cache fair share limit for server machines such as SFS and SQL. With VM/ESA Release 2.2, guests such as VSE may also benefit from minidisk caching and you might want to add NOMDCFS to the directory entries of production guests.

- Reformat some minidisks to smaller blocksize.

  Each CMS file consists of at least one DASD record even if the file is very small. The capacity of 4K byte formatted minidisks that consist of mostly small files may be increased by reformatting them using 2K, 1K or 512 byte sizes. Prior to VM/ESA Release 2.2, reformatting to a smaller size would have made the minidisk ineligible for minidisk caching. Care must be taken, however, because as the record size gets smaller, the number of available bytes per track is reduced.

- Prepare for minidisk caching on devices shared between first and second level systems.

  Prior to VM/ESA Release 2.2, most second level systems did not use minidisk cache because it required expanded storage, which most second level test systems do not have. Now that minidisk caching can use real storage instead of expanded storage, second level systems benefit. Care should be used for minidisks shared with first level systems so that all changes are reflected to the second level system. For example, a minidisk is cached by a second level system and a change is made to the minidisk by the first level system. The change will not be reflected on the second level system if the data had been in the second level system′s cache. In order to see the changes, one must purge the cache on the second level system with the SET MDCACHE MDISK FLUSH command and reaccess the minidisk.

- Avoid mixing standard format and non-standard format records on the same cylinder.

  Mixing different format records on the same cylinder makes the minidisk cache feature less efficient. Minidisk caching is more efficient for standard format records. However, when both formats exist on the same cylinder, CP treats all I/O to that cylinder as non-standard.

VSE Guest Environment

This section examines the benefits of minidisk caching (MDC) in VM/ESA Release 2.2 when running the PACEX8 and VSECICS workloads. The graphs in the following sections illustrate performance improvements through the use of minidisk caching and through the use of virtual disk in storage for temporary work files.

Workload: PACEX8

Hardware Configuration

Processor models: 9121-320[7]
Storage
Real: 96MB-256MB[8]
Expanded: 0MB-32MB[8]

DASD:

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

Software Configuration

VSE version: 1.3.2

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
VSEVR	1	VSE V=R	16MB/ESA	100		IOASSIST ON CCWTRANS OFF
or						
VSEVV	1	VSE V=V	16MB/ESA	100		IOASSIST OFF
SMART	1	RTM	16MB/370	100		
WRITER	1	CP monitor	2MB/XA	100		

[7] See “Hardware Used” on page 17 for an explanation of how this processor model was defined.

[8] The amount of real and expanded storage was adjusted for specific test environments. Refer to the detailed data for specific measurement configurations.

Benefits as a Function of Minidisk Cache Size

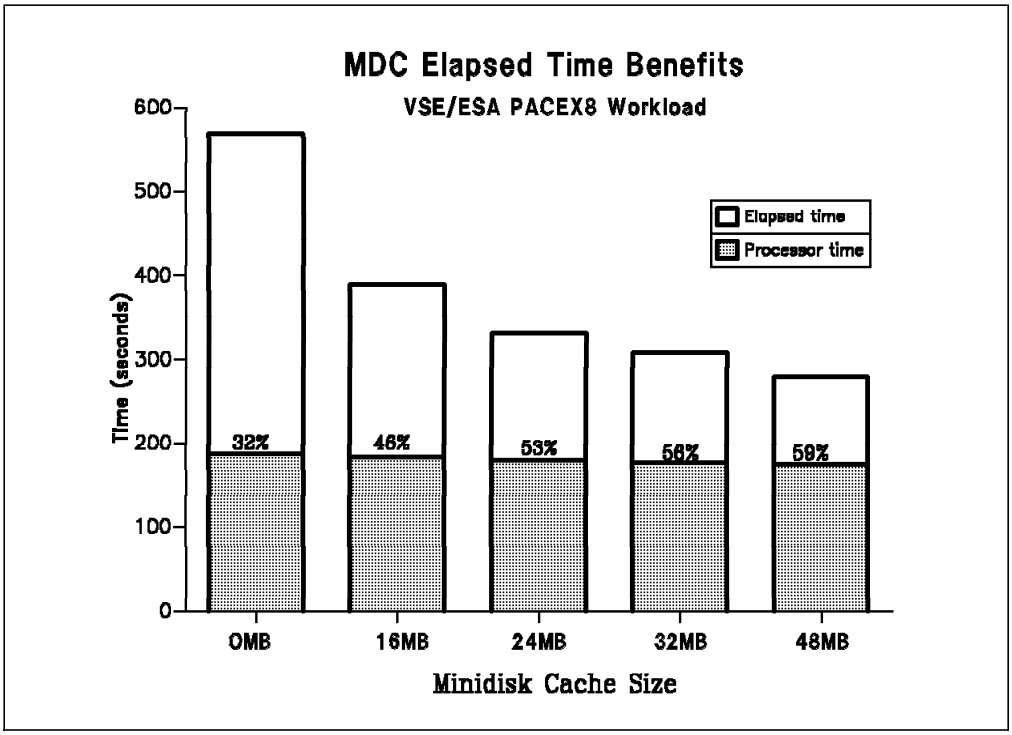

Figure 6. Elapsed time benefits from using minidisk caching. PACEX8 workload on a single VSE/ESA guest of VM/ESA Release 2.2 on the 9121-320 processor.

Figure 6 illustrates batch job elapsed run time reductions and processor utilization effects when the amount of storage allocated for MDC is increased. The initial measurement point did not use minidisk caching. The amount of storage for MDC was increased in steps to 48MB. Processor utilization nearly doubled across this range because of the elapsed time reduction.

The minidisk cache can either be set to a fixed size or the size can be defaulted, in which case CP dynamically determines the cache size based on storage demands and availability. In this series of measurements, the MDC was set to a fixed size. The guest mode was V=V with all guest DASD defined as full pack minidisks (denoted FPM in the following results tables). Overall, a 51% elapsed time reduction was achieved for the PACEX8 workload through the use of minidisk caching. The average PACEX8 job experienced a large elapsed time reduction because a large fraction of its original (pre-MDC) elapsed time was spent waiting for DASD read I/Os that are eligible for minidisk caching.

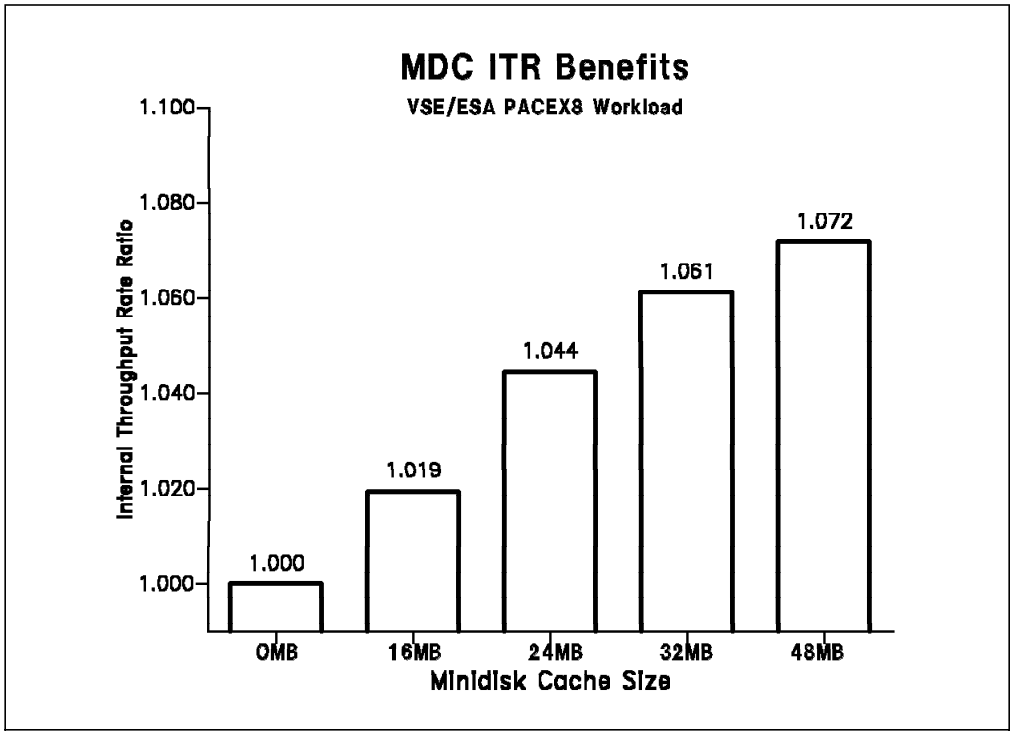

Figure 7. ITR benefit from using minidisk caching. PACEX8 workload on a single VSE/ESA guest of VM/ESA Release 2.2 on the 9121-320 processor.

Increasing the size of the MDC also resulted in significant ITR improvements for the V=V guest. Figure 7 illustrates the successive gains in ITR compared to the base measurement point where MDC was not used. The base measurement point ITR was normalized to a value of one. A 7% increase in internal throughput was achieved with the PACEX8 workload when the MDC size was set at 48MB.

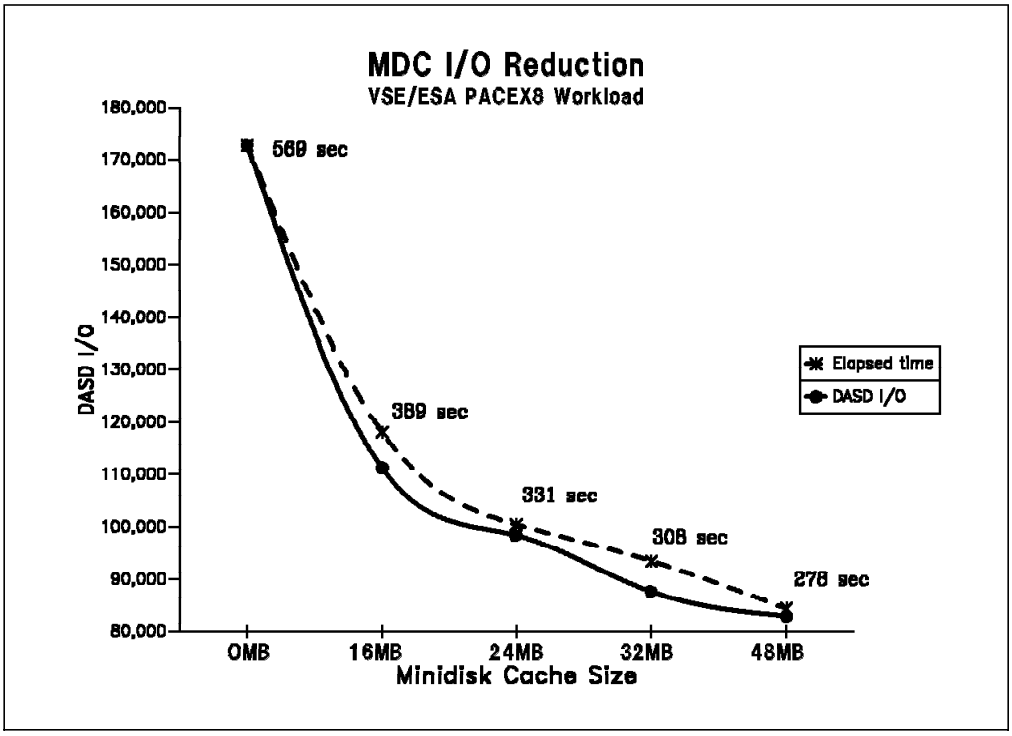

Figure 8. I/O benefits from using minidisk caching. PACEX8 workload on a single VSE/ESA guest of VM/ESA Release 2.2 on the 9121-320 processor.

Increasing the size of the MDC resulted in reductions in elapsed run time for the workload. Figure 8 shows the correlation between DASD I/O reduction by MDC and the effects on elapsed time. Since the PACEX8 workload is very I/O-intensive, this environment benefits significantly from DASD I/O reduction.

Table 31 (Page 1 of 2). VSE/ESA guest MDC benefits, PACEX8, VM/ESA 2.2 on the 9121-320

MDC Setting	**Off**	**16MB**	**24MB**	**32MB**	**48MB**
Release	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
Run ID	**LB78VEXC**	**LB78VEXD**	**LB78VEXF**	**LB78VEXE**	**LB78VEY6**
Environment					
Real Storage	**64MB**	**64MB**	**64MB**	**64MB**	**96MB**
Exp. Storage	**0MB**	**0MB**	**0MB**	**0MB**	**0MB**
VM Mode	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
VM Size	**16MB**	**16MB**	**16MB**	**16MB**	**16MB**
Guest Setting	**V=V**	**V=V**	**V=V**	**V=V**	**V=V**
Guest DASD	**FPM**	**FPM**	**FPM**	**FPM**	**FPM**
VSE Supervisor	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
Processors	**1**	**1**	**1**	**1**	**1**
Throughput (Min)					
Elapsed Time (C)	569.0	389.0	331.0	308.0	278.0
ETR (C)	5.91	8.64	10.15	10.91	12.08
ITR (H)	17.99	18.34	18.79	19.09	19.28
ITR	17.71	18.10	18.46	18.83	19.05
ITRR (H)	1.000	1.020	1.045	1.061	1.072
ITRR	1.000	1.022	1.043	1.064	1.076

Enhanced Minidisk Cache

Table 31 (Page 2 of 2). VSE/ESA guest MDC benefits, PACEX8, VM/ESA 2.2 on the 9121-320

MDC Setting	**Off**	**16MB**	**24MB**	**32MB**	**48MB**
Release	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
Run ID	**LB78VEXC**	**LB78VEXD**	**LB78VEXF**	**LB78VEXE**	**LB78VEY6**
Environment					
Real Storage	**64MB**	**64MB**	**64MB**	**64MB**	**96MB**
Exp. Storage	**0MB**	**0MB**	**0MB**	**0MB**	**0MB**
VM Mode	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
VM Size	**16MB**	**16MB**	**16MB**	**16MB**	**16MB**
Guest Setting	**V=V**	**V=V**	**V=V**	**V=V**	**V=V**
Guest DASD	**FPM**	**FPM**	**FPM**	**FPM**	**FPM**
VSE Supervisor	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
Processors	**1**	**1**	**1**	**1**	**1**
Proc. Usage (Sec)					
PBT/CMD (H)	3.336	3.272	3.193	3.144	3.112
PBT/CMD	3.388	3.315	3.250	3.186	3.150
CP/CMD (H)	1.236	1.229	1.169	1.135	1.117
CP/CMD	1.129	1.129	1.083	1.062	1.050
EMUL/CMD (H)	2.100	2.043	2.024	2.009	1.995
EMUL/CMD	2.259	2.187	2.166	2.124	2.100
Processor Util.					
TOTAL (H)	32.49	46.39	53.06	56.24	59.28
TOTAL	33.00	47.00	54.00	57.00	60.00
TOTAL EMUL (H)	20.45	28.96	33.63	35.94	38.01
TOTAL EMUL	22.00	31.00	36.00	38.00	40.00
TVR(H)	1.59	1.60	1.58	1.57	1.56
TVR	1.50	1.52	1.50	1.50	1.50
Storage					
NUCLEUS SIZE (V)	2544KB	2544KB	2544KB	2544KB	2544KB
TRACE TABLE (V)	200KB	200KB	200KB	200KB	200KB
PGBLPGS	14484	14467	14455	14456	22494
FREEPGS	72	75	75	76	75
FREE UTIL	0.45	0.46	0.48	0.48	0.47
SHRPGS	93	1053	297	375	582
Paging					
PAGE/CMD	0.000	0.000	0.000	0.000	5.250
XSTOR/CMD	0.000	0.000	0.000	0.000	0.000
FAST CLR/CMD	82.143	183.393	180.536	184.446	183.750
I/O					
VIO RATE	308.000	449.000	525.000	568.000	606.000
VIO/CMD	3162.500	3167.054	3159.375	3174.714	3181.500
RIO RATE (V)	320.000	266.000	274.000	293.000	278.000
RIO/CMD (V)	3285.714	1876.250	1648.893	1637.661	1459.500
DASD IO TOTAL (V)	172747	111311	98373	87662	82967
DASD IO RATE (V)	319.90	265.03	273.26	292.21	276.56
DASD IO/CMD (V)	3284.71	1869.38	1644.43	1633.23	1451.92
MDC REAL SIZE (MB)	0	16.0	23.1	31.6	47.8
MDC READS (I/Os)	0	255	298	323	344
MDC WRITES (I/Os)	0	139	163	176	189
MDC AVOID	1.00	167	233	280	321
MDC HIT RATIO	0	0.58	0.69	0.77	0.83
PRIVOPs					
PRIVOP/CMD	3167.310	3169.489	3163.861	3176.632	3182.351
DIAG/CMD	641.003	506.486	460.438	462.066	441.489
SIE/CMD	13687.054	13613.393	13594.339	13621.089	13644.750
SIE INTCPT/CMD	12592.089	12524.321	12506.792	12531.402	12553.170
FREE TOTL/CMD	3901.786	5480.625	4826.321	4504.964	3806.250

Note: V=VMPRF, H=Hardware Monitor, C=VSE Console Timestamps, Unmarked=RTM

Minidisk Caching Used with Virtual Disk in Storage

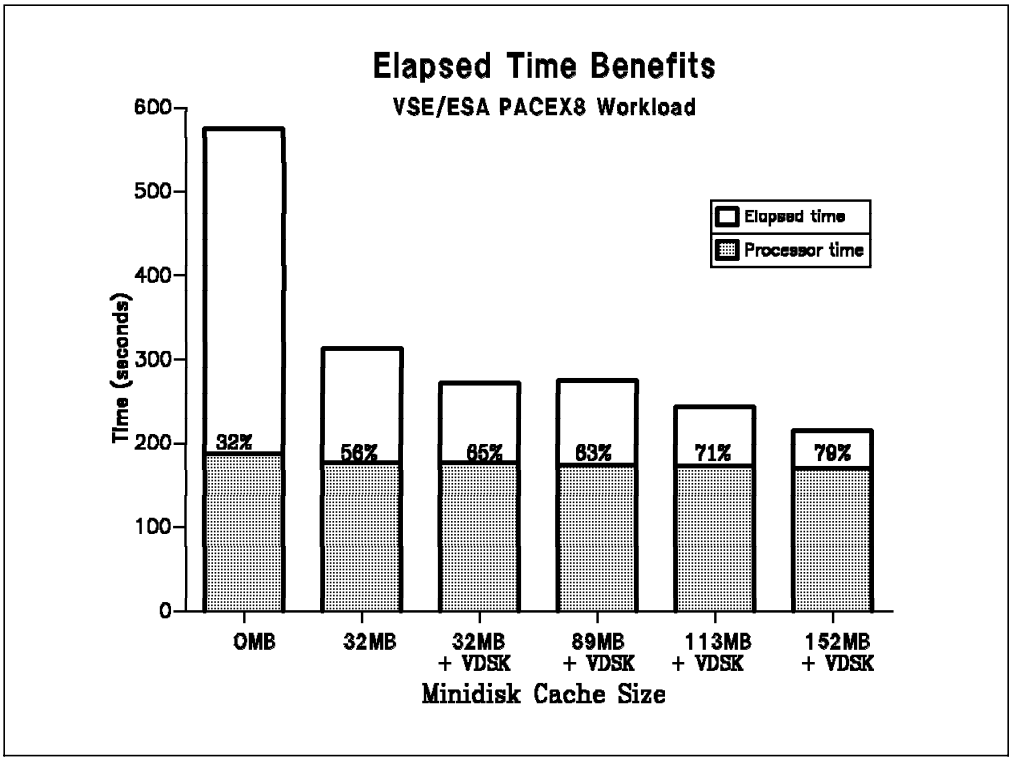

Figure 9. Combined benefits of using minidisk caching and virtual disk in storage. PACEX8 workload on a single VSE/ESA guest of VM/ESA Release 2.2 on the 9121-320 processor.

This series of measurements illustrates how the use of minidisk caching in conjunction with virtual disk in storage can benefit batch job run times. For all measurements, the guest mode was V=V with all guest DASD defined as full pack minidisks.

The initial measurement point in Figure 9 (left side of chart) shows the PACEX8 workload running without MDC. Total job run time for the workload was 569 seconds. In the next series of measurement points, MDC was introduced and storage allocations for MDC were gradually increased as well as utilizing virtual disk in storage for temporary batch work files.

The last three measurement points on the chart were run with the default MDC size (dynamically determined by CP). For those points, the MDC size shown is the average size, as determined by the real storage arbiter.

Overall, a 63% elapsed time reduction was achieved by utilizing MDC and virtual disk in storage. The internal throughput rate increased by over 10% across the range of measurements illustrated in Figure 9.

Enhanced Minidisk Cache

Table 32 (Page 1 of 2). MDC and virtual disk in storage benefits, PACEX8, VM/ESA 2.2 on the 9121-320

MDC Setting	**Off**	**32MB**	**32MB**	**0-125MB**	**0-157MB**	**0-221MB**
Virtual Disk in Storage	**not used**	**not used**	**yes**	**yes**	**yes**	**yes**
Release	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
Run ID	**LB78VEXC**	**LB78VEXE**	**LB78VEXO**	**LB78VEY2**	**LB78VEY3**	**LB78VEY1**
Environment						
Real Storage	**64MB**	**64MB**	**64MB**	**128MB**	**160MB**	**224MB**
Exp. Storage	**0MB**	**0MB**	**0MB**	**0MB**	**0MB**	**0MB**
VM Mode	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
VM Size	**16MB**	**16MB**	**16MB**	**16MB**	**16MB**	**16MB**
Guest Setting	**V=V**	**V=V**	**V=V**	**V=V**	**V=V**	**V=V**
Guest DASD	**FPM**	**FPM**	**FPM**	**FPM**	**FPM**	**FPM**
VSE Supervisor	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
Processors	**1**	**1**	**1**	**1**	**1**	**1**
Throughput (Min)						
Elapsed Time (C)	569.0	308.0	269.0	268.0	237.0	208.0
ETR (C)	5.91	10.91	12.49	12.53	14.17	16.15
ITR (H)	17.99	19.09	19.11	19.39	19.51	19.83
ITR	17.71	18.83	19.00	19.09	19.20	19.38
ITRR (H)	1.000	1.061	1.063	1.078	1.084	1.102
ITRR	1.000	1.064	1.073	1.078	1.085	1.095
Proc. Usage (Sec)						
PBT/CMD (H)	3.336	3.144	3.140	3.094	3.076	3.026
PBT/CMD	3.388	3.186	3.157	3.143	3.124	3.095
CP/CMD (H)	1.236	1.135	1.160	1.129	1.108	1.066
CP/CMD	1.129	1.062	1.069	1.031	1.041	0.994
EMUL/CMD (H)	2.100	2.009	1.980	1.965	1.968	1.960
EMUL/CMD	2.259	2.124	2.089	2.112	2.083	2.102
Processor Util.						
TOTAL (H)	32.49	56.24	64.64	63.00	70.89	79.20
TOTAL	33.00	57.00	65.00	64.00	72.00	81.00
TOTAL EMUL (H)	20.45	35.94	40.76	40.01	45.35	51.30
TOTAL EMUL	22.00	38.00	43.00	43.00	48.00	55.00
TVR(H)	1.59	1.57	1.59	1.57	1.56	1.54
TVR	1.50	1.50	1.51	1.49	1.50	1.47
Storage						
NUCLEUS SIZE (V)	2544KB	2544KB	2544KB	2544KB	2544KB	2544KB
TRACE TABLE (V)	200KB	200KB	200KB	200KB	200KB	200KB
PGBLPGS	14484	14456	14394	30631	38772	55071
FREEPGS	72	76	76	77	76	74
FREE UTIL	0.45	0.48	0.44	0.43	0.46	0.46
SHRPGS	93	375	109	116	306	1080
Paging						
PAGE/CMD	0.000	0.000	238.000	422.321	177.911	0.000
XSTOR/CMD	0.000	0.000	0.000	0.000	0.000	0.000
FAST CLR/CMD	82.143	184.446	272.000	181.696	190.929	191.071
I/O						
VIO RATE	308.000	568.000	626.000	631.000	716.000	815.000
VIO/CMD	3162.500	3174.714	3040.571	3098.661	3106.929	3114.464
RIO RATE (V)	320.000	293.000	306.000	284.000	292.000	265.000
RIO/CMD (V)	3285.714	1637.661	1486.286	1394.643	1267.071	1012.679
DASD IO TOTAL (V)	172747	87662	73174	67867	69893	63200
DASD IO RATE (V)	319.90	292.21	304.89	282.78	291.22	263.33
DASD IO/CMD (V)	3284.71	1633.23	1480.90	1388.65	1263.69	1006.31
MDC REAL SIZE (MB)	0	31.6	31.0	88.8	112.8	151.9
MDC READS (I/Os)	0	323	365	368	417	475
MDC WRITES (I/Os)	0	176	198	200	228	260
MDC AVOID	1.00	280	311	352	395	475
MDC HIT RATIO	0	0.77	0.80	0.90	0.89	0.94

Table 32 (Page 2 of 2). MDC and virtual disk in storage benefits, PACEX8, VM/ESA 2.2 on the 9121-320

MDC Setting	**Off**	**32MB**	**32MB**	**0-125MB**	**0-157MB**	**0-221MB**
Virtual Disk in Storage	**not used**	**not used**	**yes**	**yes**	**yes**	**yes**
Release	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
Run ID	**LB78VEXC**	**LB78VEXE**	**LB78VEXO**	**LB78VEY2**	**LB78VEY3**	**LB78VEY1**
Environment						
Real Storage	**64MB**	**64MB**	**64MB**	**128MB**	**160MB**	**224MB**
Exp. Storage	**0MB**	**0MB**	**0MB**	**0MB**	**0MB**	**0MB**
VM Mode	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
VM Size	**16MB**	**16MB**	**16MB**	**16MB**	**16MB**	**16MB**
Guest Setting	**V=V**	**V=V**	**V=V**	**V=V**	**V=V**	**V=V**
Guest DASD	**FPM**	**FPM**	**FPM**	**FPM**	**FPM**	**FPM**
VSE Supervisor	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
Processors	**1**	**1**	**1**	**1**	**1**	**1**
PRIVOPs						
PRIVOP/CMD	3167.310	3176.632	3031.302	3088.803	3094.927	3117.298
DIAG/CMD	641.003	462.066	411.416	419.023	396.654	377.224
SIE/CMD	13687.054	13621.089	13148.286	13366.964	13330.286	13355.893
SIE INTCPT/CMD	12592.089	12531.402	11964.940	12297.607	12263.863	12287.421
FREE TOTL/CMD	3901.786	4504.964	3914.857	3501.339	3549.536	3259.679

Note: V=VMPRF, H=Hardware Monitor, C=VSE Console Timestamps, Unmarked=RTM

Relationship to VSE Native

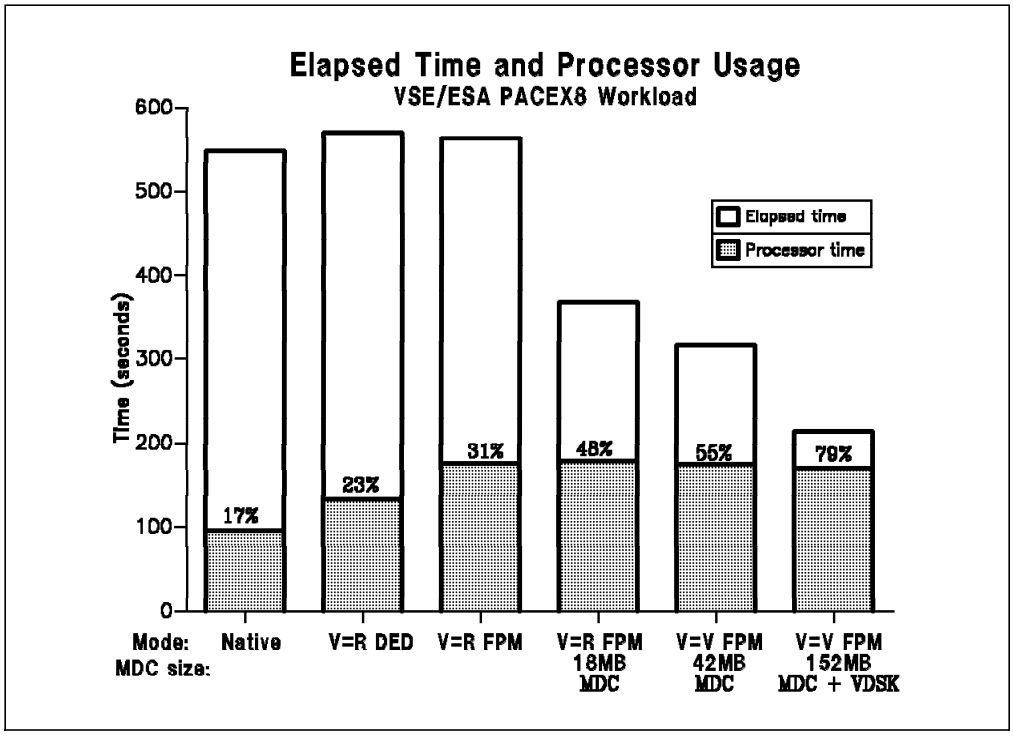

Figure 10. Comparison to VSE native. PACEX8 workload on a single VSE/ESA guest of VM/ESA Release 2.2 on the 9121-320 processor.

These measurements demonstrate how various guest run modes compare to VSE native running the same workload. In this chart, the first measurement point is VSE/ESA running native. The PACEX8 batch jobs completed in 545 seconds and processor utilization for this workload on the 9121-320 is very low at 17%.

The next point in the progression is the same workload, but running as a preferred guest with dedicated DASD (DED). As would be expected, the batch run time and processor utilization increased somewhat.

The third bar in the graph illustrates the V=R guest with DASD defined as full pack minidisks (FPM). With the DASD defined as full pack minidisks, the elapsed time is relatively unchanged when compared to the preferred guest with dedicated DASD. However, the processor utilization increased because the preferred guest without dedicated DASD is not receiving the benefit of I/O assist.

For the fourth measurement point, MDC was used with a range of 0MB to 29MB available (18MB average size). The introduction of MDC decreased batch run time over the native mode measurement by nearly 33%. However, processor time and processor utilization were significantly higher. As the figure shows, the increase in processor time and the decrease in I/O wait time both contributed to the increased processor utilization.

In the fifth point of the measurement progression, the guest mode was switched from V=R to V=V, and the size of the MDC was increased. Increasing the size of the MDC provided further benefits in elapsed time reduction.

The last point in the chart shows a further reduction in elapsed run time by increasing the size of the MDC and using virtual disk in storage for the temporary work files.

Overall, a reduction of 61% of the batch job run time was achieved by utilizing MDC and virtual disk in storage when compared to VSE native running the same workload. At the same time, however, processor utilization increased from 17% to 79% because of the increased processor time and decreased I/O wait time. These results illustrate that minidisk caching and virtual disks in storage can be especially valuable in situations where it is important to minimize the elapsed time of I/O-intensive VSE batch workloads.

Table 33 (Page 1 of 2). Various VSE guest cases relative to VSE/ESA native, PACEX8, VM/ESA 2.2 on the 9121-320

MDC Setting	**not used**	**not used**	**not used**	**0-29MB**	**0-61MB**	**0-221MB**
Virtual Disk in Storage	**not used**	**not used**	**not used**	**not used**	**not used**	**yes**
Release	**not used**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
Run ID	**LBNAT132**	**LB78REX1**	**LB78REX6**	**LB78REX2**	**LB78VEXN**	**LB78VEY1**
Environment						
Real Storage	**64MB**	**64MB**	**64MB**	**64MB**	**64MB**	**224MB**
Exp. Storage	**0MB**	**0MB**	**0MB**	**0MB**	**0MB**	**0MB**
VM Mode	**na**	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
VM Size	**na**	**16MB**	**16MB**	**16MB**	**16MB**	**16MB**
Guest Setting	**na**	**V=R**	**V=R**	**V=R**	**V=V**	**V=V**
Guest DASD	**na**	**dedicated**	**FPM**	**FPM**	**FPM**	**FPM**
VSE Supervisor	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
Processors	**1**	**1**	**1**	**1**	**1**	**1**
Throughput (Min)						
Elapsed Time (C)	545.0	566.0	562.0	361.0	314.0	208.0
ETR (C)	6.16	5.94	5.98	9.31	10.70	16.15
ITR (H)	35.46	25.38	19.14	18.93	19.21	19.83
ITR	na	25.58	19.22	18.68	19.33	19.38
ITRR (H)	1.000	0.716	0.540	0.534	0.542	0.559
Proc. Usage (Sec)						
PBT/CMD (H)	1.69	2.364	3.135	3.170	3.123	3.026
PBT/CMD	na	2.345	3.122	3.211	3.104	3.095
CP/CMD (H)	na	0.386	1.075	1.148	1.134	1.066
CP/CMD	na	0.306	1.007	1.049	1.016	0.994
EMUL/CMD (H)	na	1.978	2.060	2.022	1.989	1.960
EMUL/CMD	na	2.039	2.115	2.163	2.088	2.102
Processor Util.						
TOTAL (H)	17.39	23.19	31.13	48.37	55.35	79.20
TOTAL	na	23.00	31.00	49.00	55.00	81.00
TOTAL EMUL (H)	na	19.40	20.45	30.86	35.25	51.30
TOTAL EMUL	na	20.00	21.00	33.00	37.00	55.00
TVR(H)	na	1.20	1.52	1.57	1.57	1.54
TVR	na	1.15	1.48	1.48	1.49	1.47
Storage						
NUCLEUS SIZE (V)	na	2544KB	2532KB	2544KB	2544KB	2544KB
TRACE TABLE (V)	na	200KB	200KB	200KB	200KB	200KB
PGBLPGS	na	6331	6332	6294	14443	55071
FREEPGS	na	69	68	73	76	74
FREE UTIL	na	0.41	0.44	0.45	0.46	0.46
SHRPGS	na	1052	1053	37	213	1080

Enhanced Minidisk Cache

Table 33 (Page 2 of 2). Various VSE guest cases relative to VSE/ESA native, PACEX8, VM/ESA 2.2 on the 9121-320

MDC Setting	**not used**	**not used**	**not used**	**0-29MB**	**0-61MB**	**0-221MB**
Virtual Disk in Storage	**not used**	**not used**	**not used**	**not used**	**not used**	**yes**
Release	**not used**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
Run ID	**LBNAT132**	**LB78REX1**	**LB78REX6**	**LB78REX2**	**LB78VEXN**	**LB78VEY1**
Environment						
Real Storage	**64MB**	**64MB**	**64MB**	**64MB**	**64MB**	**224MB**
Exp. Storage	**0MB**	**0MB**	**0MB**	**0MB**	**0MB**	**0MB**
VM Mode	**na**	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
VM Size	**na**	**16MB**	**16MB**	**16MB**	**16MB**	**16MB**
Guest Setting	**na**	**V=R**	**V=R**	**V=R**	**V=V**	**V=V**
Guest DASD	**na**	**dedicated**	**FPM**	**FPM**	**FPM**	**FPM**
VSE Supervisor	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
Processors	**1**	**1**	**1**	**1**	**1**	**1**
Paging						
PAGE/CMD	na	0.000	0.000	0.000	62.071	0.000
XSTOR/CMD	na	0.000	0.000	0.000	0.000	0.000
FAST CLR/CMD	na	0.000	0.000	0.000	180.571	191.071
I/O						
VIO RATE	na	5.000	312.000	483.000	555.000	815.000
VIO/CMD	na	50.982	3142.286	3165.375	3131.786	3114.464
RIO RATE (V)	na	1.000	318.000	317.000	263.000	265.000
RIO/CMD (V)	na	10.196	3202.714	2077.482	1484.071	1012.679
DASD IO TOTAL (V)	na	167661	171305	94810	78659	63200
DASD IO RATE (V)	na	310.48	317.23	316.03	262.20	263.33
DASD IO/CMD (V)	na	3165.82	3194.97	2071.15	1479.54	1006.31
MDC REAL SIZE (MB)	na	0	0	18.2	41.5	151.9
MDC READS (I/Os)	na	0.03	0	275	315	475
MDC WRITES (I/Os)	na	0.01	0	149	173	260
MDC AVOID	na	0.01	0	191	288	475
MDC HIT RATIO	na	0.50	0	0.62	0.81	0.94
PRIVOPs						
PRIVOP/CMD (R)	na	55.730	3152.180	3170.412	3136.571	3117.298
DIAG/CMD (R)	na	636.315	627.709	484.748	456.946	377.224
SIE/CMD	na	3456.589	10444.071	10361.196	13435.643	13355.893
SIE INTCPT/CMD	na	3110.930	9608.546	9532.301	12360.791	12287.421
FREE TOTL/CMD	na	978.857	7110.429	8316.482	4102.357	3259.679

Note: V=VMPRF, H=Hardware Monitor, C=VSE Console Timestamps, Unmarked=RTM

Real Storage Arbiter

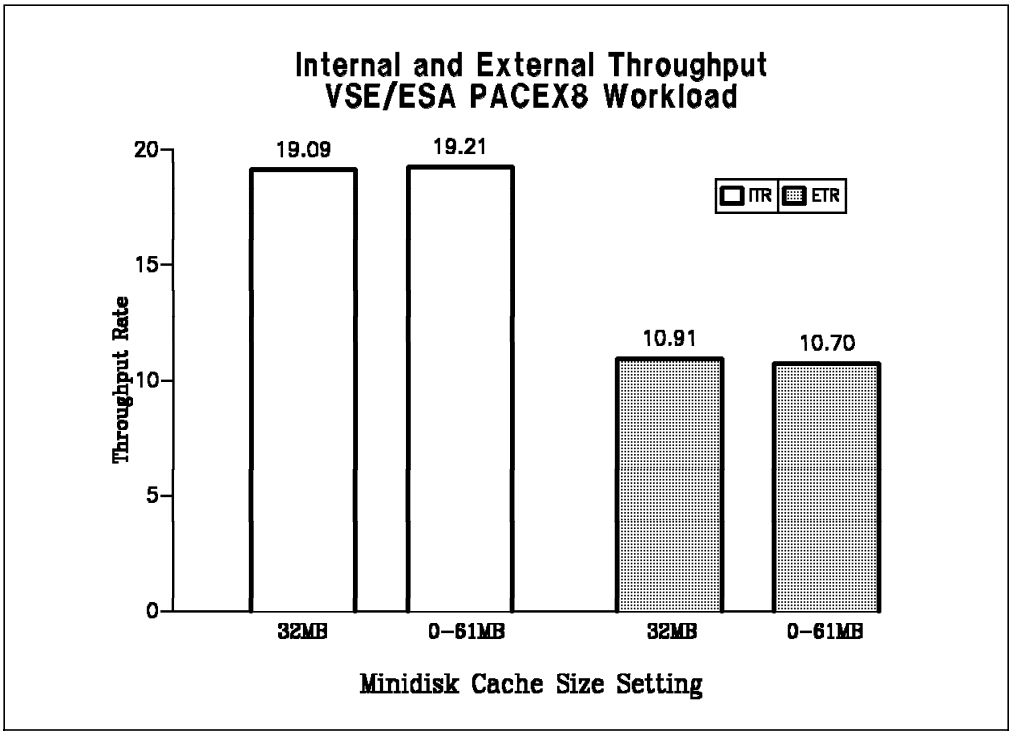

Figure 11. MDC fixed cache size vs. variable setting. PACEX8 workload on a single VSE/ESA guest of VM/ESA Release 2.2 on the 9121-320 processor.

VM/ESA Release 2.2 provides controls that allow the setting of the MDC to be a fixed size or allowed to default to a variable range. The PACEX8 workload was run as a V=V guest with a 32MB fixed size MDC and then with the default size. The variable range setting was 0MB - 61MB. The MDC for either mode used real storage only (no expanded storage allocated). In both measurements, the effective real storage size configured was 64MB.

The purpose of this experiment was to verify the efficiency of the real storage arbiter. Through a series of previous tests, it was determined that a fixed MDC size of 32MB was a favorable setting for the PACEX8 workload. When the MDC size was allowed to be variable, the real storage arbiter increased the MDC size beyond 32MB to 41.5MB. By increasing the MDC size to 41.5MB, the arbiter balanced the additional benefits of reduced DASD I/O with some degree of paging. The net result was a 9% decrease in real I/Os (RIO/CMD (V)) and similar overall performance. These results indicate that the real storage arbiter was able to find a good balance between using real storage for demand paging and using it for MDC.

Enhanced Minidisk Cache

Table 34 (Page 1 of 2). Real storage arbiter evaluation, PACEX8, VM/ESA 2.2 on the 9121-320

MDC Setting **Release** **Run ID**	**32MB** **VM/ESA 2.2** **LB78VEXE**	**0-61MB** **VM/ESA 2.2** **LB78VEXN**	**Difference**	**%Difference**
Environment				
Real Storage	**64MB**	**64MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**16MB**	**16MB**		
Guest Setting	**V=V**	**V=V**		
Guest DASD	**FPM**	**FPM**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (C)	308.0	314.0	6.0	1.95%
ETR (C)	10.91	10.70	-0.21	-1.96%
ITR (H)	19.09	19.21	0.12	0.65%
ITR	18.83	19.33	0.50	2.65%
ITRR (H)	1.000	1.007	0.007	0.65%
ITRR	1.000	1.027	0.027	2.65%
Proc. Usage (Sec)				
PBT/CMD (H)	3.144	3.123	-0.020	-0.65%
PBT/CMD	3.186	3.104	-0.082	-2.58%
CP/CMD (H)	1.135	1.134	-0.001	-0.10%
CP/CMD	1.062	1.016	-0.046	-4.36%
EMUL/CMD (H)	2.009	1.989	-0.019	-0.96%
EMUL/CMD	2.124	2.088	-0.036	-1.70%
Processor Util.				
TOTAL (H)	56.24	55.35	-0.90	-1.59%
TOTAL	57.00	55.00	-2.00	-3.51%
TOTAL EMUL (H)	35.94	35.25	-0.68	-1.90%
TOTAL EMUL	38.00	37.00	-1.00	-2.63%
TVR(H)	1.57	1.57	0.00	0.32%
TVR	1.50	1.49	-0.01	-0.90%
Storage				
NUCLEUS SIZE (V)	2544KB	2544KB	0KB	0.00%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	14456	14443	-13	-0.09%
FREEPGS	76	76	0	0.00%
FREE UTIL	0.48	0.46	-0.01	-2.76%
SHRPGS	375	213	-162	-43.20%
Paging				
PAGE/CMD	0.000	62.071	62.071	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	184.446	180.571	-3.875	-2.10%
I/O				
VIO RATE	568.000	555.000	-13.000	-2.29%
VIO/CMD	3174.714	3131.786	-42.929	-1.35%
RIO RATE (V)	293.000	263.000	-30.000	-10.24%
RIO/CMD (V)	1637.661	1484.071	-153.589	-9.38%
DASD IO TOTAL (V)	87662	78659	-9003	-10.27%
DASD IO RATE (V)	292.21	262.20	-30.01	-10.27%
DASD IO/CMD (V)	1633.23	1479.54	-153.69	-9.41%
MDC REAL SIZE (MB)	31.6	41.5	9.9	31.33%
MDC READS (I/Os)	323	315	-8	-2.48%
MDC WRITES (I/Os)	176	173	-3	-1.70%
MDC AVOID	280	288	8	2.86%
MDC HIT RATIO	0.77	0.81	0.04	5.19%

Table 34 (Page 2 of 2). Real storage arbiter evaluation, PACEX8, VM/ESA 2.2 on the 9121-320

MDC Setting **Release** **Run ID**	**32MB** **VM/ESA 2.2** **LB78VEXE**	**0-61MB** **VM/ESA 2.2** **LB78VEXN**	**Difference**	**%Difference**
Environment				
Real Storage	**64MB**	**64MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**16MB**	**16MB**		
Guest Setting	**V=V**	**V=V**		
Guest DASD	**FPM**	**FPM**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
PRIVOPs				
PRIVOP/CMD	3176.632	3136.571	-40.060	-1.26%
DIAG/CMD	462.066	456.946	-5.119	-1.11%
SIE/CMD	13621.089	13435.643	-185.446	-1.36%
SIE INTCPT/CMD	12531.402	12360.791	-170.611	-1.36%
FREE TOTL/CMD	4504.964	4102.357	-402.607	-8.94%

Note: V=VMPRF, H=Hardware Monitor, C=VSE Console Timestamps, Unmarked=RTM

Minidisk Caching in Expanded Storage

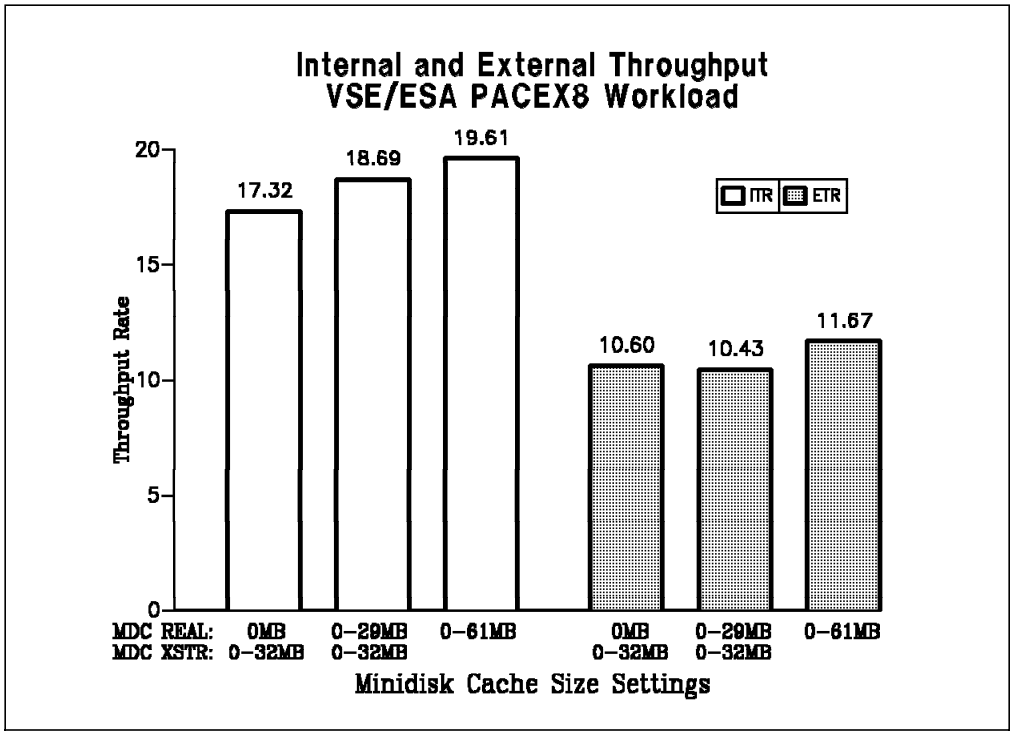

Figure 12. Real and expanded storage MDC combinations. PACEX8 workload on a single VSE/ESA guest of VM/ESA Release 2.2 on the 9121-320 processor.

Minidisk cache can be defined in real storage, expanded storage, or both. The PACEX8 workload was run as a V=V guest with MDC defined in three combinations of real and expanded storage:

- MDC in expanded storage only
- MDC in real and expanded storage
- MDC in real only

The total processor storage size was 64MB for all three measurements.

Figure 12 illustrates the throughput differences for the MDC combinations listed. When MDC was defined in real and expanded storage vs. expanded storage only, a 7.9% improvement in internal throughput was noted. By placing all MDC in real storage, an additional 4.9% gain in internal throughput occurred.

Using MDC in expanded storage only or a combination of real plus expanded, resulted in near equivalent external throughput rates. Allocating MDC entirely in real storage resulted in a 10% gain in external throughput compared to MDC in expanded storage.

Recommendations:

- Use MDC in real storage exclusively when practical.
- If expanded storage is to be used for MDC, some real storage should also be allocated. Usually, the best way to do this is to let the MDC real storage size default (MDC size determined by the real storage arbiter).

Table 35 (Page 1 of 2). MDC in expanded storage, PACEX8, VM/ESA 2.2 on the 9121-320

MDC setting (real)	**0MB**	**0-29MB**	**0-61MB**
MDC setting (exp)	**0-32MB**	**0-32MB**	**na**
Release	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
Run ID	**LB78VEXL**	**LB78VEXI**	**LB78VEXH**
Environment			
Real Storage	**32MB**	**32MB**	**64MB**
Exp. Storage	**32MB**	**32MB**	**0MB**
VM Mode	**ESA**	**ESA**	**ESA**
VM Size	**16MB**	**16MB**	**16MB**
Guest Setting	**V=V**	**V=V**	**V=V**
Guest DASD	**FPM**	**FPM**	**FPM**
VSE Supervisor	**ESA**	**ESA**	**ESA**
Processors	**1**	**1**	**1**
Throughput (Min)			
Elapsed Time (C)	317.0	323.0	288.0
ETR (C)	10.60	10.43	11.67
ITR (H)	17.32	18.69	19.61
ITR	17.21	18.80	19.64
ITRR (H)	1.000	1.079	1.132
ITRR	1.000	1.092	1.141
Proc. Usage (Sec)			
PBT/CMD (H)	3.464	3.210	3.059
PBT/CMD	3.486	3.192	3.055
CP/CMD (H)	1.463	1.225	1.092
CP/CMD	1.371	1.103	0.984
EMUL/CMD (H)	2.001	1.985	1.967
EMUL/CMD	2.114	2.089	2.071
Processor Util.			
TOTAL (H)	60.62	55.32	59.07
TOTAL	61.00	55.00	59.00
TOTAL EMUL (H)	35.02	34.20	37.97
TOTAL EMUL	37.00	36.00	40.00
TVR(H)	1.73	1.62	1.56
TVR	1.65	1.53	1.48
Storage			
NUCLEUS SIZE (V)	2536KB	2536KB	2536KB
TRACE TABLE (V)	200KB	200KB	200KB
PGBLPGS	6313	6310	14467
FREEPGS	87	84	84
FREE UTIL	0.51	0.54	0.53
SHRPGS	1052	92	215
Paging			
PAGE/CMD	0.000	69.643	15.536
XSTOR/CMD	0.000	40.625	0.000
FAST CLR/CMD	177.143	179.911	176.071

Enhanced Minidisk Cache

Table 35 (Page 2 of 2). MDC in expanded storage, PACEX8, VM/ESA 2.2 on the 9121-320

MDC setting (real)	**0MB**	**0-29MB**	**0-61MB**
MDC setting (exp)	**0-32MB**	**0-32MB**	**na**
Release	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
Run ID	**LB78VEXL**	**LB78VEXI**	**LB78VEXH**
Environment			
Real Storage	**32MB**	**32MB**	**64MB**
Exp. Storage	**32MB**	**32MB**	**0MB**
VM Mode	**ESA**	**ESA**	**ESA**
VM Size	**16MB**	**16MB**	**16MB**
Guest Setting	**V=V**	**V=V**	**V=V**
Guest DASD	**FPM**	**FPM**	**FPM**
VSE Supervisor	**ESA**	**ESA**	**ESA**
Processors	**1**	**1**	**1**
I/O			
VIO RATE	550.000	541.000	608.000
VIO/CMD	3142.857	3139.732	3148.571
RIO RATE (V)	246.000	254.000	273.000
RIO/CMD (V)	1405.714	1474.107	1413.750
DASD IO TOTAL (V)	88122	91005	81652
DASD IO RATE (V)	244.78	252.79	272.17
DASD IO/CMD (V)	1398.76	1467.09	1409.47
MDC REAL SIZE (MB)	1.2	17.6	42.2
MDC XSTOR SIZE (MB)	28.2	29.8	0
MDC READS (I/Os)	313	307	346
MDC WRITES (I/Os)	168	167	190
MDC AVOID	277	264	327
MDC HIT RATIO	0.79	0.76	0.84
PRIVOPs			
PRIVOP/CMD	3148.518	3141.821	3152.567
DIAG/CMD	452.768	454.018	431.990
SIE/CMD	13394.286	13441.071	13469.464
SIE INTCPT/CMD	12322.743	12365.786	12391.907
FREE TOTL/CMD	5994.286	4787.946	3966.786

Note: V=VMPRF, H=Hardware Monitor, C=VSE Console Timestamps, Unmarked=RTM

VSECICS Workload Results

Workload: VSECICS

Hardware Configuration

Processor models: 9121-320[9]

Storage

Real: 128MB

Expanded: 0MB

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*VSAM*	*VSE Sys.*	*VM Sys.*
3380-A	3880-03	2				18	2	2

Software Configuration

VSE version: 1.3.2

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
VSEVR	1	VSE V=R	32MB/ESA	100		IOASSIST ON CCWTRANS OFF
or						
VSEVV	1	VSE V=V	32MB/ESA	100		IOASSIST OFF
SMART	1	RTM	16MB/370	100		
WRITER	1	CP monitor	2MB/XA	100		

[9] See “Hardware Used” on page 17 for an explanation of how this processor model was defined.

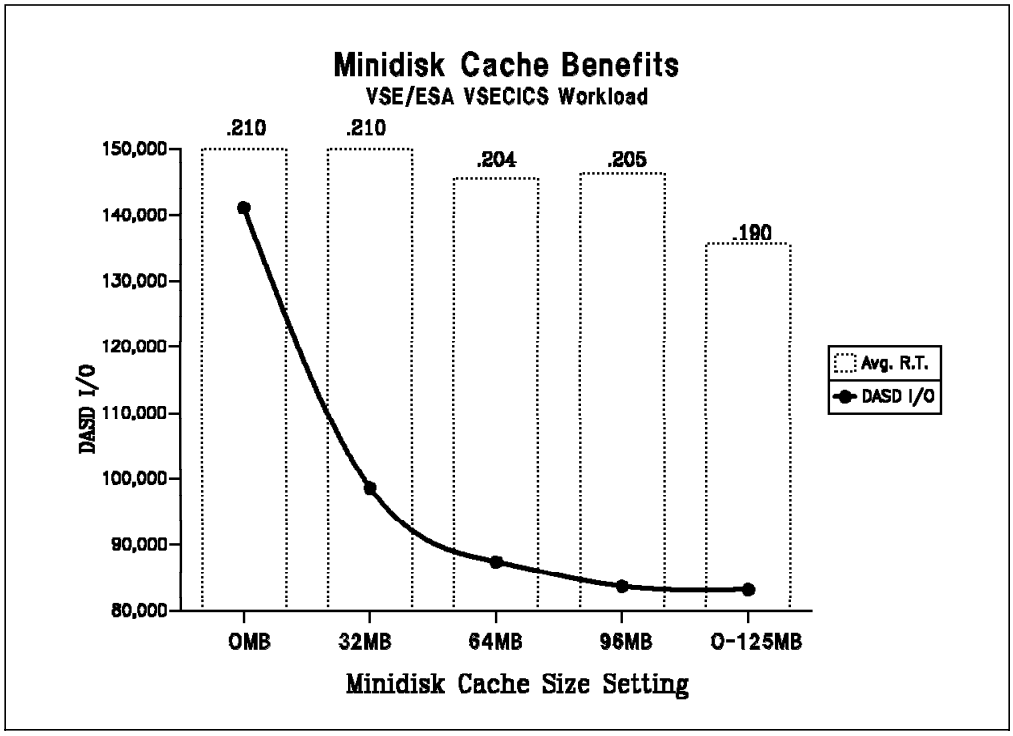

Figure 13. VSE guest MDC benefits of VM/ESA Release 2.2. VSECICS workload on a single VSE/ESA guest of VM/ESA Release 2.2 on the 9121-320 processor.

In this series of measurement points, the storage allocated to MDC was increased in steps. The I/O rate dropped substantially when MDC was introduced. However, this had a limited effect on performance because, for this workload, waiting for DASD I/O is a relatively small component of overall response time.

Table 36 (Page 1 of 2). VSE/ESA guest MDC benefits, VSECICS, VM/ESA 2.2 on the 9121-320

MDC Setting	**Off**	**32MB**	**64MB**	**96MB**	**0-125MB**
RELEASE	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
RUN ID	**VVF3207D**	**VV13207C**	**VV13207D**	**VV13207E**	**VV13207F**
Environment					
REAL STORAGE	**128M**	**128M**	**128M**	**128M**	**128M**
EXP. STORAGE	**0M**	**0M**	**0M**	**0M**	**0M**
USERS	**960**	**950**	**950**	**950**	**950**
VM MODE	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
VM SIZE	**32MB**	**32MB**	**32MB**	**32MB**	**32MB**
GUEST SETTING	**V=V**	**V=V**	**V=V**	**V=V**	**V=V**
GUEST DASD	**FPM**	**FPM**	**FPM**	**FPM**	**FPM**
VSE MODE	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
PROCESSORS	**1**	**1**	**1**	**1**	**1**
Response Time					
AVG RESP (C)	0.210	0.210	0.204	0.205	0.190
Throughput					
ETR (C)	48.68	48.14	48.21	48.23	48.24
ITR (H)	69.09	68.08	69.38	69.64	69.64
ITRR (H)	1.000	0.985	1.004	1.008	1.008
Proc. Usage					
PBT/CMD (H)	14.473	14.689	14.414	14.360	14.359
PBT/CMD	14.380	14.750	14.312	14.307	14.302
CP/CMD (H)	2.079	2.272	2.164	2.152	2.144
CP/CMD	1.849	2.077	1.867	1.866	1.865
EMUL/CMD (H)	12.394	12.417	12.250	12.209	12.214
EMUL/CMD	12.531	12.673	12.445	12.441	12.437
Processor Util.					
TOTAL (H)	70.45	70.71	69.49	69.26	69.27
TOTAL	70.00	71.00	69.00	69.00	69.00
TOTAL EMUL (H)	60.33	59.77	59.06	58.88	58.93
TOTAL EMUL	61.00	61.00	60.00	60.00	60.00
TVR(H)	1.17	1.18	1.18	1.18	1.18
TVR	1.15	1.16	1.15	1.15	1.15
Storage					
NUCLEUS SIZE (V)	2532K	2544K	2544K	2544K	2544K
TRACE TABLE (V)	200K	200K	200K	200K	200K
PGBLPGS	23232	30761	30743	30745	30743
FREEPGS	71	78	75	74	74
FREE UTIL	0.49	0.47	0.49	0.46	0.49
SHRPGS	119	1080	1080	50	37
Paging					
PAGE/CMD	0.000	0.000	0.000	0.062	0.000
XSTOR/CMD	0.000	0.000	0.000	0.000	0.000
FAST CLR/CMD	3.143	3.220	3.111	3.131	3.109
I/O					
VIO RATE	156.000	160.000	155.000	153.000	154.000
VIO/CMD	3.205	3.324	3.215	3.173	3.192
RIO RATE (V)	157.000	110.000	97.000	93.000	93.000
RIO/CMD (V)	3.225	2.285	2.012	1.928	1.928
MDC REAL SIZE (MB)	0	31.3	63.1	95.4	89.5
MDC READS (I/Os)	0	71	70	69	70
MDC WRITES (I/Os)	0	55	55	53	53
MDC AVOID	0	50	58	62	62
MDC HIT RATIO	0	0.70	0.82	0.88	0.87

Enhanced Minidisk Cache

Table 36 (Page 2 of 2). VSE/ESA guest MDC benefits, VSECICS, VM/ESA 2.2 on the 9121-320

MDC Setting	**Off**	**32MB**	**64MB**	**96MB**	**0-125MB**
RELEASE	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
RUN ID	**VVF3207D**	**VV13207C**	**VV13207D**	**VV13207E**	**VV13207F**
Environment					
REAL STORAGE	**128M**	**128M**	**128M**	**128M**	**128M**
EXP. STORAGE	**0M**	**0M**	**0M**	**0M**	**0M**
USERS	**960**	**950**	**950**	**950**	**950**
VM MODE	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
VM SIZE	**32MB**	**32MB**	**32MB**	**32MB**	**32MB**
GUEST SETTING	**V=V**	**V=V**	**V=V**	**V=V**	**V=V**
GUEST DASD	**FPM**	**FPM**	**FPM**	**FPM**	**FPM**
VSE MODE	**ESA**	**ESA**	**ESA**	**ESA**	**ESA**
PROCESSORS	**1**	**1**	**1**	**1**	**1**
PRIVOPs					
PRIVOP/CMD	3.218	3.325	3.220	3.188	3.201
DIAG/CMD	4.038	4.120	4.023	4.032	4.019
SIE/CMD	23.378	23.953	23.356	23.286	23.257
SIE INTCPT/CMD	16.131	17.246	17.050	16.999	16.977
FREE TOTL/CMD	4.232	5.464	4.791	4.603	4.622

Note: V=VMPRF, H=Hardware Monitor, C=CICSPARS, Unmarked=RTM. Per command (/CMD), for this workload, means per CICS transaction.

CMS Environment

This section discusses using minidisk caching (MDC) in VM/ESA Release 2.2 when running the CMS workload FS8F. Comparisons are shown for using real storage in place of, or in conjunction with, expanded storage for MDC. Measurements showing the effects of using an option to bias the arbiter against MDC are discussed. The last set of measurements in this section shows the advantage of using cached control units and MDC together.

Minidisk Caching in Real Storage

The following three sets of measurements show the advantage of using real storage in place of, or in conjunction with, expanded storage for MDC. The first comparison shows real storage being used instead of expanded storage by reconfiguring storage. The second comparison shows real storage being used along with expanded storage. The last comparison shows using only real storage instead of both real and expanded storage.

On a 9121-480, the processor storage can be partitioned between real and expanded storage. A base measurement was obtained with the 256MB of processor storage partitioned as 224MB real storage and 32MB expanded storage. This is compared to an equivalent measurement that was run with all 256MB as real storage. This comparison illustrates the improved performance that is potentially available when processor storage is used as real storage instead of expanded storage.

On the 9121-742, the measurements were made without changing the hardware configuration of storage, but by changing how the various storage were used by MDC. A measurement with MDC using just expanded storage is compared to a measurement where both real and expanded storage were used. Again, the results show that performance can be improved by using real storage for MDC, this time in conjunction with minidisk caching in expanded storage. When only real storage was used for MDC and expanded storage was used only for paging, the performance did not improve significantly.

Minidisk cache is more efficient when real storage is used instead of expanded storage for various reasons. Moving the data from expanded storage to real storage can be avoided when caching in real storage. While granularity of storage configured as expanded storage can be large (16MB on some systems), caching in real storage can be specified in increments of a megabyte (by the SET MDCACHE command). Finer granularity for MDC cache size can be achieved by allowing the arbiter to dynamically adjust the cache size. This allows storage to be used for paging or MDC, whichever is more beneficial at any given time. Because of these factors, it is usually beneficial to configure processor storage as real storage.

Enhanced Minidisk Cache

Workload: FS8F0R

Hardware Configuration

Processor model: 9121-480
Processors used: 2
Storage, base run:
  Real: 224MB
  Expanded: 32MB (all for MDC)
Storage, comparison run:
  Real: 256MB (default MDC)
  Expanded: 0MB
Tape: 3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2	16 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*
3088-08	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XC	10000	512	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1860	Users	3MB/XC	100		

Measurement Discussion: When expanded storage is configured as real storage, improved system performance is evident by the 8.7% improvement in response time (AVG LAST(T)) and the 1.4% improvement in internal throughput rate (ITR(H)). These improvements result from the greater efficiency of caching in real storage and a larger available cache size. Note that while total amount of storage used for MDC increased, there was a decrease in paging (PAGE IO/CMD(V)).

Table 37 (Page 1 of 2). Benefits of using real storage for minidisk caching on the 9121-480

MDC setting (real) **MDC setting (exp)** **Release** **Run ID**	**0MB** **32MB** **VM/ESA 2.2** **L27E186D**	**0-254MB** **na** **VM/ESA 2.2** **L27E186A**	**Difference**	**%Difference**
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.130	0.128	-0.002	-1.54%
NONTRIV INT	0.433	0.397	-0.036	-8.31%
TOT INT	0.331	0.307	-0.024	-7.25%
TOT INT ADJ	0.292	0.269	-0.022	-7.69%
AVG FIRST (T)	0.250	0.234	-0.016	-6.59%
AVG LAST (T)	0.355	0.325	-0.031	-8.72%
Throughput				
AVG THINK (T)	26.20	26.18	-0.03	-0.10%
ETR	57.63	57.30	-0.33	-0.57%
ETR (T)	65.40	65.33	-0.07	-0.10%
ETR RATIO	0.881	0.877	-0.004	-0.47%
ITR (H)	74.32	75.36	1.04	1.40%
ITR	32.75	33.06	0.31	0.95%
EMUL ITR	47.98	47.99	0.01	0.02%
ITRR (H)	1.000	1.014	0.014	1.40%
ITRR	1.000	1.009	0.009	0.95%
Proc. Usage				
PBT/CMD (H)	26.910	26.538	-0.372	-1.38%
PBT/CMD	26.911	26.479	-0.431	-1.60%
CP/CMD (H)	9.117	8.824	-0.293	-3.22%
CP/CMD	8.562	8.265	-0.297	-3.47%
EMUL/CMD (H)	17.792	17.714	-0.078	-0.44%
EMUL/CMD	18.348	18.214	-0.134	-0.73%
Processor Util.				
TOTAL (H)	176.00	173.38	-2.61	-1.48%
TOTAL	176.00	173.00	-3.00	-1.70%
UTIL/PROC (H)	88.00	86.69	-1.31	-1.48%
UTIL/PROC	88.00	86.50	-1.50	-1.70%
TOTAL EMUL (H)	116.37	115.73	-0.63	-0.54%
TOTAL EMUL	120.00	119.00	-1.00	-0.83%
MASTER TOTAL (H)	87.67	86.25	-1.42	-1.62%
MASTER TOTAL	88.00	86.00	-2.00	-2.27%
MASTER EMUL (H)	51.45	50.95	-0.50	-0.97%
MASTER EMUL	53.00	53.00	0.00	0.00%
TVR(H)	1.51	1.50	-0.01	-0.94%
TVR	1.47	1.45	-0.01	-0.88%
Storage				
NUCLEUS SIZE (V)	2536KB	2536KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	84	84	0	0.00%
PGBLPGS	47277	55497	8220	17.39%
PGBLPGS/USER	25.4	29.8	4.4	17.39%
FREEPGS	5332	5247	-85	-1.59%
FREE UTIL	0.96	0.93	-0.03	-3.46%
SHRPGS	1254	1281	27	2.15%

Enhanced Minidisk Cache

Table 37 (Page 2 of 2). Benefits of using real storage for minidisk caching on the 9121-480

MDC setting (real) **MDC setting (exp)** **Release** **Run ID**	**0MB** **32MB** **VM/ESA 2.2** **L27E186D**	**0-254MB** **na** **VM/ESA 2.2** **L27E186A**	**Difference**	**%Difference**
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	625	628	3	0.48%
WRITES/SEC	436	433	-3	-0.69%
PAGE/CMD	16.223	16.240	0.017	0.10%
PAGE IO RATE (V)	177.600	174.100	-3.500	-1.97%
PAGE IO/CMD (V)	2.716	2.665	-0.051	-1.87%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	8.792	8.770	-0.021	-0.24%
Queues				
DISPATCH LIST	37.69	39.07	1.38	3.67%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	665	664	-1	-0.15%
VIO/CMD	10.168	10.163	-0.005	-0.05%
RIO RATE (V)	396	362	-34	-8.59%
RIO/CMD (V)	6.055	5.541	-0.514	-8.49%
NONPAGE RIO/CMD (V)	3.339	2.876	-0.463	-13.88%
DASD RESP TIME (V)	17.900	19.700	1.800	10.06%
MDC REAL SIZE (MB)	0.2	39.1	38.9	19450.00%
MDC XSTOR SIZE (MB)	31.4	0	-31.4	-100.00%
MDC READS (I/Os)	185	190	5	2.70%
MDC WRITES (I/Os)	7.84	9.45	1.61	20.54%
MDC AVOID	164	179	15	9.15%
MDC HIT RATIO	0.86	0.94	0.08	9.30%
PRIVOPs				
PRIVOP/CMD	13.933	13.904	-0.029	-0.21%
DIAG/CMD	27.631	27.709	0.078	0.28%
DIAG 04/CMD	2.461	2.471	0.011	0.44%
DIAG 08/CMD	0.737	0.738	0.001	0.10%
DIAG 0C/CMD	1.126	1.126	0.001	0.07%
DIAG 14/CMD	0.025	0.025	0.000	0.44%
DIAG 58/CMD	1.248	1.249	0.001	0.08%
DIAG 98/CMD	1.144	1.168	0.024	2.06%
DIAG A4/CMD	3.619	3.616	-0.003	-0.07%
DIAG A8/CMD	2.834	2.817	-0.017	-0.60%
DIAG 214/CMD	13.292	13.350	0.058	0.44%
SIE/CMD	53.989	53.541	-0.449	-0.83%
SIE INTCPT/CMD	35.093	34.801	-0.292	-0.83%
FREE TOTL/CMD	53.194	49.653	-3.541	-6.66%
VTAM Machines				
WKSET (V)	540	542	2	0.37%
TOT CPU/CMD (V)	3.8735	3.8520	-0.0215	-0.56%
CP CPU/CMD (V)	1.4611	1.4541	-0.0070	-0.48%
VIRT CPU/CMD (V)	2.4124	2.3979	-0.0145	-0.60%
DIAG 98/CMD (V)	1.145	1.167	0.022	1.90%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Workload: FS8F0R

Hardware Configuration

Processor model:	9121-742
Processors used:	4
Storage	
Real:	1024MB
Expanded:	1024MB (64MB minimum for MDC)
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-3	4	6	7	7	32 R		2 R
3390-2	3990-2	4	16	6	6			

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	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	500	QUICKDSP ON
VSCSn	3	VSCS	64MB/XA	10000	1200	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	550	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	5500	Users	3MB/XC	100		

Measurement Discussion: For these measurements, a system with 1GB of real storage and 1GB of expanded storage was used. A measurement with MDC turned off for real storage is compared to one with MDC turned on for real storage (default size). In both measurements, expanded storage was used for MDC with a 64 to 1024MB minimum and maximum setting.

The use of real storage in addition to expanded storage for MDC improved system performance. There was a 10.7% improvement in response time (AVG LAST(T)) and a 0.5% improvement in internal throughput rate (ITR(H)). There were slightly more I/Os avoided because of the extra real storage for MDC. However, paging I/Os increased (PAGE IO/CMD(V)).

Enhanced Minidisk Cache

These measurements were done with the default bias setting (1.0) for the expanded storage arbiter. The results for both measurements would have been better if the expanded storage arbiter had been biased against minidisk cache (see the following section for details).

Table 38 (Page 1 of 3). Combined use of real and expanded storage for MDC on the 9121-742

MDC setting (real) **MDC setting (exp)** **Release** **Run ID**	**Off** **64MB-1GB** **VM/ESA 2.2** **S47E5501**	**0MB-1015MB** **64MB-1GB** **VM/ESA 2.2** **S47E5500**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Response Time				
TRIV INT	0.113	0.126	0.013	11.50%
NONTRIV INT	0.368	0.371	0.003	0.82%
TOT INT	0.262	0.276	0.014	5.34%
TOT INT ADJ	0.260	0.260	-0.001	-0.23%
AVG FIRST (T)	0.297	0.263	-0.034	-11.44%
AVG LAST (T)	0.384	0.343	-0.041	-10.68%
Throughput				
AVG THINK (T)	26.12	26.21	0.10	0.37%
ETR	191.79	181.87	-9.92	-5.17%
ETR (T)	192.95	193.18	0.23	0.12%
ETR RATIO	0.994	0.941	-0.053	-5.29%
ITR (H)	213.20	214.27	1.07	0.50%
ITR	53.00	50.47	-2.54	-4.79%
EMUL ITR	81.05	76.80	-4.25	-5.24%
ITRR (H)	1.000	1.005	0.005	0.50%
ITRR	1.000	0.952	-0.048	-4.79%
Proc. Usage				
PBT/CMD (H)	18.762	18.668	-0.094	-0.50%
PBT/CMD	18.762	18.687	-0.075	-0.40%
CP/CMD (H)	6.901	6.809	-0.092	-1.33%
CP/CMD	6.478	6.419	-0.060	-0.92%
EMUL/CMD (H)	11.861	11.860	-0.002	-0.01%
EMUL/CMD	12.283	12.268	-0.015	-0.12%
Processor Util.				
TOTAL (H)	362.01	360.64	-1.37	-0.38%
TOTAL	362.00	361.00	-1.00	-0.28%
UTIL/PROC (H)	90.50	90.16	-0.34	-0.38%
UTIL/PROC	90.50	90.25	-0.25	-0.28%
TOTAL EMUL (H)	228.86	229.11	0.25	0.11%
TOTAL EMUL	237.00	237.00	0.00	0.00%
MASTER TOTAL (H)	92.67	92.32	-0.35	-0.38%
MASTER TOTAL	93.00	92.00	-1.00	-1.08%
MASTER EMUL (H)	37.54	37.25	-0.29	-0.78%
MASTER EMUL	39.00	39.00	0.00	0.00%
TVR(H)	1.58	1.57	-0.01	-0.49%
TVR	1.53	1.52	0.00	-0.28%

Table 38 (Page 2 of 3). Combined use of real and expanded storage for MDC on the 9121-742

MDC setting (real) **MDC setting (exp)** **Release** **Run ID**	**Off** **64MB-1GB** **VM/ESA 2.2** **S47E5501**	**0MB-1015MB** **64MB-1GB** **VM/ESA 2.2** **S47E5500**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Storage				
NUCLEUS SIZE (V)	2536KB	2536KB	0KB	0.00%
TRACE TABLE (V)	800KB	800KB	0KB	0.00%
WKSET (V)	78	78	0	0.00%
PGBLPGS	233k	233k	0k	0.00%
PGBLPGS/USER	42.4	42.4	0.0	0.00%
FREEPGS	15755	15574	-181	-1.15%
FREE UTIL	0.93	0.92	-0.01	-0.61%
SHRPGS	1819	1761	-58	-3.19%
Paging				
READS/SEC	1100	1177	77	7.00%
WRITES/SEC	862	926	64	7.42%
PAGE/CMD	10.169	10.886	0.718	7.06%
PAGE IO RATE (V)	325.800	349.300	23.500	7.21%
PAGE IO/CMD (V)	1.689	1.808	0.120	7.08%
XSTOR IN/SEC	531	473	-58	-10.92%
XSTOR OUT/SEC	1573	1586	13	0.83%
XSTOR/CMD	10.904	10.658	-0.246	-2.26%
FAST CLR/CMD	8.857	8.836	-0.021	-0.24%
Queues				
DISPATCH LIST	113.84	108.86	-4.98	-4.37%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	1801	1807	6	0.33%
VIO/CMD	9.334	9.354	0.020	0.21%
RIO RATE (V)	697	719	22	3.16%
RIO/CMD (V)	3.612	3.722	0.110	3.03%
NONPAGE RIO/CMD (V)	1.924	1.914	-0.010	-0.52%
DASD RESP TIME (V)	20.300	20.700	0.400	1.97%
MDC REAL SIZE (MB)	0.2	30.4	30.2	15100.00%
MDC XSTOR SIZE (MB)	207.0	199.2	-7.8	-3.79%
MDC READS (I/Os)	557	560	3	0.54%
MDC WRITES (I/Os)	26	27	1	3.85%
MDC AVOID	523	527	4	0.76%
MDC HIT RATIO	0.93	0.93	0.00	0.00%

Enhanced Minidisk Cache

Table 38 (Page 3 of 3). Combined use of real and expanded storage for MDC on the 9121-742

MDC setting (real) **MDC setting (exp)** **Release** **Run ID**	**Off** **64MB-1GB** **VM/ESA 2.2** **S47E5501**	**0MB-1015MB** **64MB-1GB** **VM/ESA 2.2** **S47E5500**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
PRIVOPs				
PRIVOP/CMD	20.507	20.568	0.062	0.30%
DIAG/CMD	25.305	25.345	0.041	0.16%
DIAG 04/CMD	0.933	0.944	0.011	1.22%
DIAG 08/CMD	0.737	0.737	0.000	0.00%
DIAG 0C/CMD	1.126	1.125	-0.001	-0.08%
DIAG 14/CMD	0.025	0.025	0.000	-0.05%
DIAG 58/CMD	1.249	1.248	0.000	-0.02%
DIAG 98/CMD	0.324	0.338	0.014	4.29%
DIAG A4/CMD	3.601	3.609	0.008	0.23%
DIAG A8/CMD	2.822	2.822	0.000	0.00%
DIAG 214/CMD	13.339	13.349	0.011	0.08%
SIE/CMD	57.010	56.941	-0.069	-0.12%
SIE INTCPT/CMD	38.197	38.151	-0.046	-0.12%
FREE TOTL/CMD	46.406	45.299	-1.107	-2.39%
VTAM Machines				
WKSET (V)	1490	1484	-6	-0.40%
TOT CPU/CMD (V)	2.7555	2.7608	0.0053	0.19%
CP CPU/CMD (V)	1.2208	1.2251	0.0043	0.35%
VIRT CPU/CMD (V)	1.5347	1.5357	0.0010	0.07%
DIAG 98/CMD (V)	0.324	0.338	0.014	4.46%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Workload: FS8F0R

Hardware Configuration

Processor model:	9121-742
Processors used:	4
Storage	
Real:	1024MB
Expanded:	1024MB
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-3	4	6	7	7	32 R		2 R
3390-2	3990-2	4	16	6	6			

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	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	500	QUICKDSP ON
VSCSn	3	VSCS	64MB/XA	10000	1200	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	550	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	5500	Users	3MB/XC	100		

Measurement Discussion: For these measurements, a system with 1GB of real storage and 1GB of expanded storage was used. Both real and expanded storage were used for MDC in the first measurement, while only real storage was used in the second. In both measurements, there was additional tuning. The first measurement used a MDC expanded storage bias of 0.1 and the second measurement used a MDC real storage bias of 0.2.

When only using real storage for MDC, the performance was not significantly different.

Enhanced Minidisk Cache

Table 39 (Page 1 of 2). Using only real storage for MDC on the 9121-742

MDC setting (real) **MDC setting (exp)** **Release** **Run ID**	**0MB-1015MB** **0MB-1GB** **VM/ESA 2.2** **S47E550A**	**0MB-1015MB** **Off** **VM/ESA 2.2** **S47E5509**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Response Time				
TRIV INT	0.108	0.107	-0.001	-0.93%
NONTRIV INT	0.339	0.337	-0.002	-0.59%
TOT INT	0.249	0.247	-0.002	-0.80%
TOT INT ADJ	0.236	0.234	-0.002	-0.77%
AVG FIRST (T)	0.240	0.241	0.001	0.42%
AVG LAST (T)	0.320	0.320	0.000	-0.10%
Throughput				
AVG THINK (T)	26.09	26.11	0.02	0.06%
ETR	182.96	182.92	-0.04	-0.02%
ETR (T)	193.13	193.02	-0.10	-0.05%
ETR RATIO	0.947	0.948	0.000	0.03%
ITR (H)	218.41	218.97	0.55	0.25%
ITR	51.78	51.93	0.15	0.29%
EMUL ITR	77.56	77.71	0.15	0.20%
ITRR (H)	1.000	1.003	0.003	0.25%
ITRR	1.000	1.003	0.003	0.29%
Proc. Usage				
PBT/CMD (H)	18.314	18.268	-0.046	-0.25%
PBT/CMD	18.278	18.236	-0.042	-0.23%
CP/CMD (H)	6.480	6.459	-0.021	-0.33%
CP/CMD	6.058	6.061	0.003	0.05%
EMUL/CMD (H)	11.834	11.809	-0.025	-0.21%
EMUL/CMD	12.220	12.175	-0.045	-0.37%
Processor Util.				
TOTAL (H)	353.69	352.61	-1.08	-0.31%
TOTAL	353.00	352.00	-1.00	-0.28%
UTIL/PROC (H)	88.42	88.15	-0.27	-0.31%
UTIL/PROC	88.25	88.00	-0.25	-0.28%
TOTAL EMUL (H)	228.55	227.94	-0.60	-0.26%
TOTAL EMUL	236.00	235.00	-1.00	-0.42%
MASTER TOTAL (H)	90.81	90.54	-0.27	-0.30%
MASTER TOTAL	91.00	90.00	-1.00	-1.10%
MASTER EMUL (H)	39.08	38.99	-0.09	-0.23%
MASTER EMUL	41.00	41.00	0.00	0.00%
TVR(H)	1.55	1.55	0.00	-0.04%
TVR	1.50	1.50	0.00	0.14%
Storage				
NUCLEUS SIZE (V)	2556KB	2556KB	0KB	0.00%
TRACE TABLE (V)	800KB	800KB	0KB	0.00%
WKSET (V)	72	72	0	0.00%
PGBLPGS	233k	233k	0k	0.00%
PGBLPGS/USER	42.4	42.4	0.0	0.00%
FREEPGS	15848	15595	-253	-1.60%
FREE UTIL	0.92	0.92	0.00	-0.16%
SHRPGS	1760	1800	40	2.27%

Table 39 (Page 2 of 2). Using only real storage for MDC on the 9121-742

MDC setting (real) **MDC setting (exp)** **Release** **Run ID**	**0MB-1015MB** **0MB-1GB** **VM/ESA 2.2** **S47E550A**	**0MB-1015MB** **Off** **VM/ESA 2.2** **S47E5509**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Paging				
READS/SEC	604	571	-33	-5.46%
WRITES/SEC	448	416	-32	-7.14%
PAGE/CMD	5.447	5.113	-0.334	-6.13%
PAGE IO RATE (V)	158.400	147.000	-11.400	-7.20%
PAGE IO/CMD (V)	0.820	0.762	-0.059	-7.15%
XSTOR IN/SEC	844	870	26	3.08%
XSTOR OUT/SEC	1413	1404	-9	-0.64%
XSTOR/CMD	11.687	11.781	0.094	0.81%
FAST CLR/CMD	8.761	8.755	-0.006	-0.06%
Queues				
DISPATCH LIST	98.85	99.76	0.91	0.92%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	1802	1800	-2	-0.11%
VIO/CMD	9.331	9.325	-0.005	-0.06%
RIO RATE (V)	542	518	-24	-4.43%
RIO/CMD (V)	2.806	2.684	-0.123	-4.38%
NONPAGE RIO/CMD (V)	1.986	1.922	-0.064	-3.23%
DASD RESP TIME (V)	20.000	20.300	0.300	1.50%
MDC REAL SIZE (MB)	33.0	39.1	6.1	18.48%
MDC XSTOR SIZE (MB)	43.0	0.0	-43.0	-100.00%
MDC READS (I/Os)	550	552	2	0.36%
MDC WRITES (I/Os)	27	26	-1	-3.70%
MDC AVOID	508	519	11	2.17%
MDC HIT RATIO	0.91	0.93	0.02	2.20%
PRIVOPs				
PRIVOP/CMD	20.563	20.544	-0.019	-0.09%
DIAG/CMD	25.217	25.224	0.007	0.03%
DIAG 04/CMD	0.931	0.932	0.000	0.05%
DIAG 08/CMD	0.739	0.737	-0.002	-0.27%
DIAG 0C/CMD	1.125	1.126	0.001	0.05%
DIAG 14/CMD	0.025	0.025	0.000	-0.65%
DIAG 58/CMD	1.248	1.248	0.000	0.01%
DIAG 98/CMD	0.344	0.353	0.009	2.60%
DIAG A4/CMD	3.567	3.569	0.002	0.05%
DIAG A8/CMD	2.829	2.820	-0.009	-0.32%
DIAG 214/CMD	13.261	13.268	0.007	0.05%
SIE/CMD	56.958	56.988	0.030	0.05%
SIE INTCPT/CMD	38.162	38.182	0.020	0.05%
FREE TOTL/CMD	44.919	44.461	-0.458	-1.02%
VTAM Machines				
WKSET (V)	4144	4144	0	0.00%
TOT CPU/CMD (V)	2.7673	2.7803	0.0130	0.47%
CP CPU/CMD (V)	1.2312	1.2405	0.0093	0.76%
VIRT CPU/CMD (V)	1.5361	1.5398	0.0037	0.24%
DIAG 98/CMD (V)	0.344	0.352	0.008	2.47%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Biasing the Expanded Storage Arbiter

The measurements discussed in this section show that performance can be improved by biasing the expanded storage arbiter against MDC in very unconstrained storage environments. Two series of measurements were made: one in an unconstrained environment and one in a constrained environment.

The unconstrained series used a 9121-742 with 1GB real storage and 1GB expanded storage. In this series, the first measurement was made with the default bias setting (1.0), while the second and third were made with expanded storage bias settings of 0.2 and 0.1, respectively. The smaller bias settings improved performance in this environment.

The constrained series used the same 9121-742 with 312MB real storage and 512MB expanded storage. For this series, the first measurement was made with the default bias setting (1.0), the second measurement with 0.5, and the third with 0.1. Unlike the unconstrained environment, performance did not improve with the smaller bias settings.

Workload: FS8F0R

Hardware Configuration

Processor model:	9121-742
Processors used:	4
Storage	
Real:	1024MB (default MDC)
Expanded:	1024MB (MDC size: 0MB to 1024MB)
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-3	4	6	7	7	32 R		2 R
3390-2	3990-2	4	16	6	6			

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	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	500	QUICKDSP ON
VSCSn	3	VSCS	64MB/XA	10000	1200	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	550	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	5500	Users	3MB/XC	100		

Measurement Discussion: In the unconstrained environment, response time (AVG LAST(T)) improved 16.6% for the 0.2 bias and 21.2% for the 0.1 bias. Internal throughput (ITR(H)) improved 2.2% and 2.5%, respectively. While biasing the expanded storage arbiter against minidisk cache caused a decrease in I/O avoided, this was more than offset by a decrease in paging I/O, as can be seen in Figure 14.

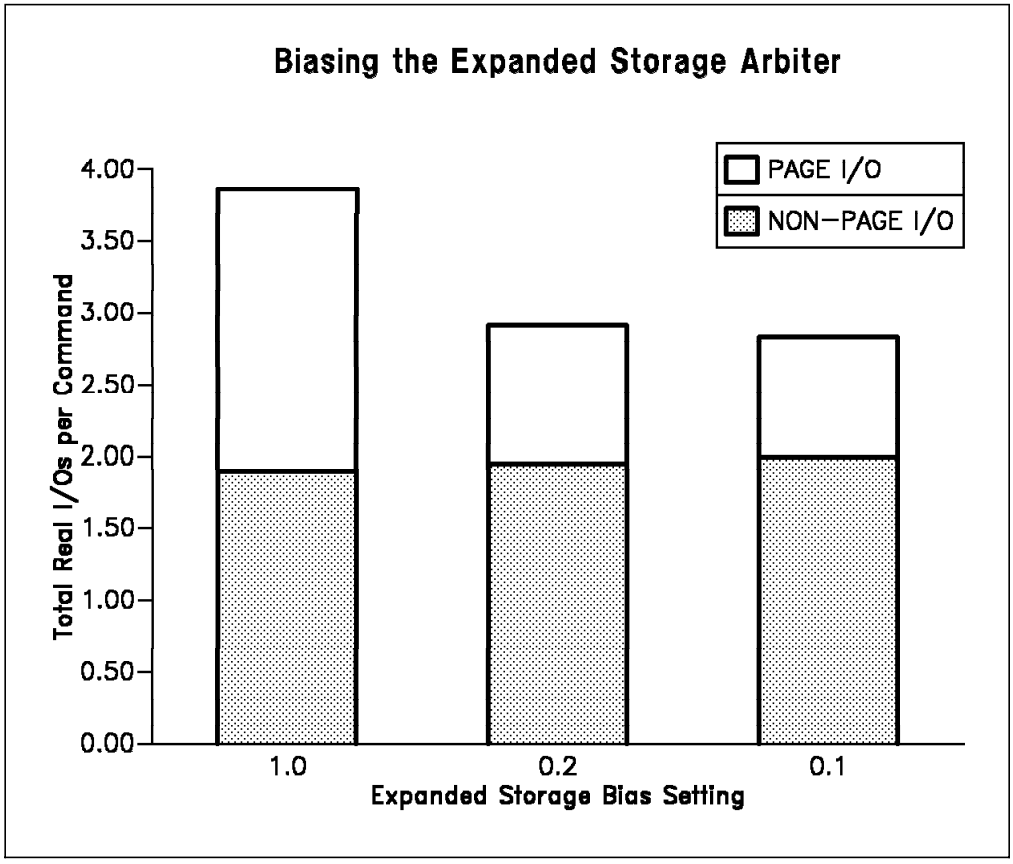

Figure 14. Biasing the Expanded Storage Arbiter. Real I/O per command with various minidisk cache bias settings for expanded storage.

Figure 15 shows the bias setting's effect on the size of the MDC expanded storage cache. A graph such as this can be used to estimate the bias setting required to reach a desired average cache size. Setting the bias to 0 effectively turns off MDC for the given type of storage. This is the basis for the first point on the graph. The three remaining points were determined by plotting the average expanded storage cache size versus the expanded storage bias setting for the three measurements. Because the relationship between the bias setting and the storage used is close to being linear, linear interpolation between 0 and a system's current bias setting should be sufficiently accurate.

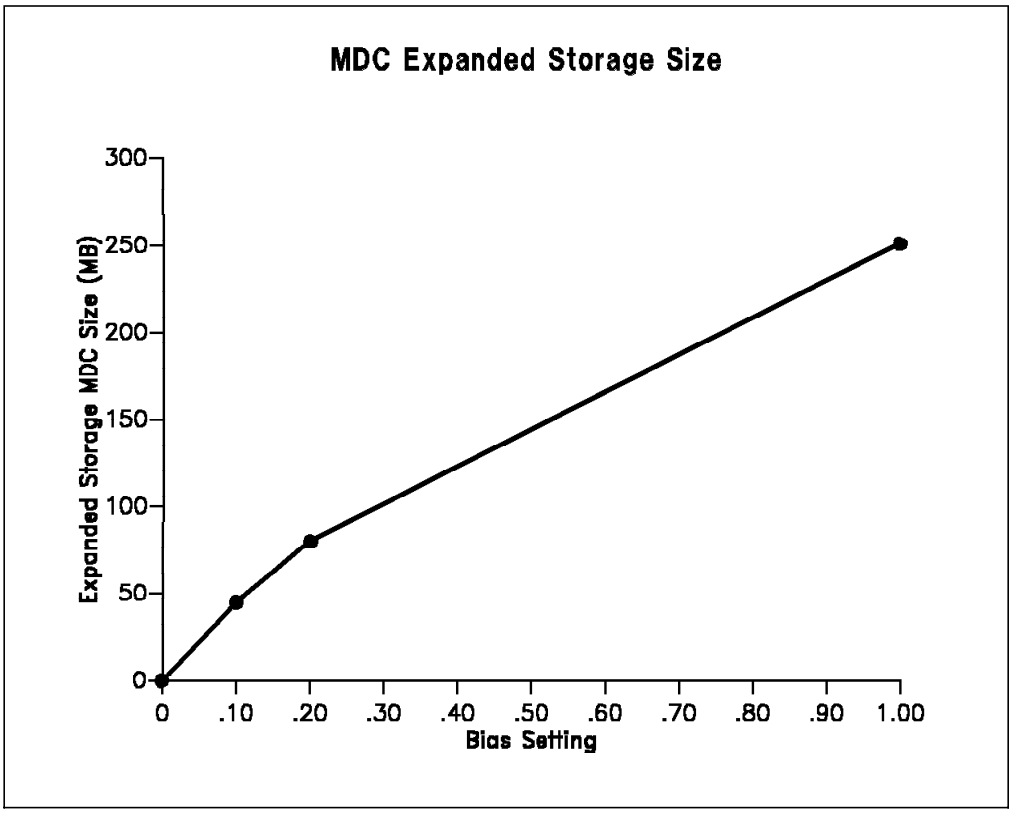

Figure 15. Biasing the Expanded Storage Arbiter. The average size of the MDC expanded storage cache with various bias settings for expanded storage.

Enhanced Minidisk Cache

Table 40 (Page 1 of 2). Use of the BIAS option with the expanded storage arbiter on the 9121-742, not storage constrained.

BIAS Setting (exp) **Release** **Run ID**	**1.0** **VM/ESA 2.2** **S47E5505**	**0.2** **VM/ESA 2.2** **S47E5503**	**0.1** **VM/ESA 2.2** **S47E5504**
Environment			
Real Storage	**1024MB**	**1024MB**	**1024MB**
Exp. Storage	**1024MB**	**1024MB**	**1024MB**
Users	**5500**	**5500**	**5500**
VTAMs	**1**	**1**	**1**
VSCSs	**3**	**3**	**3**
Processors	**4**	**4**	**4**
Response Time			
TRIV INT	0.120	0.109	0.110
NONTRIV INT	0.383	0.344	0.337
TOT INT	0.274	0.252	0.250
TOT INT ADJ	0.271	0.239	0.233
AVG FIRST (T)	0.303	0.247	0.231
AVG LAST (T)	0.391	0.326	0.308
Throughput			
AVG THINK (T)	26.10	26.10	26.07
ETR	190.87	183.16	179.76
ETR (T)	192.83	192.97	193.12
ETR RATIO	0.990	0.949	0.931
ITR (H)	213.11	217.78	218.42
ITR	52.76	51.72	50.91
EMUL ITR	80.70	77.73	76.29
ITRR (H)	1.000	1.022	1.025
ITRR	1.000	0.980	0.965
Proc. Usage			
PBT/CMD (H)	18.770	18.367	18.313
PBT/CMD	18.773	18.345	18.279
CP/CMD (H)	6.915	6.551	6.490
CP/CMD	6.482	6.115	6.058
EMUL/CMD (H)	11.855	11.816	11.823
EMUL/CMD	12.291	12.230	12.220
Processor Util.			
TOTAL (H)	361.93	354.43	353.66
TOTAL	362.00	354.00	353.00
UTIL/PROC (H)	90.48	88.61	88.42
UTIL/PROC	90.50	88.50	88.25
TOTAL EMUL (H)	228.60	228.02	228.33
TOTAL EMUL	237.00	236.00	236.00
MASTER TOTAL (H)	92.76	91.04	90.85
MASTER TOTAL	93.00	91.00	91.00
MASTER EMUL (H)	36.96	38.61	38.95
MASTER EMUL	39.00	40.00	40.00
TVR(H)	1.58	1.55	1.55
TVR	1.53	1.50	1.50
Storage			
NUCLEUS SIZE (V)	2540KB	2540KB	2540KB
TRACE TABLE (V)	800KB	800KB	800KB
WKSET (V)	79	73	72
PGBLPGS	233k	233k	233k
PGBLPGS/USER	42.4	42.4	42.4
FREEPGS	15570	15684	15585
FREE UTIL	0.92	0.93	0.92
SHRPGS	1787	1795	1765

Table 40 (Page 2 of 2). Use of the BIAS option with the expanded storage arbiter on the 9121-742, not storage constrained.

BIAS Setting (exp)	**1.0**	**0.2**	**0.1**
Release	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
Run ID	**S47E5505**	**S47E5503**	**S47E5504**
Environment			
Real Storage	**1024MB**	**1024MB**	**1024MB**
Exp. Storage	**1024MB**	**1024MB**	**1024MB**
Users	**5500**	**5500**	**5500**
VTAMs	**1**	**1**	**1**
VSCSs	**3**	**3**	**3**
Processors	**4**	**4**	**4**
Paging			
READS/SEC	1265	692	623
WRITES/SEC	1005	521	455
PAGE/CMD	11.772	6.286	5.582
PAGE IO RATE (V)	379.700	186.200	162.300
PAGE IO/CMD (V)	1.969	0.965	0.840
XSTOR IN/SEC	349	777	828
XSTOR OUT/SEC	1548	1433	1409
XSTOR/CMD	9.838	11.452	11.583
FAST CLR/CMD	8.842	8.815	8.813
Queues			
DISPATCH LIST	109.00	100.35	102.55
ELIGIBLE LIST	0.00	0.00	0.02
I/O			
VIO RATE	1799	1806	1806
VIO/CMD	9.330	9.359	9.352
RIO RATE (V)	744	562	547
RIO/CMD (V)	3.858	2.912	2.832
NONPAGE RIO/CMD (V)	1.889	1.947	1.992
DASD RESP TIME (V)	21.100	20.200	19.900
MDC REAL SIZE (MB)	31.9	33.1	33.3
MDC XSTOR SIZE (MB)	250.0	78.1	43.0
MDC READS (I/Os)	561	558	557
MDC WRITES (I/Os)	28	27	26
MDC AVOID	528	521	515
MDC HIT RATIO	0.94	0.92	0.91
PRIVOPs			
PRIVOP/CMD	20.449	20.551	20.570
DIAG/CMD	25.295	25.341	25.324
DIAG 04/CMD	0.933	0.932	0.948
DIAG 08/CMD	0.737	0.736	0.738
DIAG 0C/CMD	1.125	1.125	1.125
DIAG 14/CMD	0.025	0.024	0.025
DIAG 58/CMD	1.248	1.248	1.249
DIAG 98/CMD	0.324	0.342	0.345
DIAG A4/CMD	3.601	3.599	3.592
DIAG A8/CMD	2.817	2.827	2.824
DIAG 214/CMD	13.339	13.358	13.333
SIE/CMD	51.859	57.003	56.959
SIE INTCPT/CMD	35.264	38.192	38.162
FREE TOTL/CMD	44.682	44.799	44.821
VTAM Machines			
WKSET (V)	4136	4144	4137
TOT CPU/CMD (V)	2.7716	2.7638	2.7703
CP CPU/CMD (V)	1.2389	1.2322	1.2312
VIRT CPU/CMD (V)	1.5327	1.5316	1.5390
DIAG 98/CMD (V)	0.324	0.343	0.345

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Enhanced Minidisk Cache

Workload: FS8F0R

Hardware Configuration

Processor model: 9121-742
Processors used: 4
Storage
  Real: 320MB (Default MDC)
  Expanded: 512MB
Tape: 3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*User*	*Server*	*System*
3390-2	3990-3	4	8	5	7	32 R		2 R
3390-2	3990-2	4	18	4	6			

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	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	500	QUICKDSP ON
VSCSn	3	VSCS	64MB/XA	10000	1200	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	550	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	5200	Users	3MB/XC	100		

Measurement Discussion: The 9121-742 environment was made storage constrained by configuring only 320MB of real storage and 512MB of expanded storage. The MDC expanded storage bias settings for the three measurements were 1.0, 0.5, and 0.1. Unlike the unconstrained measurements, performance degraded when the lower bias settings were used. This is seen in response time (AVG LAST(T)) and internal throughput (ITR(H)). For the measurement with the 0.1 bias setting, the smaller cache size resulted in a lower minidisk cache hit ratio (MDC HIT RATIO). This data is the basis for the guideline that the expanded storage arbiter does a good job in all but the unconstrained storage environments.

Table 41 (Page 1 of 2). Use of the BIAS option with expanded storage on the 9121-742, storage constrained

BIAS Setting (exp)	**1.0**	**0.5**	**0.1**
Release	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
Run ID	**S47E5202**	**S47E5204**	**S47E5203**
Environment			
Real Storage	**320MB**	**320MB**	**320MB**
Exp. Storage	**512MB**	**512MB**	**512MB**
Users	**5200**	**5200**	**5200**
VTAMs	**1**	**1**	**1**
VSCSs	**3**	**3**	**3**
Processors	**4**	**4**	**4**
Response Time			
TRIV INT	0.156	0.159	0.152
NONTRIV INT	0.413	0.436	0.416
TOT INT	0.319	0.332	0.315
TOT INT ADJ	0.293	0.309	0.294
AVG FIRST (T)	0.283	0.303	0.294
AVG LAST (T)	0.365	0.392	0.379
Throughput			
AVG THINK (T)	26.15	26.15	26.14
ETR	167.41	169.74	170.55
ETR (T)	182.45	182.52	182.42
ETR RATIO	0.918	0.930	0.935
ITR (H)	206.67	206.19	205.32
ITR	47.45	47.98	48.05
EMUL ITR	74.02	74.98	75.27
ITRR (H)	1.000	0.998	0.993
ITRR	1.000	1.011	1.013
Proc. Usage			
PBT/CMD (H)	19.354	19.399	19.481
PBT/CMD	19.347	19.396	19.460
CP/CMD (H)	7.371	7.412	7.472
CP/CMD	6.961	6.958	7.017
EMUL/CMD (H)	11.983	11.987	12.010
EMUL/CMD	12.387	12.437	12.444
Processor Util.			
TOTAL (H)	353.12	354.07	355.39
TOTAL	353.00	354.00	355.00
UTIL/PROC (H)	88.28	88.52	88.85
UTIL/PROC	88.25	88.50	88.75
TOTAL EMUL (H)	218.63	218.79	219.09
TOTAL EMUL	226.00	227.00	227.00
MASTER TOTAL (H)	90.77	90.96	91.22
MASTER TOTAL	91.00	91.00	91.00
MASTER EMUL (H)	35.10	35.41	35.14
MASTER EMUL	37.00	37.00	37.00
TVR(H)	1.62	1.62	1.62
TVR	1.56	1.56	1.56
Storage			
NUCLEUS SIZE (V)	2556KB	2556KB	2556KB
TRACE TABLE (V)	800KB	800KB	800KB
WKSET (V)	69	69	69
PGBLPGS	55045	55106	54812
PGBLPGS/USER	10.6	10.6	10.5
FREEPGS	14900	14855	15137
FREE UTIL	0.91	0.91	0.91
SHRPGS	1695	1690	1679

Enhanced Minidisk Cache

Table 41 (Page 2 of 2). Use of the BIAS option with expanded storage on the 9121-742, storage constrained

BIAS Setting (exp) **Release** **Run ID**	**1.0** **VM/ESA 2.2** **S47E5202**	**0.5** **VM/ESA 2.2** **S47E5204**	**0.1** **VM/ESA 2.2** **S47E5203**
Environment			
Real Storage	**320MB**	**320MB**	**320MB**
Exp. Storage	**512MB**	**512MB**	**512MB**
Users	**5200**	**5200**	**5200**
VTAMs	**1**	**1**	**1**
VSCSs	**3**	**3**	**3**
Processors	**4**	**4**	**4**
Paging			
READS/SEC	1637	1681	1606
WRITES/SEC	1296	1333	1271
PAGE/CMD	16.075	16.514	15.771
PAGE IO RATE (V)	531.700	543.500	523.200
PAGE IO/CMD (V)	2.914	2.978	2.868
XSTOR IN/SEC	378	373	403
XSTOR OUT/SEC	1531	1458	1669
XSTOR/CMD	10.463	10.032	11.358
FAST CLR/CMD	8.791	8.744	8.782
Queues			
DISPATCH LIST	111.38	117.85	115.25
ELIGIBLE LIST	0.22	0.75	0.07
I/O			
VIO RATE	1708	1707	1705
VIO/CMD	9.361	9.353	9.346
RIO RATE (V)	894	912	924
RIO/CMD (V)	4.900	4.997	5.065
NONPAGE RIO/CMD (V)	1.986	2.019	2.197
DASD RESP TIME (V)	21.600	22.200	20.800
MDC REAL SIZE (MB)	7.9	8.4	7.2
MDC XSTOR SIZE (MB)	117.2	58.6	13.7
MDC READS (I/Os)	522	521	512
MDC WRITES (I/Os)	25	25	24
MDC AVOID	488	484	454
MDC HIT RATIO	0.92	0.92	0.86
PRIVOPs			
PRIVOP/CMD	20.954	20.916	20.894
DIAG/CMD	25.306	25.294	25.304
DIAG 04/CMD	0.982	0.981	0.982
DIAG 08/CMD	0.737	0.737	0.736
DIAG 0C/CMD	1.126	1.125	1.125
DIAG 14/CMD	0.025	0.025	0.024
DIAG 58/CMD	1.248	1.248	1.249
DIAG 98/CMD	0.376	0.384	0.380
DIAG A4/CMD	3.563	3.561	3.563
DIAG A8/CMD	2.827	2.815	2.818
DIAG 214/CMD	13.276	13.271	13.278
SIE/CMD	54.808	54.790	54.818
SIE INTCPT/CMD	36.722	36.709	36.180
FREE TOTL/CMD	45.036	45.087	45.438
VTAM Machines			
WKSET (V)	4010	4009	4006
TOT CPU/CMD (V)	2.8941	2.8947	2.8901
CP CPU/CMD (V)	1.2903	1.2906	1.2852
VIRT CPU/CMD (V)	1.6038	1.6041	1.6049
DIAG 98/CMD (V)	0.376	0.384	0.380

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Biasing the Real Storage Arbiter

The guidelines for biasing the real storage arbiter are similar to those for the expanded storage arbiter. In all but the most unconstrained environments, the arbiter appears to make good decisions with the default bias setting. However, in very unconstrained environments, performance can be improved by biasing the real storage arbiter against MDC.

Workload: FS8F0R

Hardware Configuration

Processor model:	9121-742
Processors used:	4
Storage	
Real:	1024MB
Expanded:	1024MB
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-3	4	6	7	7	32 R		2 R
3390-2	3990-2	4	16	6	6			

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	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	500	QUICKDSP ON
VSCSn	3	VSCS	64MB/XA	10000	1200	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	550	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	5500	Users	3MB/XC	100		

Measurement Discussion: These measurements were done in an unconstrained environment. For these measurements, only real storage was used for MDC. The first measurement used the default real storage bias setting (1.0), while the second measurement used a setting of 0.2. With the lower bias setting, the response time (AVG LAST(T)) improved 8% and the internal throughput rate

(ITR(H)) improved 1.6%. While biasing the real storage arbiter against minidisk cache caused an increase in non-page I/O (NONPAGE RIO/CMD(V)), this was more than offset by a decrease in paging I/O.

Table 42 (Page 1 of 3). Use of the BIAS option with the real storage arbiter on the 9121-742

BIAS Setting (real) **Release** **Run ID**	**1.0** **VM/ESA 2.2** **S47E5508**	**0.2** **VM/ESA 2.2** **S47E5509**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Response Time				
TRIV INT	0.114	0.107	-0.007	-6.14%
NONTRIV INT	0.358	0.337	-0.021	-5.87%
TOT INT	0.260	0.247	-0.013	-5.00%
TOT INT ADJ	0.250	0.234	-0.016	-6.45%
AVG FIRST (T)	0.266	0.241	-0.024	-9.17%
AVG LAST (T)	0.349	0.320	-0.029	-8.32%
Throughput				
AVG THINK (T)	26.07	26.11	0.03	0.13%
ETR	185.75	182.92	-2.83	-1.52%
ETR (T)	193.01	193.02	0.01	0.01%
ETR RATIO	0.962	0.948	-0.015	-1.53%
ITR (H)	215.61	218.97	3.35	1.56%
ITR	51.90	51.93	0.04	0.07%
EMUL ITR	78.67	77.71	-0.96	-1.22%
ITRR (H)	1.000	1.016	0.016	1.56%
ITRR	1.000	1.001	0.001	0.07%
Proc. Usage				
PBT/CMD (H)	18.552	18.268	-0.284	-1.53%
PBT/CMD	18.548	18.236	-0.312	-1.68%
CP/CMD (H)	6.716	6.459	-0.257	-3.83%
CP/CMD	6.321	6.061	-0.259	-4.10%
EMUL/CMD (H)	11.836	11.809	-0.027	-0.23%
EMUL/CMD	12.227	12.175	-0.052	-0.43%
Processor Util.				
TOTAL (H)	358.08	352.61	-5.47	-1.53%
TOTAL	358.00	352.00	-6.00	-1.68%
UTIL/PROC (H)	89.52	88.15	-1.37	-1.53%
UTIL/PROC	89.50	88.00	-1.50	-1.68%
TOTAL EMUL (H)	228.45	227.94	-0.51	-0.22%
TOTAL EMUL	236.00	235.00	-1.00	-0.42%
MASTER TOTAL (H)	91.80	90.54	-1.26	-1.37%
MASTER TOTAL	92.00	90.00	-2.00	-2.17%
MASTER EMUL (H)	37.87	38.99	1.12	2.97%
MASTER EMUL	39.00	41.00	2.00	5.13%
TVR(H)	1.57	1.55	-0.02	-1.31%
TVR	1.52	1.50	-0.02	-1.26%

Table 42 (Page 2 of 3). Use of the BIAS option with the real storage arbiter on the 9121-742

BIAS Setting (real) **Release** **Run ID**	**1.0** **VM/ESA 2.2** **S47E5508**	**0.2** **VM/ESA 2.2** **S47E5509**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Storage				
NUCLEUS SIZE (V)	2540KB	2556KB	16KB	0.63%
TRACE TABLE (V)	800KB	800KB	0KB	0.00%
WKSET (V)	76	72	-4	-5.26%
PGBLPGS	233k	233k	0k	0.00%
PGBLPGS/USER	42.4	42.4	0.0	0.00%
FREEPGS	15568	15595	27	0.17%
FREE UTIL	0.92	0.92	0.00	-0.17%
SHRPGS	1801	1800	-1	-0.06%
Paging				
READS/SEC	940	571	-369	-39.26%
WRITES/SEC	732	416	-316	-43.17%
PAGE/CMD	8.663	5.113	-3.549	-40.97%
PAGE IO RATE (V)	270.100	147.000	-123.100	-45.58%
PAGE IO/CMD (V)	1.399	0.762	-0.638	-45.58%
XSTOR IN/SEC	594	870	276	46.46%
XSTOR OUT/SEC	1486	1404	-82	-5.52%
XSTOR/CMD	10.777	11.781	1.005	9.32%
FAST CLR/CMD	8.828	8.755	-0.073	-0.83%
Queues				
DISPATCH LIST	104.89	99.76	-5.13	-4.89%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	1802	1800	-2	-0.11%
VIO/CMD	9.336	9.325	-0.011	-0.12%
RIO RATE (V)	634	518	-116	-18.30%
RIO/CMD (V)	3.285	2.684	-0.601	-18.30%
NONPAGE RIO/CMD (V)	1.885	1.922	0.037	1.95%
DASD RESP TIME (V)	20.500	20.300	-0.200	-0.98%
MDC REAL SIZE (MB)	191.4	39.1	-152.3	-79.57%
MDC READS (I/Os)	559	552	-7	-1.25%
MDC WRITES (I/Os)	27	26	-1	-3.70%
MDC AVOID	527	519	-8	-1.52%
MDC HIT RATIO	0.93	0.93	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	20.490	20.544	0.054	0.26%
DIAG/CMD	25.314	25.224	-0.090	-0.36%
DIAG 04/CMD	0.933	0.932	-0.002	-0.17%
DIAG 08/CMD	0.736	0.737	0.001	0.07%
DIAG 0C/CMD	1.125	1.126	0.001	0.08%
DIAG 14/CMD	0.024	0.025	0.000	0.20%
DIAG 58/CMD	1.249	1.248	0.000	-0.03%
DIAG 98/CMD	0.328	0.353	0.024	7.43%
DIAG A4/CMD	3.604	3.569	-0.035	-0.97%
DIAG A8/CMD	2.818	2.820	0.001	0.05%
DIAG 214/CMD	13.349	13.268	-0.081	-0.60%
SIE/CMD	56.991	56.988	-0.003	-0.01%
SIE INTCPT/CMD	38.184	38.182	-0.002	-0.01%
FREE TOTL/CMD	44.474	44.461	-0.013	-0.03%

Table 42 (Page 3 of 3). Use of the BIAS option with the real storage arbiter on the 9121-742

BIAS Setting (real) **Release** **Run ID**	**1.0** **VM/ESA 2.2** **S47E5508**	**0.2** **VM/ESA 2.2** **S47E5509**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
VTAM Machines				
WKSET (V)	4134	4144	10	0.24%
TOT CPU/CMD (V)	2.7834	2.7803	-0.0031	-0.11%
CP CPU/CMD (V)	1.2377	1.2405	0.0028	0.23%
VIRT CPU/CMD (V)	1.5457	1.5398	-0.0059	-0.38%
DIAG 98/CMD (V)	0.328	0.352	0.024	7.36%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Use of Cached Control Units with Minidisk Caching

With the improvements to minidisk cache, the question arises as to whether cached control units are still beneficial. The answer is yes. MDC only benefits read I/O, so cached control units with DASD fast write can provide improvements for write I/O. MDC does not apply to CP data (such as spool volumes) or to shared DASD. As a result, cached control units can provide significant additional benefit in these situations. Lastly, cached control units shorten DASD access time for cache misses in MDC. The measurements in this section illustrate the benefit of using both MDC and cached control units.

Workload: FS8F0R

Hardware Configuration

Processor model:	9121-480
Processors used:	2
Storage	
Real:	256MB (default MDC)
Expanded:	0MB
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes -* *PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	2						2 R,W
3390-2	3990-3	4		2	2	16 R,W		

Note: *R* or *W* next to the DASD counts means basic cache enabled or DASD fast write (and basic cache) enabled, respectively. Control unit caching was only enabled for some of the measurements (see the results table).

Communications:

Control Unit	*Number*	*Lines per Control Unit*	*Speed*
3088-08	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XC	10000	512	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1860	Users	3MB/XC	100		

Measurement Discussion: Running with the control units enabled for basic (read) caching showed almost a 5% improvement in response time (AVG LAST(T)). Internal throughput (ITR(H)) was equivalent. While overall DASD response time (DASD RESP TIME(V)) was equivalent, the DASD response time on key volumes improved significantly. For example, the system resident volume response time improved 13%.

Enabling both basic and DASD fast write cache resulted in over a 6% improvement in response time (AVG LAST(T)). Internal throughput (ITR(H)) was again equivalent. In this comparison, the average DASD response time (DASD RESP TIME(V)) was reduced 11%.

In this environment, all user DASD I/O is eligible for minidisk caching and there is sufficient real storage for the cache. This is a best case scenario for using MDC alone. In cases where either of these is not true, the additional benefits of cached control units can be much greater. Also, scenarios that include server machines doing synchronous I/O can greatly benefit from cached control units with DASD fast write. This would include OfficeVision* environments.

Table 43 (Page 1 of 2). Use of MDC and control unit caching on the 9121-480

Cache Control Unit **Release** **Run ID**	**Off** **VM/ESA 2.2** **L27E186B**	**Read** **VM/ESA 2.2** **L27E186A**	**Read + Write** **VM/ESA 2.2** **L27E186M**
Environment			
Real Storage	**256MB**	**256MB**	**256MB**
Exp. Storage	**0MB**	**0MB**	**0MB**
Users	**1860**	**1860**	**1860**
VTAMs	**1**	**1**	**1**
VSCSs	**0**	**0**	**0**
Processors	**2**	**2**	**2**
Response Time			
TRIV INT	0.128	0.128	0.123
NONTRIV INT	0.401	0.397	0.391
TOT INT	0.308	0.307	0.301
TOT INT ADJ	0.273	0.269	0.264
AVG FIRST (T)	0.250	0.234	0.236
AVG LAST (T)	0.341	0.325	0.320
Throughput			
AVG THINK (T)	26.17	26.18	26.21
ETR	57.81	57.30	57.38
ETR (T)	65.34	65.33	65.40
ETR RATIO	0.885	0.877	0.877
ITR (H)	75.38	75.36	75.35
ITR	33.36	33.06	33.06
EMUL ITR	48.41	47.99	47.99
ITRR (H)	1.000	1.000	1.000
ITRR	1.000	0.991	0.991
Proc. Usage			
PBT/CMD (H)	26.533	26.538	26.544
PBT/CMD	26.478	26.479	26.604
CP/CMD (H)	8.822	8.824	8.833
CP/CMD	8.265	8.265	8.256
EMUL/CMD (H)	17.712	17.714	17.711
EMUL/CMD	18.213	18.214	18.348
Processor Util.			
TOTAL (H)	173.36	173.38	173.61
TOTAL	173.00	173.00	174.00
UTIL/PROC (H)	86.68	86.69	86.80
UTIL/PROC	86.50	86.50	87.00
TOTAL EMUL (H)	115.72	115.73	115.84
TOTAL EMUL	119.00	119.00	120.00
MASTER TOTAL (H)	86.28	86.25	86.35
MASTER TOTAL	86.00	86.00	86.00
MASTER EMUL (H)	51.05	50.95	50.99
MASTER EMUL	53.00	53.00	53.00
TVR(H)	1.50	1.50	1.50
TVR	1.45	1.45	1.45
Storage			
NUCLEUS SIZE (V)	2536KB	2536KB	2536KB
TRACE TABLE (V)	400KB	400KB	400KB
WKSET (V)	84	84	84
PGBLPGS	55476	55497	55496
PGBLPGS/USER	29.8	29.8	29.8
FREEPGS	5252	5247	5252
FREE UTIL	0.93	0.93	0.93
SHRPGS	1268	1281	1222

Table 43 (Page 2 of 2). Use of MDC and control unit caching on the 9121-480

Cache Control Unit	**Off**	**Read**	**Read + Write**
Release	**VM/ESA 2.2**	**VM/ESA 2.2**	**VM/ESA 2.2**
Run ID	**L27E186B**	**L27E186A**	**L27E186M**
Environment			
Real Storage	**256MB**	**256MB**	**256MB**
Exp. Storage	**0MB**	**0MB**	**0MB**
Users	**1860**	**1860**	**1860**
VTAMs	**1**	**1**	**1**
VSCSs	**0**	**0**	**0**
Processors	**2**	**2**	**2**
Paging			
READS/SEC	627	628	630
WRITES/SEC	434	433	433
PAGE/CMD	16.239	16.240	16.253
PAGE IO RATE (V)	176.300	174.100	175.200
PAGE IO/CMD (V)	2.698	2.665	2.679
XSTOR IN/SEC	0	0	0
XSTOR OUT/SEC	0	0	0
XSTOR/CMD	0.000	0.000	0.000
FAST CLR/CMD	8.770	8.770	8.761
Queues			
DISPATCH LIST	38.11	39.07	34.60
ELIGIBLE LIST	0.00	0.00	0.00
I/O			
VIO RATE	663	664	666
VIO/CMD	10.147	10.163	10.183
RIO RATE (V)	363	362	365
RIO/CMD (V)	5.556	5.541	5.581
NONPAGE RIO/CMD (V)	2.858	2.876	2.902
DASD RESP TIME (V)	19.500	19.700	17.300
MDC READS (I/Os)	189	190	189
MDC WRITES (I/Os)	9.30	9.45	9.52
MDC AVOID	178	179	179
MDC HIT RATIO	0.94	0.94	0.94
PRIVOPs			
PRIVOP/CMD	13.905	13.904	13.889
DIAG/CMD	27.684	27.709	27.711
DIAG 04/CMD	2.471	2.471	2.469
DIAG 08/CMD	0.737	0.738	0.738
DIAG 0C/CMD	1.126	1.126	1.126
DIAG 14/CMD	0.025	0.025	0.025
DIAG 58/CMD	1.248	1.249	1.249
DIAG 98/CMD	1.146	1.168	1.181
DIAG A4/CMD	3.610	3.616	3.615
DIAG A8/CMD	2.823	2.817	2.825
DIAG 214/CMD	13.349	13.350	13.336
SIE/CMD	53.491	53.541	53.651
SIE INTCPT/CMD	34.769	34.801	34.873
FREE TOTL/CMD	49.665	49.653	49.691
VTAM Machines			
WKSET (V)	505	542	547
TOT CPU/CMD (V)	3.8346	3.8520	3.8394
CP CPU/CMD (V)	1.4539	1.4541	1.4525
VIRT CPU/CMD (V)	2.3807	2.3979	2.3869
DIAG 98/CMD (V)	1.146	1.167	1.181

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Scheduler Share Capping and Proportional Distribution

The scheduler enhancements for this release consisted of two major changes and a few minor ones. The first major change is for surplus (or excess) processor resources to be distributed proportionally to share settings. The second major change is to add an upper bound on share, since the current share setting is a floor, not a ceiling. Customer requests were the driving force behind these changes. These requests were in the form of APARs that had been closed as suggestions (SUG) and a number of user group requirements. These changes should lead to more efficient and consistent use of system resources, easier tuning, and fewer scheduler problems. The minor changes are discussed later in this section.

Before discussing these changes further, a review of some terminology is helpful.

Regular Share: The existing setting for SET SHARE command and SHARE directory option statement. This is a target minimum share of system resources for given virtual machine.

Normalized Share: The value arrived at when compared to the share setting of other users in the dispatch list. If the sum of absolute shares is greater than 99%, then absolute shares are normalized with each other to 99%. Relative shares are normalized with each other to 100% minus the absolutes shares (99% at most). Because of the dynamics of the number and type of users in the dispatch list, the normalized share value is dynamic. The scheduler uses this value in prioritizing users.

Surplus Share: When a virtual machine's actual resource usage is less than its normalized share, the delta is considered surplus. The virtual machine can not use the minimum resources we want to assign it. This is also known as excess share.

Maximum Share: This is the new share setting introduced with VM/ESA Release 2.2 and is discussed in more detail later. It is some times referred to as the limit share.

The share capping enhancement was implemented through a maximum share setting on the SET SHARE command or SHARE directory option. The SET SHARE command and Directory option now have parameters that allow you to set a maximum share (limit share). As with regular share (target minimum), this can be done as either a relative or absolute setting. The maximum share setting also has a limit type associated with it. Three types of limits are supported: NOLIMIT, LIMITSOFT, and LIMITHARD.

- NOLIMIT (the default) means that the user will not be limited.
- LIMITSOFT means that the user will not get more than its share if there is any other user that can use the surplus without exceeding its own limit. However, if no other user can use the processing time, the limit will be overridden to give the user more than its share rather than letting the processor time go to waste.

- LIMITHARD means that the user will not get more than its share, even if no one else can use the surplus. Extensive use of LIMITHARD could result in processor resources going unused when they could have been used.

Some example settings are as follows:

```
SET SHARE LARRY   REL 2000 REL 4000  LIMITHARD
SET SHARE MOE     REL 100  ABS 5%    LIMITSOFT
SET SHARE ALBERT  REL 1000 NOLIMIT
```

Users that have exceeded their maximum share are kept in a subset of the dispatch list known as the limit list. There is a priority associated with users in the limit list similar to the priority used in other scheduler lists. The number of users on the limit list can be determined from the INDICATE LOAD command. The INDICATE QUEUE command can be used to identify the users on the limit list and their associated priority. Users are removed from the limit list (and therefore become dispatchable again) when enough time has elapsed such that they fall below their maximum share. Most limited users are expected to have very short stays on the limit list. Therefore, the snapshot from the INDICATE commands may not show these users.

The other significant change to the scheduler is the distribution of surplus share. In prior releases, surplus processing favored the ready-to-run user having the highest share. In VM/ESA Release 2.2, the surplus is given out to other users in a manner that is proportional to their shares. Surplus share is often generated by setting very high relative shares (10000) for server virtual machines. The normalized share for those virtual machines can often be 30% or higher, which is more processor than many server virtual machines can use.

In previous releases, the distribution of surplus to the next largest share often made a slightly favored virtual machine look like it was running away with the system. If you were looking for share capping to solve a problem like that, the proportional distribution of surplus may be enough to correct this behavior.

One of the minor enhancements is the ability to use decimal values on the absolute share settings. While the SET SHARE command accepts several digits to the right of the decimal point, it is recommended that only tenths be used. The accuracy of values with smaller granularity is difficult to guarantee at this time. With this change the smallest absolute setting is ABS 0.1%.

Another minor change was for the SET SRM STORBUF default values. They have been changed from 100 85 75 to 125 105 95.

The last minor change worth noting is the change to ratios for elapsed timeslice (ETS). The ETS determines the length of a user's stay in the dispatch list. Prior to VM/ESA Release 2.2, the ratio was 1 to 6 to 48 which meant a class 2 user (Q2) was 6 times the Q1 ETS and a Q3 was 48 times the Q1 ETS. Now the Q2 ETS is 8 times the Q1 ETS, while the Q3 ETS remains unchanged. This change was made to improve the distribution of users in the three scheduler classes. Observation of various systems showed that there were typically fewer users in Q2 than Q1 or Q3.

Measurements were made to illustrate the two major enhancements: share capping and proportional distribution of surplus. The results of these

measurements are shown in the sections that follow. Measurements were not made to explicitly show the impact of the minor changes.

Share Capping

Workload: FS8F0R

Hardware Configuration

Processor model: 9121-480
Processors used: 2
Storage:
Real: 256MB (default MDC)
Expanded: 0MB
Tape: 3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2	16 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*
3088-08	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	512	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1860	Users	3MB/XC	100		
MARK	1	Runaway	24MB/XA	100		

Measurement Discussion: To illustrate the effect of share capping, a runaway user (user in an infinite loop) was added to the standard CMS workload. This was done on VM/ESA Release 2.1, VM/ESA Release 2.2, and VM/ESA Release 2.2 with a maximum share setting of ABSOLUTE 1% LIMITHARD.

The maximum share setting worked as expected. In the last row of the table, you see that the processor usage by the runaway user was limited to under 1%

of the system while in the unlimited measurements it used over 10% of the system. This improved system performance significantly as is seen in a decrease for response time (AVG LAST(T)) and an increase in internal throughput rate (ITR(H)).

Table 44 (Page 1 of 3). Use of maximum share setting on the 9121-480

Maximum Share **Release** **Run ID**	**No Limit** **VM/ESA 2.2** **L27E186G**	**Abs 1% Hard** **VM/ESA 2.2** **L27E186E**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.143	0.129	-0.014	-9.79%
NONTRIV INT	0.523	0.418	-0.105	-20.08%
TOT INT	0.395	0.320	-0.075	-18.99%
TOT INT ADJ	0.350	0.282	-0.068	-19.52%
AVG FIRST (T)	0.289	0.244	-0.046	-15.89%
AVG LAST (T)	0.427	0.343	-0.084	-19.77%
Throughput				
AVG THINK (T)	26.28	26.25	-0.03	-0.11%
ETR	57.97	57.73	-0.24	-0.41%
ETR (T)	65.36	65.52	0.16	0.24%
ETR RATIO	0.887	0.881	-0.006	-0.66%
ITR (H)	66.13	74.21	8.09	12.23%
ITR	29.33	32.70	3.38	11.51%
EMUL ITR	40.32	47.30	6.98	17.32%
ITRR (H)	1.000	1.122	0.122	12.23%
ITRR	1.000	1.115	0.115	11.51%
Proc. Usage				
PBT/CMD (H)	30.246	26.950	-3.296	-10.90%
PBT/CMD	30.294	27.015	-3.279	-10.82%
CP/CMD (H)	8.826	8.895	0.069	0.79%
CP/CMD	8.262	8.394	0.132	1.60%
EMUL/CMD (H)	21.420	18.055	-3.365	-15.71%
EMUL/CMD	22.032	18.620	-3.411	-15.48%
Processor Util.				
TOTAL (H)	197.69	176.57	-21.11	-10.68%
TOTAL	198.00	177.00	-21.00	-10.61%
UTIL/PROC (H)	98.84	88.29	-10.56	-10.68%
UTIL/PROC	99.00	88.50	-10.50	-10.61%
TOTAL EMUL (H)	140.00	118.29	-21.71	-15.51%
TOTAL EMUL	144.00	122.00	-22.00	-15.28%
MASTER TOTAL (H)	98.09	87.96	-10.12	-10.32%
MASTER TOTAL	98.00	88.00	-10.00	-10.20%
MASTER EMUL (H)	62.30	52.28	-10.03	-16.09%
MASTER EMUL	64.00	54.00	-10.00	-15.63%
TVR(H)	1.41	1.49	0.08	5.71%
TVR	1.38	1.45	0.08	5.51%

Table 44 (Page 2 of 3). Use of maximum share setting on the 9121-480

Maximum Share **Release** **Run ID**	**No Limit** **VM/ESA 2.2** **L27E186G**	**Abs 1% Hard** **VM/ESA 2.2** **L27E186E**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Storage				
NUCLEUS SIZE (V)	2540KB	2540KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	83	83	0	0.00%
PGBLPGS	55389	55413	24	0.04%
PGBLPGS/USER	29.8	29.8	0.0	0.04%
FREEPGS	5349	5342	-7	-0.13%
FREE UTIL	0.91	0.91	0.00	0.13%
SHRPGS	1277	1215	-62	-4.86%
Paging				
READS/SEC	598	628	30	5.02%
WRITES/SEC	434	434	0	0.00%
PAGE/CMD	15.789	16.209	0.419	2.66%
PAGE IO RATE (V)	172.300	173.400	1.100	0.64%
PAGE IO/CMD (V)	2.636	2.647	0.010	0.39%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	8.859	8.806	-0.052	-0.59%
Queues				
DISPATCH LIST	45.98	41.35	-4.64	-10.08%
ELIGIBLE LIST	0.04	0.04	0.00	-3.64%
I/O				
VIO RATE	656	667	11	1.68%
VIO/CMD	10.037	10.180	0.143	1.43%
RIO RATE (V)	346	360	14	4.05%
RIO/CMD (V)	5.294	5.495	0.201	3.79%
NONPAGE RIO/CMD (V)	2.658	2.848	0.190	7.16%
DASD RESP TIME (V)	19.600	19.500	-0.100	-0.51%
MDC READS (I/Os)	190	191	1	0.53%
MDC WRITES (I/Os)	9.30	9.36	0.06	0.65%
MDC AVOID	179	180	1	0.56%
MDC HIT RATIO	0.94	0.94	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.986	13.960	-0.026	-0.18%
DIAG/CMD	27.476	27.758	0.281	1.02%
DIAG 04/CMD	2.382	2.466	0.084	3.52%
DIAG 08/CMD	0.736	0.738	0.002	0.22%
DIAG 0C/CMD	1.126	1.127	0.001	0.09%
DIAG 14/CMD	0.024	0.025	0.000	0.66%
DIAG 58/CMD	1.249	1.249	0.000	0.00%
DIAG 98/CMD	1.022	1.139	0.117	11.44%
DIAG A4/CMD	3.618	3.631	0.013	0.36%
DIAG A8/CMD	2.820	2.831	0.011	0.38%
DIAG 214/CMD	13.345	13.397	0.052	0.39%
SIE/CMD	54.330	53.678	-0.651	-1.20%
SIE INTCPT/CMD	34.228	34.891	0.663	1.94%
FREE TOTL/CMD	49.219	49.573	0.353	0.72%

Table 44 (Page 3 of 3). Use of maximum share setting on the 9121-480

Maximum Share Release Run ID	No Limit VM/ESA 2.2 L27E186G	Abs 1% Hard VM/ESA 2.2 L27E186E	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
VTAM Machines				
WKSET (V)	506	534	28	5.53%
TOT CPU/CMD (V)	3.7931	3.8580	0.0649	1.71%
CP CPU/CMD (V)	1.4423	1.4669	0.0246	1.71%
VIRT CPU/CMD (V)	2.3508	2.3911	0.0403	1.71%
DIAG 98/CMD (V)	1.016	1.140	0.123	12.13%
Runaway Machine				
CPU PCT BUSY (V)	11.7	0.9	10.8	-92.31%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Proportional Distribution of Surplus CPU

Workload: FS8F0R

Hardware Configuration

Processor model: 9121-480
Processors used: 2
Storage, VM/ESA 2.1:
  Real: 224MB
  Expanded: 32MB (all for MDC)
Storage, VM/ESA 2.2:
  Real: 256MB (default MDC)
  Expanded: 0MB
Tape: 3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2	16 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*
3088-08	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	1024	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1860	Users	3MB/XC	100		
U1950	1	Runaway	3MB/XC	200		
U1951	1	Runaway	3MB/XC	300		
U1952	1	Runaway	3MB/XC	400		
U1953	1	Runaway	3MB/XC	500		
U1954	1	Runaway	3MB/XC	600		

Measurement Discussion: To illustrate the proportional distribution of surplus share, five users in an infinite loop were added to the standard CMS workload. These users were given relative share settings of 200, 300, 400, 500, and 600. Using a scenario such as this illustrates the benefit of distributing surplus share proportionally, but does make it difficult to compare these measurements to others in this report.

Figure 16 on page 146 shows the processor resources used by the five runaway users. The line for VM/ESA Release 2.1 has a sharp curve in it which shows how the users with higher relative share settings benefitted from getting more than their fair share of the surplus. The straighter VM/ESA Release 2.2 line shows how the surplus was more proportionally distributed. If the distribution was perfectly proportional, the line would be straight.

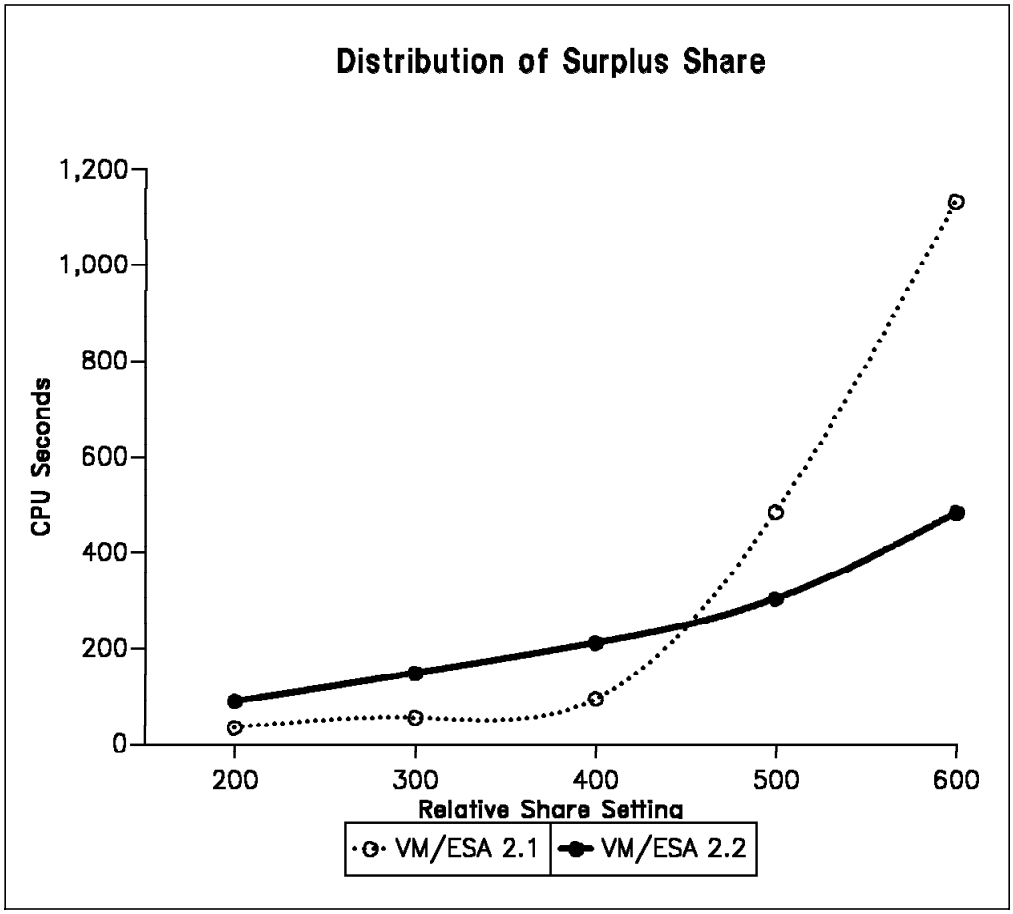

Figure 16. Distribution of surplus share. Processor resources consumed by virtual machines with various relative share settings.

System performance also improved significantly as seen in the 68% improvement in response time (AVG LAST(T)) and the 63% improvement in internal throughput rate (ITR(H)). These improvements were due to more resources being given to end users and VTAM as seen in Figure 17 on page 147 and Figure 18 on page 147.

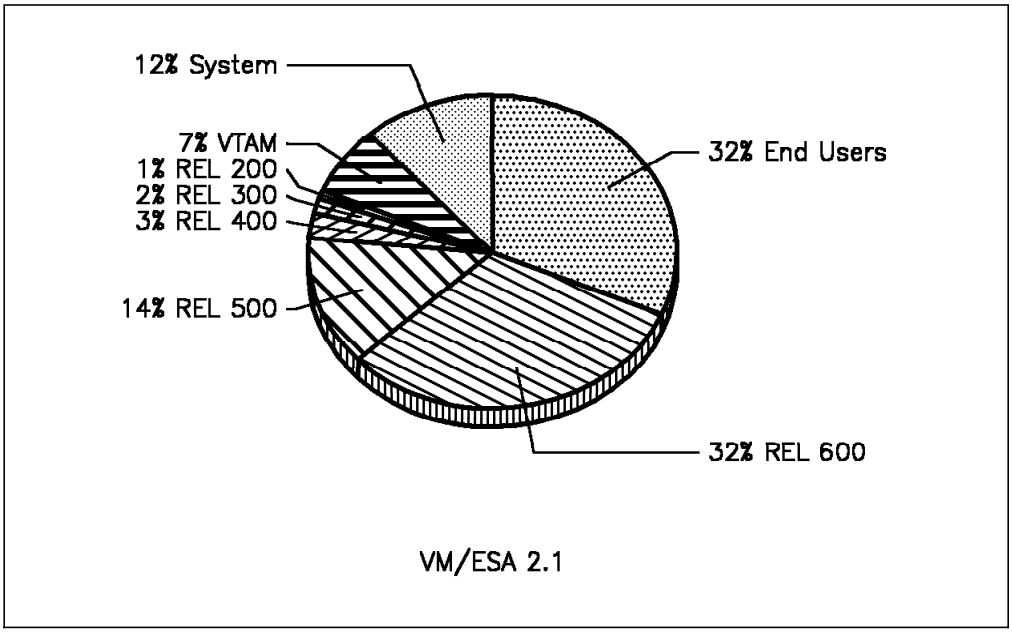

Figure 17. Processor Resource Distribution in VM/ESA 2.1. Percentage of system processor resources given to various virtual machines.

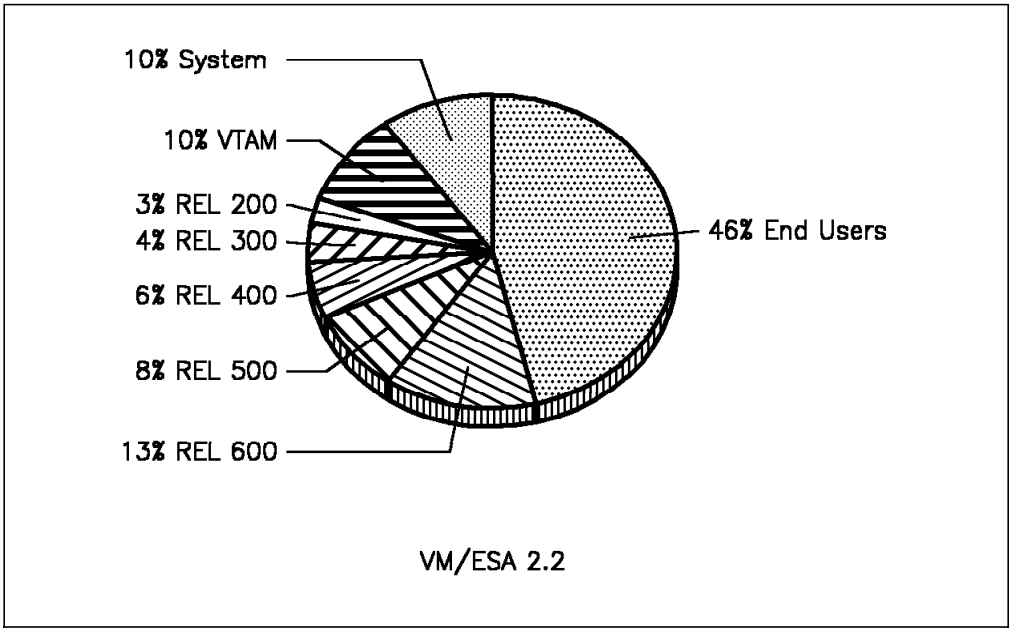

Figure 18. Processor Resource Distribution in VM/ESA 2.2. Percentage of system processor resources given to various virtual machines.

While the response times did improve significantly, a response time of over 12 seconds is not ideal. Some tuning could be applied to improve this. In particular, an eligible list forms in the VM/ESA Release 2.2 measurement due to loading users. Increasing the SRM LDUBUF parameters may help response time. Even with additional tuning, the test case is extreme and not typical of a normal system.

Scheduler Improvements

Table 45 (Page 1 of 2). Proportional distribution of surplus share on the 9121-480

Release Run ID	VM/ESA 2.1 L26E186P	VM/ESA 2.2 L27E186F	Difference	%Difference
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.299	0.030	-0.269	-89.97%
NONTRIV INT	1.925	1.265	-0.660	-34.29%
TOT INT	1.190	0.446	-0.744	-62.52%
TOT INT ADJ	3.870	1.112	-2.758	-71.28%
AVG FIRST (T)	31.854	10.121	-21.733	-68.23%
AVG LAST (T)	39.400	12.567	-26.833	-68.10%
Throughput				
AVG THINK (T)	26.33	26.33	0.00	0.00%
ETR	92.59	116.00	23.41	25.28%
ETR (T)	28.47	46.54	18.07	63.47%
ETR RATIO	3.252	2.492	-0.760	-23.36%
ITR (H)	28.47	46.54	18.07	63.47%
ITR	46.32	58.03	11.71	25.27%
EMUL ITR	60.21	76.17	15.96	26.52%
ITRR (H)	1.000	1.635	0.635	63.47%
ITRR	1.000	1.253	0.253	25.27%
Proc. Usage				
PBT/CMD (H)	70.248	42.974	-27.274	-38.83%
PBT/CMD	70.248	42.974	-27.274	-38.83%
CP/CMD (H)	17.081	10.836	-6.245	-36.56%
CP/CMD	16.157	10.314	-5.843	-36.17%
EMUL/CMD (H)	53.167	32.138	-21.029	-39.55%
EMUL/CMD	54.091	32.660	-21.431	-39.62%
Processor Util.				
TOTAL (H)	200.00	200.00	0.00	0.00%
TOTAL	200.00	200.00	0.00	0.00%
UTIL/PROC (H)	100.00	100.00	0.00	0.00%
UTIL/PROC	100.00	100.00	0.00	0.00%
TOTAL EMUL (H)	151.37	149.57	-1.80	-1.19%
TOTAL EMUL	154.00	152.00	-2.00	-1.30%
MASTER TOTAL (H)	100.00	100.00	0.00	0.00%
MASTER TOTAL	100.00	100.00	0.00	0.00%
MASTER EMUL (H)	70.14	69.84	-0.30	-0.43%
MASTER EMUL	71.00	71.00	0.00	0.00%
TVR(H)	1.32	1.34	0.02	1.20%
TVR	1.30	1.32	0.02	1.32%
Storage				
NUCLEUS SIZE (V)	2364KB	2540KB	176KB	7.45%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	41	63	22	53.66%
PGBLPGS	46916	55266	8350	17.80%
PGBLPGS/USER	25.2	29.7	4.5	17.80%
FREEPGS	5591	5430	-161	-2.88%
FREE UTIL	0.96	0.94	-0.02	-1.94%
SHRPGS	1258	1306	48	3.82%

Table 45 (Page 2 of 2). Proportional distribution of surplus share on the 9121-480

Release Run ID	VM/ESA 2.1 L26E186P	VM/ESA 2.2 L27E186F	Difference	%Difference
Environment				
Real Storage	**224MB**	**256MB**		
Exp. Storage	**32MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	1019	641	-378	-37.10%
WRITES/SEC	832	527	-305	-36.66%
PAGE/CMD	65.014	25.097	-39.918	-61.40%
PAGE IO RATE (V)	376.400	227.400	-149.000	-39.59%
PAGE IO/CMD (V)	13.221	4.886	-8.335	-63.04%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	7.025	9.433	2.408	34.28%
Queues				
DISPATCH LIST	103.66	79.91	-23.75	-22.91%
ELIGIBLE LIST	0.00	7.26	7.26	na
I/O				
VIO RATE	285	438	153	53.68%
VIO/CMD	10.010	9.411	-0.599	-5.98%
RIO RATE (V)	440	321	-119	-27.05%
RIO/CMD (V)	15.455	6.897	-8.557	-55.37%
NONPAGE RIO/CMD (V)	2.234	2.011	-0.223	-9.97%
DASD RESP TIME (V)	21.900	20.700	-1.200	-5.48%
MDC READS (blks)	158	na		
MDC READS (I/Os)	na	134		
MDC WRITES (blks)	62	na		
MDC WRITES (I/Os)	na	6.37		
MDC MODS	46	na		
MDC AVOID	75	126	51	68.00%
MDC HIT RATIO	0.90	0.93	0.03	3.33%
PRIVOPs				
PRIVOP/CMD	13.769	13.703	-0.066	-0.48%
DIAG/CMD	30.915	27.098	-3.817	-12.35%
DIAG 04/CMD	5.462	2.694	-2.768	-50.67%
DIAG 08/CMD	0.741	0.734	-0.007	-0.99%
DIAG 0C/CMD	1.328	1.122	-0.206	-15.54%
DIAG 14/CMD	0.025	0.024	0.000	-1.02%
DIAG 58/CMD	1.723	1.249	-0.473	-27.46%
DIAG 98/CMD	0.425	0.447	0.023	5.31%
DIAG A4/CMD	3.696	3.602	-0.095	-2.56%
DIAG A8/CMD	2.863	2.817	-0.046	-1.61%
DIAG 214/CMD	13.460	13.252	-0.209	-1.55%
SIE/CMD	75.165	58.014	-17.151	-22.82%
SIE INTCPT/CMD	34.576	31.908	-2.668	-7.72%
FREE TOTL/CMD	70.704	49.377	-21.328	-30.16%
VTAM Machines				
WKSET (V)	1014	1046	32	3.16%
TOT CPU/CMD (V)	4.8897	4.0586	-0.8311	-17.00%
CP CPU/CMD (V)	2.8832	2.0293	-0.8539	-29.62%
VIRT CPU/CMD (V)	2.0065	2.0293	0.0228	1.14%
DIAG 98/CMD (V)	0.424	0.447	0.023	5.51%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

SPXTAPE

The SPXTAPE CP command is provided in VM/ESA Release 2.2 as a high-performance alternative to the SPTAPE command. These commands are primarily used to dump and load spool files to/from tape. See "SPXTAPE Command" on page 7 for a discussion of the techniques SPXTAPE uses to achieve this improved performance.

A set of measurements was collected in order to assess the performance of the SPXTAPE command. Most of these measurements were done in a dedicated environment (no other system activity) and were used to compare SPXTAPE to the SPTAPE command. The primary metric was the elapsed time required to process a given number of spool files.

One of the SPXTAPE DUMP cases was then repeated while the system was running the FS8F CMS-intensive workload at 90% processor utilization in order to observe the interactions between SPXTAPE and concurrent system activity.

Dedicated Measurements

Workload: SPTAPE or SPXTAPE

Hardware Configuration

Processor model:	9121-742
Storage:	
Real:	1024MB
Expanded:	1024MB
Tapes:	3490 (1, 2, or 4 used)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	6	4		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.

Software Configuration

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
OPERATOR	1	SPXTAPE	24MB/XA	100		
SMART	1	RTM	16MB/370	3%	500	QUICKDSP ON

Note: RTM was only active for certain measurements (see results).

Additional Information

Spool files dumped/loaded:	34,000
Spool pages dumped/loaded:	353,000
Average pages per file:	10.4

Measurement Discussion: For the SPTAPE DUMP and SPXTAPE DUMP measurements, the system was prepared in advance by creating (approximately) 34,000 reader files on four 3390 spool volumes. The reader files ranged in size from 100 to 600,000 80-byte records and occupied, on average, 10.4 spool pages. The distribution of sizes was chosen so as to approximate the distribution of spool file sizes that was observed on a local CMS-intensive production system. The reader files were created concurrently by 8000 autologged users. SPTAPE or SPXTAPE were then used to dump these spool files to one or more 3490 tape drives.

For the SPTAPE and SPXTAPE LOAD measurements, the spool files were first purged from these same four spool volumes. The SPTAPE LOAD measurement was done by loading the tapes that had been created by SPTAPE DUMP. The SPXTAPE LOAD measurement was done by loading the tapes that had been created by the 1-tape SPXTAPE DUMP.

All required tape cartridges were preloaded into the 3490 cartridge stacker so that tape change time was as short and consistent as possible. Tape change time was observed to require about 1 minute and was mainly due to the time required to rewind the tape.

For each measurement, the timestamps provided in the summary log file created by SPTAPE or SPXTAPE were used to calculate the elapsed time. Additional timestamps in the summary log file were used to calculate the amount of time during which a tape drive was idle because a tape cartridge was being changed (*Tape Change Time*) and the amount of time this caused the system to be idle (*System Idle Time*).

RTM VM/ESA was active for some of the measurements. In those cases, the results table (Table 46) also includes processor time and processor utilization data.

Table 46. SPTAPE and SPXTAPE results — dedicated 9121-742

Command **Function** **Tape Drives**	**SPTAPE DUMP 1**	**SPXTAPE DUMP 1**	**SPXTAPE DUMP 2**	**SPXTAPE DUMP 4**	**SPTAPE LOAD 1**	**SPXTAPE LOAD 1**
Elapsed Time	3:24:39	22:42	19:27	18:55	2:09:32	1:07:42
Elapsed Time (sec)	12279	1362	1167	1135	7772	4062
Elapsed Time Ratio	1.000	0.111	0.095	0.092	1.000	0.52
Rate Ratio	1.00	9.02	10.52	10.82	1.00	1.91
Tape Change Time (sec)	653	371	309	247	671	371
System Idle Time (sec)	653	230	0	0	671	0
Overlapped Tape Change Time (sec)	0	141	309	247	0	371
Percentage Overlap	0	38	100	100	0	100
Tapes	12	7	7	8	12	7
Percentage Reduction		42	42	33		42
Processor time (sec)			79			902
Processor utilization (%)[10]			1.7			5.6

Elapsed Time Ratio was calculated as SPXTAPE *Elapsed Time* divided by SPTAPE *Elapsed Time*. *Rate Ratio* was calculated as the reciprocal of *Elapsed*

Time Ratio. *Rate Ratio* represents how many times faster SPXTAPE ran relative to the corresponding SPTAPE measurement.

Overlapped Tape Change Time was calculated by subtracting *System Idle Time* from *Tape Change Time*. *Percentage Overlap* was calculated as 100% * *Overlapped Tape Change Time* divided by *Tape Change Time*. This represents the percentage of time that changing tape cartridges was overlapped with other processing.

As shown in the above table, SPXTAPE was 9.0 to 10.8 times faster than SPTAPE for dumping spool files for the measured cases. The 2-tape and 4-tape cases were somewhat better than the 1-tape case because tape changes were completely overlapped with other processing. SPXTAPE was 1.9 times faster than SPTAPE for loading spool files for the measured 1-tape case.

The number of spool volumes is an important variable in determining how long it will take SPXTAPE to DUMP or LOAD a set of spool files. The more spool volumes there are, the faster SPXTAPE will tend to run because the spool I/Os are done in parallel across a larger number of devices. By contrast, SPTAPE only works with one spool volume at a time. Consequently, the performance relationship between SPTAPE and SPXTAPE will depend, among other things, upon the number of spool volumes. Four spool volumes were used in this test and resulted in roughly a 10:1 elapsed time improvement for dumping spool files. Cases where fewer spool volumes are involved will tend to see a smaller improvement than this, while cases with more spool volumes will tend to see a larger improvement.

SPXTAPE DUMP time is also sensitive to the distribution of spool file blocks across the spool volumes. This was illustrated by the results obtained from a 2-tape SPXTAPE DUMP of the same 34,000 files that had been built onto the 4 spool volumes by restoring from the SPTAPE dump tapes. Even though the files were the same, the elapsed time was 70% longer than the 1167 seconds shown in Table 46 on page 151 because those files were less favorably distributed across the spool volumes.

One of the ways in which SPXTAPE provides better performance relative to SPTAPE is by reducing the amount of time it is idle waiting for a tape to be changed. This is shown by the reduction in *System Idle Time* in the results table. Three factors contribute to this improvement:

1. SPXTAPE uses fewer tapes, resulting in fewer tape changes.
2. Unlike SPTAPE, SPXTAPE supports the use of multiple tape drives. When two or more tape drives are made available, SPXTAPE is able to continue processing with another tape while one tape is being changed.
3. SPXTAPE uses available real memory to buffer data. For example, during a SPXTAPE DUMP to one tape drive, when that tape drive becomes temporarily unavailable due to a tape cartridge change, SPXTAPE continues

[10] In the table, processor utilization is calculated as processor time divided by total processing capacity, where total processing capacity is elapsed time * 4. Because the 9121-742 has 4 processors, its total capacity is 4 processor-seconds per elapsed second.

to read spool files from DASD and stages them in real storage. This buffering continues until there is no more available real storage or the tape drive becomes ready again.

SPXTAPE looks at system paging indicators to determine how much of the system′s total real storage it can use. In these dedicated measurements, SPXTAPE was able to use large amounts of real storage for this buffering, resulting in substantial overlapping of processing with tape changes. For the 1-tape DUMP, it was able to overlap 38% of the tape change time. For the 1-tape LOAD, it overlapped all of the tape change time. The degree of overlap would be much less, of course, on systems with less real memory or in cases where there is concurrent activity that has significant memory requirements.

The results show that SPXTAPE elapsed time was relatively insensitive to the number of tape drives that were used. Elapsed time decreased by 14% when going from 1 tape to 2 tapes. There was very little further elapsed time reduction when going from 2 tapes to 4 tapes. This is because, with only four spool volumes, the limiting factor tended to be the time required to transfer the spool files from/to DASD.

SPXTAPE DUMP and LOAD processing result in low processor utilizations. Even though the I/O processing is heavily overlapped, the SPXTAPE DUMP and LOAD functions are I/O-bound.

The SPXTAPE LOAD measurement was run with default options in effect. SPXTAPE offers a NODUP option. If selected, SPXTAPE ensures that each loaded spool file is not a duplicate of any spool file that is already on the system. Use of this option can increase processing requirements significantly because each incoming spool file must be checked against all existing spool files.

SPTAPE writes a tape mark between each backed up spool file. The smaller the files, the more tape space is taken up by these tape marks. SPXTAPE writes the spool file data as one tape file consisting of 32K blocks. This reduces the number of tape volumes required to hold the spool files. The smaller the average spool file size, the larger the reduction relative to SPTAPE. For the measured environment, the average spool file was 10.4 pages and a 42% reduction in tapes was observed (1-tape case).

One disadvantage of using multiple tape drives with SPXTAPE is that it can increase the number of tapes required. For example, the SPXTAPE DUMP to 1 tape drive required 7 tapes, while the SPXTAPE DUMP to 4 tape drives required 8 tapes. Using “n” tape drives means that there are “n” partially filled tapes when the dump has completed. Because of this, it is better to use no more tape drives than is necessary to keep DASD I/O as the limiting factor, with a minimum of two tapes in order to get the tape change overlap benefits. For the measured case, 2 tape drives was a good number. For a case where there are far more spool volumes, using more than two tape drives can be beneficial. The suggested rule-of-thumb is one tape drive for every 4 to 6 spool volumes, with a minimum of two.

SPXTAPE will tend to perform especially well relative to SPTAPE when one or more of the following conditions apply:

- The spool files are spread across many spool volumes.
- Two tape drives are used (or more, if appropriate).
- Tape change times are long.
- The spool files are small.

Interaction with Concurrent Activity

A measurement was done to explore the degree to which SPXTAPE DUMP can affect the response times of concurrent CMS interactive work and the degree to which that activity can affect the elapsed time required to complete the SPXTAPE DUMP function.

Workload: FS8F0R + SPXTAPE DUMP

Hardware Configuration

Processor model:	9121-742
Processors used:	4
Storage:	
Real:	1024MB
Expanded:	1024MB
Tapes:	3480 (Monitor)
	3490 (2 tapes, used by SPXTAPE)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-3	4	6	7	7	32 R		2 R
3390-2	3990-2	4	16	6	6			

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

Note: The spool files backed up by SPXTAPE were all contained on 4 of the spool volumes behind a 3990-2 control unit.

Communications:

Control Unit	*Number*	*Lines per Control Unit*	*Speed*
3088	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*	*SHARE*	*RESERVED*	*Other Options*
OPERATOR	1	SPXTAPE	24MB/XA	100		
SMART	1	RTM	16MB/370	3%	500	QUICKDSP ON
VSCSn	3	VSCS	64MB/XA	10000	1200	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	550	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	5500	Users	3MB/XC	100		

Additional Information

Spool files backed up: 26,000
Spool pages backed up: 264,000
Average pages per file: 10.2
Tapes required: 6

Measurement Discussion: The 34,000 spool files were created on the system in advance, as described in “Dedicated Measurements” on page 150. The same four spool volumes were used. The system was then configured for a standard FS8F CMS-intensive workload measurement. This configuration had 9 additional spool volumes, for a total of 13 spool volumes. These additional spool volumes accommodate the spool activity generated by the FS8F workload.

The stabilization period for a 5500 user FS8F measurement was allowed to complete. An SPXTAPE DUMP to two 3490 tape drives was then started and remained active during the entire time that a FS8F measurement was obtained. A subset (26,000) of the 34,000 spool files was dumped. The FS8F measurement interval is usually 30 minutes for 9121-742 measurements. However, in this case, the measurement interval was ended when the SPXTAPE DUMP command completed, resulting in an FS8F measurement interval of about 15 minutes.[11]

Table 47 on page 156 compares the results of this measurement to a corresponding measurement without SPXTAPE activity. The results show the impact of the SPXTAPE DUMP activity on the performance of the overall system. Average external response time (AVG LAST(T)) increased from 0.3 seconds to about 1.7 seconds when the SPXTAPE DUMP was active. This increase was mainly due to I/O contention on the spool volumes. SPXTAPE processor and real storage usage were not significant factors. In both cases, they were a small percentage of the total system resources.

The rate at which SPXTAPE was able to dump spool pages decreased from 303 pages/second in the dedicated case to 237 pages/second when the system was running at 90% processor utilization with the FS8F workload, a 22% decrease. This decrease is primarily due to contention with the CMS users for processor time and the spool volumes.

[11] This shorter run interval is also the reason for the apparent decrease in think time between the two measurements.

SPXTAPE

Table 47 (Page 1 of 2). Performance interactions between SPXTAPE DUMP and concurrent CMS users

SPXTAPE DUMP **Release** **Run ID**	**NO** **VM/ESA 2.2** **S47E5504**	**YES** **VM/ESA 2.2** **S47E5506**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Response Time				
TRIV INT	0.110	0.079	-0.031	-28.18%
NONTRIV INT	0.337	0.466	0.129	38.28%
TOT INT	0.250	0.292	0.042	16.80%
TOT INT ADJ	0.233	0.341	0.108	46.53%
AVG FIRST (T)	0.231	0.531	0.300	130.11%
AVG LAST (T)	0.308	1.661	1.353	439.88%
Throughput				
AVG THINK (T)	26.07	21.57	-4.50	-17.28%
ETR	179.76	218.37	38.61	21.48%
ETR (T)	193.12	187.01	-6.12	-3.17%
ETR RATIO	0.931	1.168	0.237	25.45%
ITR (H)	218.42	213.07	-5.36	-2.45%
ITR	50.91	62.29	11.38	22.35%
EMUL ITR	76.29	95.86	19.56	25.64%
ITRR (H)	1.000	0.975	-0.025	-2.45%
ITRR	1.000	1.224	0.224	22.35%
Proc. Usage				
PBT/CMD (H)	18.313	18.773	0.461	2.51%
PBT/CMD	18.279	18.769	0.491	2.69%
CP/CMD (H)	6.490	6.994	0.505	7.78%
CP/CMD	6.058	6.577	0.519	8.57%
EMUL/CMD (H)	11.823	11.779	-0.044	-0.37%
EMUL/CMD	12.220	12.192	-0.028	-0.23%
Processor Util.				
TOTAL (H)	353.66	351.08	-2.59	-0.73%
TOTAL	353.00	351.00	-2.00	-0.57%
UTIL/PROC (H)	88.42	87.77	-0.65	-0.73%
UTIL/PROC	88.25	87.75	-0.50	-0.57%
TOTAL EMUL (H)	228.33	220.28	-8.06	-3.53%
TOTAL EMUL	236.00	228.00	-8.00	-3.39%
MASTER TOTAL (H)	90.85	89.80	-1.06	-1.16%
MASTER TOTAL	91.00	90.00	-1.00	-1.10%
MASTER EMUL (H)	38.95	36.42	-2.53	-6.50%
MASTER EMUL	40.00	38.00	-2.00	-5.00%
TVR(H)	1.55	1.59	0.04	2.90%
TVR	1.50	1.54	0.04	2.92%
Storage				
NUCLEUS SIZE (V)	2540KB	2544KB	4KB	0.16%
TRACE TABLE (V)	800KB	800KB	0KB	0.00%
WKSET (V)	72	72	0	0.00%
PGBLPGS	233K	230K	-3K	-1.29%
PGBLPGS/USER	42.4	41.8	-0.5	-1.29%
FREEPGS	15585	18180	2595	16.65%
FREE UTIL	0.92	0.90	-0.02	-2.03%
SHRPGS	1765	1605	-160	-9.07%

Table 47 (Page 2 of 2). Performance interactions between SPXTAPE DUMP and concurrent CMS users

SPXTAPE DUMP Release Run ID	NO VM/ESA 2.2 S47E5504	YES VM/ESA 2.2 S47E5506	Difference	%Difference
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Paging				
READS/SEC	623	502	-121	-19.42%
WRITES/SEC	455	453	-2	-0.44%
PAGE/CMD	5.582	5.107	-0.475	-8.51%
PAGE IO RATE (V)	162.300	158.300	-4.000	-2.46%
PAGE IO/CMD (V)	0.840	0.846	0.006	0.73%
XSTOR IN/SEC	828	797	-31	-3.74%
XSTOR OUT/SEC	1409	1302	-107	-7.59%
XSTOR/CMD	11.583	11.224	-0.359	-3.10%
FAST CLR/CMD	8.813	8.545	-0.268	-3.04%
Queues				
DISPATCH LIST	102.55	126.30	23.74	23.15%
ELIGIBLE LIST	0.02	0.00	-0.02	-100.00%
I/O				
VIO RATE	1806	1772	-34	-1.88%
VIO/CMD	9.352	9.476	0.124	1.33%
RIO RATE (V)	547	793	246	44.97%
RIO/CMD (V)	2.832	4.240	1.408	49.71%
NONPAGE RIO/CMD (V)	1.992	3.394	1.402	70.38%
DASD RESP TIME (V)	19.900	18.400	-1.500	-7.54%
MDC READS (I/Os)	557	545	-12	-2.15%
MDC WRITES (I/Os)	26	26	0	0.00%
MDC AVOID	515	505	-10	-1.94%
MDC HIT RATIO	0.91	0.91	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	20.570	20.465	-0.104	-0.51%
DIAG/CMD	25.324	25.411	0.087	0.34%
DIAG 04/CMD	0.948	0.988	0.041	4.30%
DIAG 08/CMD	0.738	0.685	-0.053	-7.17%
DIAG 0C/CMD	1.125	1.106	-0.019	-1.68%
DIAG 14/CMD	0.025	0.025	0.000	0.99%
DIAG 58/CMD	1.249	1.261	0.012	0.92%
DIAG 98/CMD	0.345	0.355	0.010	2.96%
DIAG A4/CMD	3.592	3.648	0.056	1.56%
DIAG A8/CMD	2.824	2.910	0.086	3.06%
DIAG 214/CMD	13.333	13.292	-0.041	-0.31%
SIE/CMD	56.959	58.821	1.863	3.27%
SIE INTCPT/CMD	38.162	38.822	0.660	1.73%
FREE TOTL/CMD	44.821	45.480	0.658	1.47%
VTAM Machines				
WKSET (V)	4137	4159	22	0.53%
TOT CPU/CMD (V)	2.7703	2.7883	0.0180	0.65%
CP CPU/CMD (V)	1.2312	1.2668	0.0356	2.89%
VIRT CPU/CMD (V)	1.5390	1.5215	-0.0175	-1.14%
DIAG 98/CMD (V)	0.345	0.356	0.012	3.40%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

ISFC Changes

There were several changes to the Inter-System Facility for Communications (ISFC) component of CP in VM/ESA Release 2.2. These include:

- more granular multiprocessor locking

  Prior to VM/ESA Release 2.2, all ISFC activity was serialized by the ISFC global lock. With VM/ESA Release 2.2, each active link has a lock associated with it and the I/O-related functions of ISFC are now serialized by this link lock instead of by the ISFC global lock. This change reduces contention for the ISFC global lock, thus improving responsiveness and increasing the maximum amount of message traffic that ISFC can handle when there are multiple active links.

- distributed IUCV support
- broadcast routing

  This removes the restriction that each node must have a direct connection to all other nodes in the CS collection. With broadcast routing, system A can communicate with system B through intermediate node C.

- structural changes for improved reliability and maintainability

These changes affected ISFC's performance characteristics. This section includes performance measurements that quantify those effects.

Workload: CONN1

Hardware Configuration

The hardware configuration consists of two identically configured 9121-621 systems connected by a CTC adapter. The following configuration information applies to both of these systems.

Processor model:	9121-621[12]
Processors used:	2
Storage:	
Real:	512MB
Expanded:	512MB
Tape:	3480 (Monitor)

DASD: Not shown; there were no DASD I/Os during the measurement interval.

Communications:

Control Unit	*Number*	*Lines per Control Unit*	*Speed*
3088	1	NA	4.5MB

[12] See "Hardware Used" on page 17 for an explanation of how this processor model was defined.

Software Configuration

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	500	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
For the user system:						
Unnnn	30	Users	24MB/XA	100		
For the server system:						
Unnnn	2	Servers	24MB/XA	100		

Measurement Discussion: Two 9121-621 systems were connected via CTCA. On one of these systems (user system), 30 user virtual machines were started. On the other system (server system), two server virtual machines were started. Half of the users communicated with one of the servers, while the other half communicated with the other server. Two server virtual machines were used (instead of one) in order to have concurrent ISFC activity on the server system′s two processors, thus creating the potential for ISFC lock contention in the server system.

After being started, each user machine did a large number of consecutive APPC/VM interactions with its assigned server and then became idle. Each APPC/VM interaction consisted of the following sequence:

- User: SENDDATA RECEIVE=YES, WAIT=YES
- Server: (implicit RECEIVE)
- Server: SENDDATA RECEIVE=YES, WAIT=NO. Sends 1 byte or 4KB back to the user.

Each such iteration is referred to as a “command” in the results tables. For additional workload information, see “Connectivity (CONN1)” on page 223.

The number of iterations was chosen such that all 30 users were active for about 20 minutes. Measurement data were collected for a 15 minute interval after all 30 users had been started. CP monitor and RTM data were collected on both systems. Hardware monitor and QUERY ISLINK data were collected on the server system.

Two pairs of measurements were obtained. Each pair compared VM/ESA Release 2.2 to VM/ESA Release 2.1. One pair of measurements was done with a 1 byte server reply, while the other pair was with a 4KB server reply. The results are summarized in the following two tables. The 1 byte reply results are in Table 48 on page 161, while the 4KB reply results are in Table 49 on page 161.

Commands/sec was calculated by adding together the IUCV receives per second for the two server machines. This information was taken from the USER_TO_USER_VMCOMM VMPRF report for the monitor data that was collected on the server system.

ISFC Changes

KB/sec was calculated using the QUERY ISLINK data. The QUERY TIME and QUERY ISLINK CP commands were issued before and after the 15 minute measurement interval. The difference in elapsed time, bytes sent, and bytes received were determined by subtraction. The delta bytes sent and received were added together, converted to KB, and divided by delta seconds. *KB/sec* includes system overhead bytes in addition to the actual message bytes.

The CPU utilizations were obtained from the PROCESSORS_BY_TIME VMPRF reports obtained for the user and server systems. Total utilization and emulation utilization were taken directly from these reports, while CP utilization was derived by subtracting emulation utilization from total utilization. Only total utilization is shown in the results tables, but all three utilizations were used, in conjunction with *Commands/sec*, to derive the corresponding CPU/command results[13].

The locking changes resulted in significant throughput improvements. *Commands/sec* increased 18.6% in the 1 byte reply case and 7.9% in the 4KB reply case. The degree of throughput improvement depends on the degree of loading. The amount of improvement will be negligible when the ISFC rate is low, while the amount of improvement can exceed what is shown here if the degree of loading is increased further.

The various changes that were made to ISFC in VM/ESA Release 2.2 increased ISFC processor usage. Total CPU/command, when summed over both systems, increased 12.7% in the 1 byte reply case and 12.2% in the 4KB case. The amount of increase should be essentially independent of the degree of loading.

---

[13] CPU/command = (1000*CPU_utilization*2/100)/commands_per_sec. The numerator is in units of CPU-milliseconds per elapsed second. The denominator is in units of commands per elapsed second. The CPU utilization is multiplied by 2 because the system has two processors.

Table 48. ISFC performance for two CTCA-attached 9121-621 processors, 1 byte reply

Release Run ID	VM/ESA 2.1 S26C0032	VM/ESA 2.2 S27C0031	Difference	%Difference
Throughput				
Commands/sec	1122.1	1330.5	208.4	18.6
KB/sec	161.8	191.2	29.4	18.2
KB/command	0.144	0.144	0.000	0.0
CPU Utilization				
User system	24.4	33.1	8.7	35.7
Server system	28.8	38.0	9.2	31.9
CPU/command (msec)				
User system				
Total	0.435	0.498	0.063	14.4
CP	0.428	0.492	0.064	14.9
Emulation	0.007	0.006	-0.001	-15.7
Server system				
Total	0.513	0.571	0.058	11.3
CP	0.456	0.501	0.044	9.7
Emulation	0.057	0.071	0.014	23.9
Overall				
Total	0.948	1.069	0.121	12.7
CP	0.884	0.992	0.108	12.2
Emulation	0.064	0.077	0.012	19.5

Table 49. ISFC performance for two CTCA-attached 9121-621 processors, 4KB reply

Release Run ID	VM/ESA 2.1 S26C0033	VM/ESA 2.2 S27C0030	Difference	%Difference
Throughput				
Commands/sec	374.4	403.9	29.5	7.9
KB/sec	1573.8	1692.3	118.5	7.5
KB/command	4.203	4.190	-0.013	-0.3
CPU Utilization				
User system	10.3	12.9	2.7	25.9
Server system	12.1	14.2	2.1	16.9
CPU/command (msec)				
User system				
Total	0.547	0.639	0.091	16.7
CP	0.539	0.631	0.092	17.0
Emulation	0.008	0.007	-0.001	-7.3
Server system				
Total	0.646	0.701	0.054	8.4
CP	0.572	0.626	0.055	9.6
Emulation	0.075	0.074	-0.001	-0.7
Overall				
Total	1.194	1.339	0.146	12.2
CP	1.111	1.258	0.147	13.2
Emulation	0.083	0.082	-0.001	-1.3

Improved CMS Block Allocation for Virtual Disk in Storage

This section documents the enhancements made to the virtual disk in storage support added in VM/ESA Release 2.1.

The CMS disk allocation for minidisks uses a moving cursor algorithm to improve disk access times. It allocates slots from the start of a minidisk and continues down the minidisk even if earlier slots are freed up as a result of erasing files. This reduces file fragmentation, resulting in reduced DASD access times. Free disk slots are not reused until a complete pass across the minidisk has been completed. This leaves previously used slots empty.

This method, however, is not best when using virtual disks in storage. With virtual disks in storage, a better method is to re-use the slots (pages) rather than to select new ones. If a file is erased, CP does not know that the contents of a page are no longer needed and keeps the page valid, causing it to eventually be paged out. By selecting slots from the beginning, when using virtual disks in storage, CMS now reuses these empty slots. This results in a smaller working set for the virtual disk address space.

This is only done for virtual disks. There is no seek or latency time required to access data in the address space, as opposed to a real disk, so fragmentation of files is not a concern. The new method reduces the storage requirements to support virtual disks in storage, thereby reducing overall system storage demand.

Workload

VM/ESA 2.1: FS7F0R
VM/ESA 2.2: FS8F0R

Hardware Configuration

Processor model:	9121-480
Processors used:	2
Storage, VM/ESA 2.1:	
Real:	224MB
Expanded:	32MB (all for MDC)
Storage, VM/ESA 2.2:	
Real:	256MB (default MDC)
Expanded:	0MB
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes -* *PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2	16 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
3088-08	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	SHARE	RESERVED	Other Options
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XC	10000	512	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1860	Users	3MB/XC	100		

Measurement Discussion: The following table was extracted from the VM/ESA Release 2.1 report. It contains the virtual disks in storage data for temporary minidisks (T-disks) held long term. The B-disk within our workload represents this type of activity. This table shows that, for VM/ESA Release 2.1, overall system performance decreased when a virtual disk in storage was substituted for the B-disk. External response time (AVG LAST (T)) increased 23% and internal throughput (ITR (H)) decreased by almost 2%. The decrease in minidisk I/O was overshadowed by the increase in page I/Os caused by the increased storage demand.

Also shown in the VM/ESA Release 2.1 report (not shown here) is data for the case where virtual disks in storage are substituted for T-disks that are held for short durations. This case showed improved overall system performance. The real storage associated with a private virtual disk in storage is released when the disk is detached.

For VM/ESA Release 2.1, the real storage pages used by virtual disks in storage were subtracted from the pageable page count (PGBLPGS). For VM/ESA Release 2.2, this has been corrected by only subtracting valid page and segment table storage.

Table 51 on page 166 contains the data from a similar set of measurements made on VM/ESA Release 2.2. Although the workloads are slightly different, the VM/ESA Release 2.1 differences can be compared to the VM/ESA Release 2.2 differences.

The number of DASD page slots used for the virtual disks in storage can be used to demonstrate the effects of this enhancement. From the DATA_SPACES_SHARED report in VMPRF (not shown in the tables), it required an average of 72 page slots to back the storage used for each B-disk in the VM/ESA Release 2.1 measurement. For VM/ESA Release 2.2, this was reduced to an average of 18 page slots.

Improved CMS Block Allocation

Because of these reduced real storage requirements, the VM/ESA Release 2.2 comparison shows equivalent performance when virtual disks in storage are substituted for minidisks that are held long term. This result indicates that it is now more practical to replace T-disks with virtual disks in storage.

A list of virtual disk in storage usage guidelines is provided in *VM/ESA: Performance*, SC24-5642.

Table 50 (Page 1 of 3). Virtual disks in storage for long term usage on VM/ESA Release 2.1 on the 9121-480

B-disk Implementation **Release** **Run ID**	**Minidisk** **VM/ESA 2.1** **L26E1868**	**Virtual Disk** **VM/ESA 2.1** **L26E1866**	**Difference**	**%Difference**
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.123	0.141	0.018	14.63%
NONTRIV INT	0.453	0.519	0.066	14.57%
TOT INT	0.344	0.389	0.045	13.08%
TOT INT ADJ	0.303	0.353	0.049	16.23%
AVG FIRST (T)	0.247	0.310	0.064	25.71%
AVG LAST (T)	0.378	0.464	0.087	22.88%
Throughput				
AVG THINK (T)	26.29	26.26	-0.03	-0.13%
ETR	57.57	58.91	1.34	2.33%
ETR (T)	65.25	64.96	-0.29	-0.45%
ETR RATIO	0.882	0.907	0.025	2.79%
ITR (H)	73.02	71.60	-1.41	-1.94%
ITR	32.22	32.47	0.25	0.79%
EMUL ITR	47.55	48.97	1.42	2.98%
ITRR (H)	1.000	0.981	-0.019	-1.94%
ITRR	1.000	1.008	0.008	0.79%
Proc. Usage				
PBT/CMD (H)	27.391	27.932	0.541	1.98%
PBT/CMD	27.431	27.862	0.431	1.57%
CP/CMD (H)	9.416	10.007	0.592	6.28%
CP/CMD	8.888	9.390	0.502	5.64%
EMUL/CMD (H)	17.975	17.924	-0.050	-0.28%
EMUL/CMD	18.543	18.472	-0.071	-0.38%
Processor Util.				
TOTAL (H)	178.74	181.46	2.72	1.52%
TOTAL	179.00	181.00	2.00	1.12%
UTIL/PROC (H)	89.37	90.73	1.36	1.52%
UTIL/PROC	89.50	90.50	1.00	1.12%
TOTAL EMUL (H)	117.29	116.44	-0.85	-0.72%
TOTAL EMUL	121.00	120.00	-1.00	-0.83%
MASTER TOTAL (H)	89.04	90.40	1.36	1.53%
MASTER TOTAL	89.00	90.00	1.00	1.12%
MASTER EMUL (H)	52.67	51.70	-0.97	-1.84%
MASTER EMUL	54.00	54.00	0.00	0.00%
TVR(H)	1.52	1.56	0.03	2.26%
TVR	1.48	1.51	0.03	1.96%

Table 50 (Page 2 of 3). Virtual disks in storage for long term usage on VM/ESA Release 2.1 on the 9121-480

B-disk Implementation Release Run ID	Minidisk VM/ESA 2.1 L26E1868	Virtual Disk VM/ESA 2.1 L26E1866	Difference	%Difference
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Storage				
NUCLEUS SIZE (V)	2364KB	2364KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	82	70	-12	-14.63%
PGBLPGS	47770	19740	-28030	-58.68%
PGBLPGS/USER	25.7	10.6	-15.1	-58.68%
FREEPGS	4953	5211	258	5.21%
FREE UTIL	0.93	0.93	0.00	0.33%
SHRPGS	1134	1117	-17	-1.50%
Paging				
READS/SEC	625	809	184	29.44%
WRITES/SEC	422	583	161	38.15%
PAGE/CMD	16.045	21.427	5.382	33.55%
PAGE IO RATE (V)	170.200	240.600	70.400	41.36%
PAGE IO/CMD (V)	2.608	3.704	1.095	42.00%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	6.467	6.896	0.429	6.64%
Queues				
DISPATCH LIST	37.02	42.13	5.11	13.80%
ELIGIBLE LIST	0.00	0.02	0.02	na
I/O				
VIO RATE	627	613	-14	-2.23%
VIO/CMD	9.608	9.436	-0.173	-1.80%
RIO RATE (V)	391	421	30	7.67%
RIO/CMD (V)	5.992	6.480	0.489	8.15%
NONPAGE RIO/CMD (V)	3.384	2.776	-0.608	-17.97%
MDC READS	361	342	-19	-5.26%
MDC WRITES	179	132	-47	-26.26%
MDC MODS	145	100	-45	-31.03%
MDC HIT RATIO	0.91	0.91	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	17.584	17.658	0.074	0.42%
DIAG/CMD	26.983	28.239	1.256	4.65%
DIAG 08/CMD	0.760	0.759	-0.001	-0.14%
DIAG 14/CMD	0.025	0.025	0.000	-0.49%
DIAG 58/CMD	1.251	1.250	0.000	-0.04%
DIAG 98/CMD	1.227	1.072	-0.156	-12.70%
DIAG A4/CMD	3.913	3.875	-0.038	-0.97%
DIAG A8/CMD	1.888	1.915	0.027	1.44%
DIAG 214/CMD	12.941	12.657	-0.284	-2.20%
SIE/CMD	54.586	57.447	2.861	5.24%
SIE INTCPT/CMD	37.664	39.064	1.400	3.72%
FREE TOTL/CMD	64.470	66.514	2.043	3.17%

Improved CMS Block Allocation

Table 50 (Page 3 of 3). Virtual disks in storage for long term usage on VM/ESA Release 2.1 on the 9121-480

B-disk Implementation Release Run ID	Minidisk VM/ESA 2.1 L26E1868	Virtual Disk VM/ESA 2.1 L26E1866	Difference	%Difference
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
VTAM Machines				
WKSET (V)	515	517	2	0.39%
TOT CPU/CMD (V)	3.8737	3.8141	-0.0596	-1.54%
CP CPU/CMD (V)	1.6431	1.6334	-0.0097	-0.59%
VIRT CPU/CMD (V)	2.2306	2.1807	-0.0499	-2.24%
DIAG 98/CMD (V)	1.227	1.072	-0.156	-12.69%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Table 51 (Page 1 of 3). Virtual disks in storage for long term usage on VM/ESA Release 2.2 on the 9121-480

B-disk Implementation Release Run ID	Minidisk VM/ESA 2.2 L27E186J	Virtual Disk VM/ESA 2.2 L27E186K	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.127	0.128	0.001	0.79%
NONTRIV INT	0.396	0.395	-0.001	-0.25%
TOT INT	0.305	0.305	0.000	0.00%
TOT INT ADJ	0.269	0.269	0.000	-0.12%
AVG FIRST (T)	0.245	0.248	0.002	0.81%
AVG LAST (T)	0.336	0.334	-0.002	-0.74%
Throughput				
AVG THINK (T)	26.19	26.15	-0.04	-0.13%
ETR	57.66	57.48	-0.18	-0.31%
ETR (T)	65.36	65.24	-0.12	-0.19%
ETR RATIO	0.882	0.881	-0.001	-0.12%
ITR (H)	75.15	74.86	-0.28	-0.38%
ITR	33.16	33.00	-0.15	-0.46%
EMUL ITR	48.09	48.13	0.04	0.09%
ITRR (H)	1.000	0.996	-0.004	-0.38%
ITRR	1.000	0.995	-0.005	-0.46%
Proc. Usage				
PBT/CMD (H)	26.615	26.716	0.100	0.38%
PBT/CMD	26.620	26.670	0.050	0.19%
CP/CMD (H)	8.849	8.975	0.126	1.42%
CP/CMD	8.261	8.430	0.169	2.04%
EMUL/CMD (H)	17.766	17.740	-0.025	-0.14%
EMUL/CMD	18.359	18.240	-0.119	-0.65%

Table 51 (Page 2 of 3). Virtual disks in storage for long term usage on VM/ESA Release 2.2 on the 9121-480

B-disk Implementation Release Run ID	Minidisk VM/ESA 2.2 L27E186J	Virtual Disk VM/ESA 2.2 L27E186K	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Processor Util.				
TOTAL (H)	173.97	174.30	0.33	0.19%
TOTAL	174.00	174.00	0.00	0.00%
UTIL/PROC (H)	86.98	87.15	0.16	0.19%
UTIL/PROC	87.00	87.00	0.00	0.00%
TOTAL EMUL (H)	116.12	115.74	-0.38	-0.33%
TOTAL EMUL	120.00	119.00	-1.00	-0.83%
MASTER TOTAL (H)	86.47	86.61	0.14	0.16%
MASTER TOTAL	86.00	87.00	1.00	1.16%
MASTER EMUL (H)	51.58	50.51	-1.07	-2.07%
MASTER EMUL	53.00	52.00	-1.00	-1.89%
TVR(H)	1.50	1.51	0.01	0.52%
TVR	1.45	1.46	0.01	0.84%
Storage				
NUCLEUS SIZE (V)	2556KB	2556KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	84	83	-1	-1.19%
PGBLPGS	55460	51539	-3921	-7.07%
PGBLPGS/USER	29.8	27.7	-2.1	-7.07%
FREEPGS	5282	5502	220	4.17%
FREE UTIL	0.92	0.93	0.01	1.05%
SHRPGS	1231	1231	0	0.00%
Paging				
READS/SEC	629	639	10	1.59%
WRITES/SEC	436	449	13	2.98%
PAGE/CMD	16.293	16.677	0.383	2.35%
PAGE IO RATE (V)	176.700	180.400	3.700	2.09%
PAGE IO/CMD (V)	2.703	2.765	0.062	2.29%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	8.720	8.737	0.016	0.19%
Queues				
DISPATCH LIST	36.13	39.48	3.35	9.28%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	663	661	-2	-0.30%
VIO/CMD	10.143	10.132	-0.012	-0.11%
RIO RATE (V)	364	341	-23	-6.32%
RIO/CMD (V)	5.569	5.227	-0.342	-6.14%
NONPAGE RIO/CMD (V)	2.865	2.462	-0.404	-14.09%
DASD RESP TIME (V)	19.500	19.700	0.200	1.03%
MDC REAL SIZE (MB)	31.3	40.9	9.6	30.67%
MDC READS (I/Os)	187	178	-9	-4.81%
MDC WRITES (I/Os)	9.38	3.04	-6.34	-67.59%
MDC AVOID	176	169	-7	-3.98%
MDC HIT RATIO	0.94	0.94	0.00	0.00%

Improved CMS Block Allocation

Table 51 (Page 3 of 3). Virtual disks in storage for long term usage on VM/ESA Release 2.2 on the 9121-480

B-disk Implementation **Release** **Run ID**	**Minidisk** **VM/ESA 2.2** **L27E186J**	**Virtual Disk** **VM/ESA 2.2** **L27E186K**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
PRIVOPs				
PRIVOP/CMD	13.905	13.844	-0.062	-0.44%
DIAG/CMD	27.608	29.017	1.408	5.10%
DIAG 04/CMD	2.472	4.134	1.662	67.26%
DIAG 08/CMD	0.739	0.737	-0.002	-0.26%
DIAG 0C/CMD	1.126	1.127	0.000	0.03%
DIAG 14/CMD	0.025	0.024	0.000	-0.93%
DIAG 58/CMD	1.249	1.246	-0.003	-0.21%
DIAG 98/CMD	1.163	1.148	-0.015	-1.25%
DIAG A4/CMD	3.585	3.594	0.009	0.26%
DIAG A8/CMD	2.834	2.824	-0.009	-0.33%
DIAG 214/CMD	13.266	13.029	-0.238	-1.79%
SIE/CMD	53.592	54.904	1.312	2.45%
SIE INTCPT/CMD	34.835	36.786	1.951	5.60%
FREE TOTL/CMD	49.538	52.221	2.684	5.42%
VTAM Machines				
WKSET (V)	508	509	1	0.20%
TOT CPU/CMD (V)	3.8587	3.8404	-0.0183	-0.47%
CP CPU/CMD (V)	1.4619	1.4561	-0.0058	-0.40%
VIRT CPU/CMD (V)	2.3968	2.3843	-0.0125	-0.52%
DIAG 98/CMD (V)	1.162	1.149	-0.014	-1.17%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Additional Evaluations

This portion of the report includes results from a variety of additional VM/ESA performance measurement evaluations that have been conducted over the past year.

- "Processor Capacity: 9021-941" on page 170 examines how VM/ESA performance scales when the CMS-intensive workload is run on the 9021-941 using 1, 2, 3, and 4 of the processors.
- "Performance in a Processor-Constrained Environment" on page 177 illustrates, for the CMS-intensive workload, how increasing levels of processor overcommitment are reflected in the various VM/ESA performance metrics.
- "VM/ESA 1.5 (370 Feature)" on page 182 provides performance information on migrating from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 1.5 (370 Feature) and shows a performance comparison of VMSES to VMSES/E.
- "Effect of IABIAS on Response Times" on page 192 illustrates the potential response time benefits that can be derived through the use of the SET SRM IABIAS command to increase the preference given to short requests.
- "VM/VTAM 3.4.1" on page 196 provides a performance comparison between VM/VTAM 3.4.0 and VM/VTAM 3.4.1.
- "3745 Comparison to CTCA" on page 199 shows how the differences between CTCA and 3745 connectivity affect the VM/ESA performance data for the CMS-intensive environment.
- "FS7F to FS8F Workload Comparison" on page 204 provides a bridge between FS7F and FS8F — the newest version of the CMS-intensive workload.

Processor Capacity: 9021-941

The measurements discussed in this section show the performance of VM/ESA Release 2.1 within the 9021-711 based family. The ITR(H) scaled well on the 9021-941 configured with 1, 2, and 3 processors online. However, it did not scale as well with the 4 processors online due to system storage constraints. Our 9021-941 had 512MB of real storage. This was not large enough to allow real storage to be scaled in proportion to the number of processors being used. Despite this limitation, these results are similar to earlier measurements made on the 9021-711 based processors with VM/ESA Release 1.1 using the LSPR workloads.[14]

The following graph represents the internal throughput rate (ITR) as a function of the number of processors online for the 9021-941. The dotted line, shown for reference, is a linear projection assuming no loss of ITR per processor when going from 1 to 4 processors online.

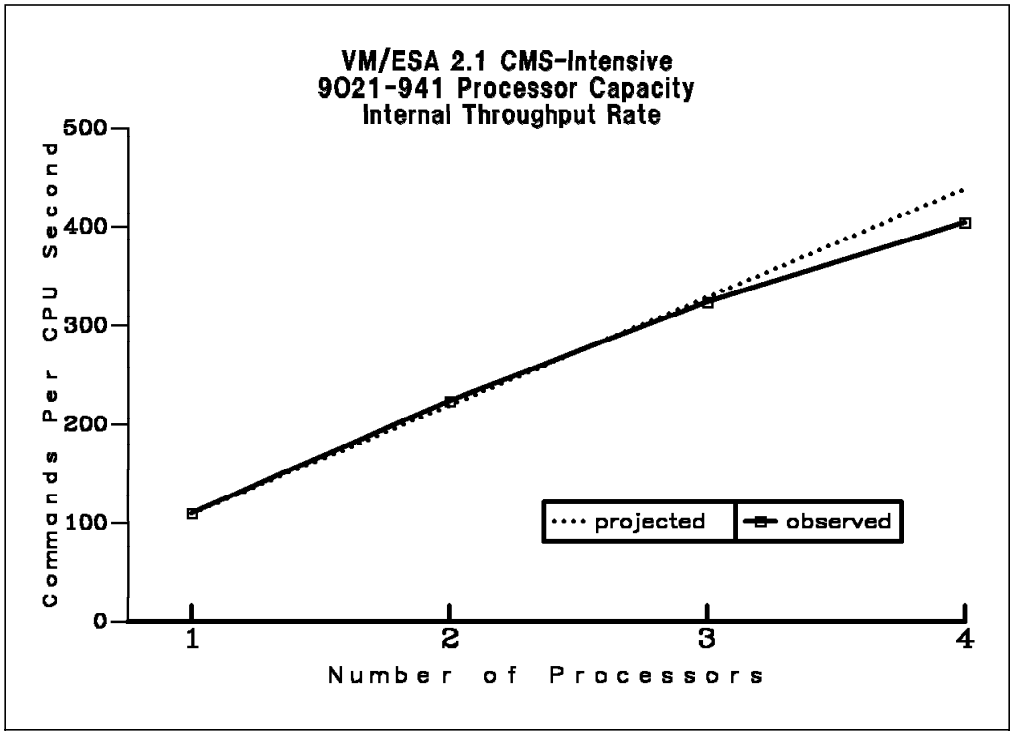

Figure 19. Internal throughput rate for the 9021-941 processor.

[14] The LSPR workloads for VM, PD4 and HT5, were developed for the purpose of comparing processors. Their key metric is ITR, with no constraints existing on storage or DASD. LSPR is an abbreviation for Large System Performance Reference.

Workload: FS7F0R

Hardware Configuration

Processor Model	*Processors Used*	*Real Storage*	*Expanded Storage*	*Retained for MDC*
9021-941	1	256MB	1024MB	400MB - 400MB
9021-941	2	512MB	2048MB	400MB - 400MB
9021-941	3	512MB	3072MB	400MB - 400MB
9021-941	4	512MB	4096MB	400MB - 400MB

Note: For each of the measurements, any expanded storage not used in the measurement was attached to an idle user.

Tape: 3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3380	3880-03	2	20					
3390-2	3990-2	4	18	6				
3390-2	3990-3	4	22	2		32W		2R

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 for SPOOL and minidisks were required or used.

Communications:

Control Unit	*Number*	*Lines per Control Unit*	*Speed*
3088-08	1	NA	4.5MB

Processor Capacity: 9021-941

Software Configuration

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

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
For the 9021-941 (1-Way):						
SMART	1	RTM	16MB/370	3%	700	QUICKDSP ON
VSCSn	2	VSCS	64MB/XA	10000	2355	QUICKDSP ON
VTAMXA	1	VTAM	64MB/XA	10000	600	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	2770	Users	3MB/XC	100		
For the 9021-941 (2-Way):						
SMART	1	RTM	16MB/370	3%	700	QUICKDSP ON
VSCSn	4	VSCS	64MB/XA	10000	2410	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	600	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	5670	Users	3MB/XC	100		
For the 9021-941 (3-Way):						
SMART	1	RTM	16MB/370	3%	700	QUICKDSP ON
VSCSn	6	VSCS	64MB/XA	10000	2340	QUICKDSP ON
VTAMXA	1	VTAM	64MB/XA	10000	600	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	8260	Users	3MB/XC	100		
For the 9021-941 (4-Way):						
SMART	1	RTM	16MB/370	3%	700	QUICKDSP ON
VSCSn	8	VSCS	64MB/XA	10000	2250	QUICKDSP ON
VTAMXA	1	VTAM	64MB/XA	10000	600	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	10300	Users	3MB/XC	100		

Measurement Discussion: The processor capacity measurements were made using a 9021-941 processor. A processor utilization of 90% was chosen as the target for each measurement. The VARY OFFLINE PROCESSORS command was used to disable excess processors. Virtual disks in storage were defined in the workload (COB417F script) instead of temporary disks.

The 1-way was configured with 256MB of real storage and the 2, 3, and 4-way measurements were configured with 512MB of real storage. The minidisk cache size was fixed at 400MB and not scaled when the number of processors increased. For the 1-way measurement, there was no internal VSCS machine.

As the number of processors and users increased for the measurements, the system became more real storage constrained. This is reflected by the fact that PGBLPBS/USER decreased. The system was most storage constrained when comparing the 3-way to the 4-way.

Response times remained low, though the system became more storage constrained, because most of the paging (PAGE/CMD + XSTOR/CMD) was done to expanded storage and not DASD. PAGE/CMD for the 1-way was higher

because the number of expanded storage blocks per user that are available for paging (XSTOR BLKS/USER) was on the low side. The 1-way, as with all the measurements, had 400 MB of expanded storage for minidisk cache. 400MB for minidisk cache in the 1-way measurement was a bigger percentage of the total amount of expanded storage than in the 2,3, or 4-way, which left less for paging. Response time for the 4-way measurement decreased slightly from the 3-way. Nothing specific was identified that contributed to this. The difference is within measurement variability.

While each measurement had a single VTAM, the number of external VSCSs active as VTAM Service Machines (VSMs) varied. Each VSCS machine had reserved pages equal to the number of users multiplied by 1.7[15] divided by the number of VSCS machines. The purpose of the reserved pages was to minimize paging in the server virtual machines.

The 2-way measurement looks especially good; it shows an ITRR(H) of 2.04. This is not unusual for VM/ESA because going from a 1-way to a 2-way system greatly reduces the impact of workload-independent processing (such as CP monitor and RTM activity) and allows for much more VTAM I/O coattailing.

The decrease in DIAG/CMD is mostly accounted for by the decreases in DIAG 04/CMD and DIAG 98/CMD. The decrease in DIAG 04/CMD is because RTM sampling is prorated across a larger number of commands. The decrease in DIAG 98/CMD reflects the more efficient VTAM I/O coattailing as total system communications traffic increases.

The following table summarizes the four processor capacity measurements. The results show that VM/ESA Release 2.1 scaled as expected on the 9021-941 as the number of processors was increased from 1 to 4. The ITRs ranged from 109.5 on the 1-way to 403.8 for the 4-way with ITRRs of 1, 2.04, 2.95, and 3.69. The table shows that these results are similar to studies of VM/ESA Release 1.1 on the 9021 711-based family using the LSPR workloads.

Table 52 (Page 1 of 2). ITR and ITR ratios for the FS7F workload on VM/ESA Release 2.1 compared to the PD4 and HT5 LSPR workloads on VM/ESA Release 1.1 on the 9021-941.

Variable	Workload	1-way	2-way	3-way	4-way
ITR(H)	FS7F	109.53	223.00	323.55	403.83
	PD4	70.71	141.14	207.82	268.96
	HT5	159.99	314.56	465.30	603.77
ITRR(H)	FS7F	1.000	2.036	2.954	3.687
	PD4	1.000	1.996	2.939	3.803
	HT5	1.000	1.966	2.908	3.773

[15] From a previous measurement, we calculated the VSCS machines to need about 1.7 pages per user.

Table 52 (Page 2 of 2). ITR and ITR ratios for the FS7F workload on VM/ESA Release 2.1 compared to the PD4 and HT5 LSPR workloads on VM/ESA Release 1.1 on the 9021-941.

Variable	Workload	1-way	2-way	3-way	4-way
ITRR(H)	FS7F	1.000	1.018	0.985	0.921
per	PD4	1.000	0.998	0.980	0.950
processor	HT5	1.000	0.983	0.969	0.943

The FS7F workload did better than the LSPR workloads when going from the 1-way to the 2-way because of the more efficient VTAM I/O coattailing. This benefit does not apply to the LSPR workloads because these workloads make use of an internal driver tool to simulate users instead of going through VTAM.

The FS7F workload experienced larger decreases in ITRR(H) per processor than the LSPR workloads when going from the 2-way to the 3-way and from the 3-way to the 4-way because the 3-way and 4-way environments were more storage constrained. For the LSPR workloads going from the 1-way to the 4-way, the real/expanded storage amounts were 1280MB/512MB, 1536MB/1GB, 1792MB/1GB, and 2GB/2GB. For the FS7F workload going from the 1-way to the 4-way, the real/expanded storage amounts were 256MB/1024MB, 512MB/2048MB, 512MB/3072MB, and 512MB/4096MB.

Table 53 (Page 1 of 2). VM/ESA Release 2.1 processor capacity measurements on the 9021-941.

Release	**VM/ESA 2.1**	**VM/ESA 2.1**	**VM/ESA 2.1**	**VM/ESA 2.1**
Run ID	**516E2770**	**526E5670**	**536E8260**	**546EA300**
Environment				
Real Storage	**256MB**	**512MB**	**512MB**	**512MB**
Exp. Storage	**1024MB**	**2048MB**	**3072MB**	**4096MB**
Users	**2770**	**5670**	**8260**	**10300**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**2**	**4**	**6**	**8**
Processors	**1**	**2**	**3**	**4**
Response Time				
TRIV INT	0.038	0.047	0.070	0.079
NONTRIV INT	0.217	0.164	0.194	0.205
TOT INT	0.161	0.126	0.147	0.159
TOT INT ADJ	0.141	0.112	0.141	0.151
AVG FIRST (T)	0.130	0.161	0.235	0.222
AVG LAST (T)	0.192	0.199	0.279	0.270
Throughput				
AVG THINK (T)	26.32	26.29	26.24	26.43
ETR	85.67	178.95	279.95	345.08
ETR (T)	98.13	200.46	291.10	362.86
ETR RATIO	0.873	0.893	0.962	0.951
ITR (H)	109.53	223.00	323.55	403.83
ITR	95.76	99.63	103.94	96.22
EMUL ITR	142.21	147.81	158.77	154.48
ITRR (H)	1.000	2.036	2.954	3.687
ITRR	1.000	1.040	1.085	1.005
Proc. Usage				
PBT/CMD (H)	9.130	8.968	9.272	9.905
PBT/CMD	9.172	8.979	9.275	9.894
CP/CMD (H)	3.185	3.123	3.400	3.964
CP/CMD	3.057	2.943	3.195	3.720
EMUL/CMD (H)	5.946	5.846	5.872	5.941
EMUL/CMD	6.114	6.036	6.080	6.173
Processor Util.				
TOTAL (H)	89.59	179.78	269.91	359.42
TOTAL	90.00	180.00	270.00	359.00
UTIL/PROC (H)	na	89.89	89.97	89.86
UTIL/PROC	na	90.00	90.00	89.75
TOTAL EMUL (H)	58.34	117.18	170.93	215.59
TOTAL EMUL	60.00	121.00	177.00	224.00
MASTER TOTAL (H)	na	90.18	91.25	92.31
MASTER TOTAL	na	90.00	91.00	92.00
MASTER EMUL (H)	na	52.60	45.69	38.20
MASTER EMUL	na	55.00	48.00	40.00
TVR(H)	1.54	1.53	1.58	1.67
TVR	1.50	1.49	1.53	1.60
Storage				
NUCLEUS SIZE (V)	2376KB	2376KB	2376KB	2376KB
TRACE TABLE (V)	200KB	400KB	600KB	800KB
WKSET (V)	79	71	67	55
PGBLPGS	47323	100K	88810	79930
PGBLPGS/USER	17.1	17.6	10.8	7.8
FREEPGS	7683	15578	22519	28029
FREE UTIL	0.90	0.90	0.90	0.90
SHRPGS	1298	1667	1926	2124

Processor Capacity: 9021-941

Table 53 (Page 2 of 2). VM/ESA Release 2.1 processor capacity measurements on the 9021-941.

Release	**VM/ESA 2.1**	**VM/ESA 2.1**	**VM/ESA 2.1**	**VM/ESA 2.1**
Run ID	**516E2770**	**526E5670**	**536E8260**	**546EA300**
Environment				
Real Storage	**256MB**	**512MB**	**512MB**	**512MB**
Exp. Storage	**1024MB**	**2048MB**	**3072MB**	**4096MB**
Users	**2770**	**5670**	**8260**	**10300**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**2**	**4**	**6**	**8**
Processors	**1**	**2**	**3**	**4**
Paging				
READS/SEC	488	392	575	703
WRITES/SEC	418	358	520	643
PAGE/CMD	9.233	3.741	3.762	3.709
PAGE IO RATE (V)	152.300	119.100	179.600	193.100
PAGE IO/CMD (V)	1.552	0.594	0.617	0.532
XSTOR IN/SEC	267	964	1613	3697
XSTOR OUT/SEC	713	1453	2407	4816
XSTOR/CMD	9.987	12.057	13.810	23.461
FAST CLR/CMD	6.991	7.069	7.042	7.036
Queues				
DISPATCH LIST	35.95	62.66	111.50	142.62
ELIGIBLE LIST	0.00	0.00	0.00	0.02
I/O				
VIO RATE	896	1760	2485	3054
VIO/CMD	9.131	8.780	8.537	8.416
RIO RATE (V)	382	512	703	822
RIO/CMD (V)	3.893	2.554	2.415	2.265
NONPAGE RIO/CMD (V)	2.341	1.960	1.798	1.733
DASD RESP TIME (V)	16.100	15.300	15.800	16.900
MDC READS (blks)	558	1136	1619	2006
MDC WRITES (blks)	201	416	632	791
MDC MODS	164	336	487	607
MDC AVOID	279	567	803	994
MDC HIT RATIO	0.93	0.93	0.91	0.91
PRIVOPs				
PRIVOP/CMD	20.266	18.179	16.677	19.932
DIAG/CMD	26.530	24.966	24.382	24.107
DIAG 04/CMD	2.425	1.241	0.886	0.737
DIAG 08/CMD	0.761	0.762	0.761	0.759
DIAG 0C/CMD	1.344	1.343	1.343	1.341
DIAG 14/CMD	0.025	0.025	0.025	0.025
DIAG 58/CMD	1.250	1.250	1.250	1.247
DIAG 98/CMD	0.910	0.530	0.295	0.214
DIAG A4/CMD	3.870	3.852	3.859	3.843
DIAG A8/CMD	1.781	1.797	1.793	1.776
DIAG 214/CMD	12.965	12.970	12.973	12.968
SIE/CMD	54.765	49.886	48.093	55.117
SIE INTCPT/CMD	39.979	34.421	31.742	34.724
FREE TOTL/CMD	72.345	69.840	65.269	63.385
VTAM Machines				
WKSET (V)	5010	10260	14250	18608
TOT CPU/CMD (V)	1.6362	1.4957	1.4924	1.6367
CP CPU/CMD (V)	0.7756	0.7119	0.7290	0.8314
VIRT CPU/CMD (V)	0.8606	0.7839	0.7634	0.8053
DIAG 98/CMD (V)	0.910	0.530	0.295	0.214

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Performance in a Processor-Constrained Environment

The purpose of this study is to illustrate the kinds of performance effects that can occur on VM systems that are highly processor-constrained.

Four VM/ESA Release 2.1 measurements were obtained using the FS7F CMS-intensive workload on a 9121-480 processor. The base measurement was run with 1860 users and resulted in a processor utilization of about 90%. The remaining three measurements were run at 1950, 2070, and 2300 users in order span the range from acceptable performance to severe processor overcommitment.

Workload: FS7F0R

Hardware Configuration

Processor model: 9121-480
Processors used: 2
Storage
  Real: 224MB
  Expanded: 32MB (all for MDC)
Tape: 3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-3	2	6	5	5			2 R
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	4		2	2	16 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*
3088-08	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
1860 user measurement:						
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	512	QUICKDSP ON
Unnnn	1860	Users	3MB/XC	100		
1950 user measurement:						
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	512	QUICKDSP ON
Unnnn	1950	Users	3MB/XC	100		
2070 user measurement:						
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	700	QUICKDSP ON
Unnnn	2070	Users	3MB/XC	100		
2300 user measurement:						
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	1100	QUICKDSP ON
Unnnn	2300	Users	3MB/XC	100		

Measurement Discussion Because of the large increase in users, average processor utilization (UTIL/PROC(H)) increased from 90.4% to 100%. The increased processor contention caused external response time (AVG LAST(T)) to increase from 0.38 seconds to 3.96 seconds.***:*** The TPNS-based measures of response time and throughput (those with the (T) or (H) designation) provide an accurate picture of what is actually happening in the system. This is because TPNS′s notion as to what constitutes a transaction (command) remains accurate and consistent, even at very high levels of processor contention.

By contrast, those measures of response time and throughput that are based on the CP scheduler′s notion of what constitutes a transaction can become misleading at very high processor utilizations. The CP scheduler has no way of definitively knowing when one transaction stops and the next one starts. As an approximation, CP assumes that when a given virtual machine has been idle for more than 300 milliseconds, the current transaction is over and any renewed activity represents a new transaction.

This approximation is usually reasonably accurate. For example, looking at the 90% measurement (L26E186C), we see that the scheduler believes that the transaction rate is 57.60 (ETR), whereas it is really 65.21 (ETR(T)) — a ratio of 0.88 (ETR RATIO). Note, however, that as we go to successively higher processor utilizations, the ETR RATIO gradually increases and finally takes a big jump to 2.32 in the 100% measurement (L26E2300). At very high utilizations, the scheduler′s method of distinguishing transactions often gets fooled into thinking that there are more transactions than there really are. This happens because delays (for services such as I/O and VTAM communications) that are normally

short-term, sometimes become elongated by processor contention to the point where they exceed the 300 millisecond criterion that the scheduler uses.

The net result is that, at very high utilizations, throughput metrics that are based on CP′s notion of a transaction are biased high, while response time metrics that are based on CP′s notion of a transaction are biased low. In the following measurement results table, examples include TRIV INT, NONTRIV INT, TOT INT, ETR, and ITR. TOT INT ADJ does not have this problem because this measure of response time, although based on CP data, has been normalized to the transaction rate as measured by TPNS.

Processor usage per command (PBT/CMD(H) held steady throughout this series of measurements, indicating that processing efficiency was not significantly affected by the increased processor contention.

The increased number of users resulted in increased real storage usage. This shows up as increases in the number of fixed pages used by CP (FREEPGS), shared pages (SHRPGS), and VTAM working set (WKSET(V) - VTAM Machines). This increased real storage usage did exert some upward pressure on paging (PAGE IO/CMD(V)), but the increase was small because the measured environment was not storage constrained. The increased paging was only a minor contributor to the increased response times.

VTAM working set was influenced by the number of pages that were reserved for the VTAM virtual machine. This was increased as the number of users increased so as to keep serial page faulting in the VTAM virtual machine to a minimum.

The decrease in non-paging I/Os (NONPAGE RIO/CMD(V)) was due to the large decrease in diagnose 98 instructions (DIAG 98/CMD(V)) issued by VTAM to do communications I/O. This occurred because the longer response times meant that more time was available to accumulate multiple terminal responses into a single communications I/O, a performance optimization known as VTAM coattailing.

VM/ESA Release 2.1 was used for these measurements. The 1860 and 2300 user measurements were repeated using VM/ESA Release 2.2 and the results were very similar. VM/ESA Release 2 and earlier VM/ESA releases should also show similar results.

A Processor-Constrained Environment

Table 54 (Page 1 of 2). Processor-constrained measurements on the 9121-480

Release	**VM/ESA 2.1**	**VM/ESA 2.1**	**VM/ESA 2.1**	**VM/ESA 2.1**
Run ID	**L26E186C**	**L26E1950**	**L26E2071**	**L26E2300**
Environment				
Real Storage	**224MB**	**224MB**	**224MB**	**224MB**
Exp. Storage	**32MB**	**32MB**	**32MB**	**32MB**
Users	**1860**	**1950**	**2070**	**2300**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**0**	**0**	**0**	**0**
Processors	**2**	**2**	**2**	**2**
Response Time				
TRIV INT	0.124	0.127	0.132	0.014
NONTRIV INT	0.458	0.536	0.799	0.802
TOT INT	0.347	0.397	0.555	0.269
TOT INT ADJ	0.307	0.358	0.540	0.624
AVG FIRST (T)	0.250	0.296	0.434	3.092
AVG LAST (T)	0.380	0.469	0.749	3.961
Throughput				
AVG THINK (T)	26.37	26.30	26.39	26.30
ETR	57.60	61.49	69.91	167.85
ETR (T)	65.21	68.21	71.83	72.30
ETR RATIO	0.883	0.902	0.973	2.321
ITR (H)	72.17	72.44	72.62	72.31
ITR	31.89	32.67	35.34	83.97
EMUL ITR	46.83	47.92	51.85	125.75
ITRR (H)	1.000	1.004	1.006	1.002
ITRR	1.000	1.025	1.108	2.633
Proc. Usage				
PBT/CMD (H)	27.711	27.607	27.542	27.660
PBT/CMD	27.756	27.563	27.566	27.661
CP/CMD (H)	9.417	9.355	9.334	9.847
CP/CMD	8.894	8.797	8.771	9.128
EMUL/CMD (H)	18.295	18.253	18.208	17.813
EMUL/CMD	18.862	18.766	18.795	18.533
Processor Util.				
TOTAL (H)	180.71	188.30	197.83	200.00
TOTAL	181.00	188.00	198.00	200.00
UTIL/PROC (H)	90.35	94.15	98.92	100.00
UTIL/PROC	90.50	94.00	99.00	100.00
TOTAL EMUL (H)	119.30	124.50	130.79	128.80
TOTAL EMUL	123.00	128.00	135.00	134.00
MASTER TOTAL (H)	90.11	94.08	98.88	100.00
MASTER TOTAL	90.00	94.00	99.00	100.00
MASTER EMUL (H)	53.59	56.21	59.41	58.93
MASTER EMUL	55.00	58.00	61.00	61.00
TVR(H)	1.51	1.51	1.51	1.55
TVR	1.47	1.47	1.47	1.49
Storage				
NUCLEUS SIZE (V)	2368KB	2368KB	2368KB	2368KB
TRACE TABLE (V)	400KB	400KB	400KB	400KB
WKSET (V)	83	82	82	78
PGBLPGS	47708	47323	46834	45865
PGBLPGS/USER	25.6	24.3	22.6	19.9
FREEPGS	5000	5253	5559	6166
FREE UTIL	0.97	0.93	0.97	0.95
SHRPGS	1181	1196	1303	1425

Table 54 (Page 2 of 2). Processor-constrained measurements on the 9121-480

Release	**VM/ESA 2.1**	**VM/ESA 2.1**	**VM/ESA 2.1**	**VM/ESA 2.1**
Run ID	**L26E186C**	**L26E1950**	**L26E2071**	**L26E2300**
Environment				
Real Storage	**224MB**	**224MB**	**224MB**	**224MB**
Exp. Storage	**32MB**	**32MB**	**32MB**	**32MB**
Users	**1860**	**1950**	**2070**	**2300**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**0**	**0**	**0**	**0**
Processors	**2**	**2**	**2**	**2**
Paging				
READS/SEC	629	639	630	654
WRITES/SEC	426	447	472	500
PAGE/CMD	16.178	15.922	15.342	15.960
PAGE IO RATE (V)	167.900	177.300	189.000	207.300
PAGE IO/CMD (V)	2.575	2.599	2.631	2.867
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	6.517	6.554	6.683	6.805
Queues				
DISPATCH LIST	37.07	43.34	58.42	79.53
ELIGIBLE LIST	0.00	0.00	0.00	0.00
I/O				
VIO RATE	626	642	665	636
VIO/CMD	9.600	9.412	9.258	8.796
RIO RATE (V)	387	391	397	382
RIO/CMD (V)	5.935	5.733	5.527	5.283
NONPAGE RIO/CMD (V)	3.360	3.133	2.896	2.416
DASD RESP TIME (V)	20.100	20.200	20.400	20.700
MDC READS (blks)	361	378	398	400
MDC WRITES (blks)	178	186	197	198
MDC MODS	145	151	160	161
MDC AVOID	180	188	198	199
MDC HIT RATIO	0.91	0.91	0.91	0.91
PRIVOPs				
PRIVOP/CMD	14.104	14.187	14.169	13.782
DIAG/CMD	27.019	26.800	26.494	26.090
DIAG 04/CMD	2.458	2.431	2.309	2.294
DIAG 08/CMD	0.762	0.758	0.761	0.761
DIAG 0C/CMD	1.343	1.340	1.340	1.340
DIAG 14/CMD	0.025	0.025	0.025	0.025
DIAG 58/CMD	1.250	1.251	1.251	1.248
DIAG 98/CMD	1.226	1.046	0.880	0.402
DIAG A4/CMD	3.916	3.912	3.915	3.915
DIAG A8/CMD	1.880	1.885	1.890	1.904
DIAG 214/CMD	12.949	12.943	12.917	12.985
SIE/CMD	51.464	50.757	49.980	47.992
SIE INTCPT/CMD	34.481	33.500	32.487	29.755
FREE TOTL/CMD	67.933	67.368	66.923	65.460
VTAM Machines				
WKSET (V)	526	518	700	1100
TOT CPU/CMD (V)	4.0041	3.8933	3.7899	3.6958
CP CPU/CMD (V)	1.5761	1.5394	1.5082	1.7058
VIRT CPU/CMD (V)	2.4280	2.3539	2.2817	1.9901
DIAG 98/CMD (V)	1.223	1.047	0.880	0.402

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

VM/ESA 1.5 (370 Feature)

This section summarizes the results of performance tests of VM/ESA Release 1.5 (370 Feature) running on a 9221-120 processor. It provides performance information based on results of tests conducted in a CMS-intensive environment as well as a performance comparison of VMSES/E to VMSES.

The CMS-intensive environment was able to support 100 users with a CMS-intensive workload and achieve external response time of less than 2 seconds at a processor utilization of about 90%. Measurements were made comparing to both the General Availability (GA) level and the current service level (Put 9305) of VM/ESA Release 1.0 (370 Feature).

The VMSES/E comparison showed that overall elapsed time required to perform service functions improved by 26% when compared to VMSES in VM/ESA Release 1.0 (370 Feature).

Migration from VM/ESA 1.0 (370 Feature)

Workload: FS7F0R

Hardware Configuration

Processor model: 9221-120
Processors used: 1
Storage
  Real: 16MB
  Expanded: 0MB
Tape: 3480 (Monitor)

DASD:

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

Communications:

Control Unit	*Number*	*Lines per Control Unit*	*Speed*
3088	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*	*PRIORITY*	*RESERVED*	*Other Options*
VTAM	1	VTAM/VSCS	8MB/370	1	OFF	QDROP OFF FAVOR
SMART	1	RTM	3MB/370	1	OFF	QDROP OFF FAVOR
Unnnn	100	Users	3MB/370	64	OFF	

Measurement Discussion: The three differences that an installation will experience when migrating from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 1.5 (370 Feature) are the new VMSES/E process, CMS Pipelines, and the difference in service between the two systems. VM/ESA Release 1.5 (370 Feature) has included in it the latest service from VM/ESA Release 1.0 (370 Feature). None of these changes has a significant impact on overall system performance; therefore you can expect to have similar system performance when migrating from VM/ESA Release 1.0 (370 Feature) to VM/ESA Release 1.5 (370 Feature).

The following tables show that VM/ESA Release 1.5 (370 Feature) compared to VMESA Release 1.0 (370 Feature) remained equivalent in external response time

(AVG LAST (T)) and internal throughput rate (ITR(H)) for the GA levels. The Release 1.0 Put level 9305 did show an increase in virtual processor busy time. However, an improvement was made to the GA level of Release 1.5 to compensate for this factor.

The measurement results are summarized in the following tables. In the tables, those metrics identified with an asterisk (*) are identical in meaning for both VM/ESA (370 Feature) and VM/ESA (ESA Feature) and are defined in the Glossary. For the derivation of the remaining metrics, refer to the Glossary of Performance Terms contained in *VM/ESA Release 2 Performance Report* .

The following table compares VM/ESA Release 1.5 (370 Feature) to VM/ESA Release 1.0 (370 Feature) at the GA level. Average response time grew by 2.1% and the internal throughput rate decreased by 1.5%.
align='L' concat=yes width=column.

Table 55 (Page 1 of 2). CMS-intensive migration from the GA level of VM/ESA Release 1.0 (370 Feature) on the 9221-120

Release Level Run ID	VM/ESA 1.0 GA H17E0106	VM/ESA 1.5 GA H17E0113	Difference	%Difference
Environment				
REAL STORAGE	**16MB**	**16MB**		
EXP. STORAGE	**0MB**	**0MB**		
USERS	**100**	**100**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
PROCESSORS	**1**	**1**		
Response Time				
TRIV INT	0.290	0.330	0.040	13.79%
NONTRIV INT	4.200	4.209	0.009	0.20%
TOT INT	1.110	1.100	-0.010	-0.90%
TOT INT ADJ	1.025	1.019	-0.006	-0.62%
AVG FIRST (T) *	1.031	1.104	0.073	7.08%
AVG LAST (T) *	1.706	1.741	0.035	2.05%
Throughput				
AVG THINK (T) *	28.21	28.26	0.05	0.18%
ETR	3.11	3.11	-0.01	-0.26%
ETR (T) *	3.37	3.35	-0.02	-0.54%
ETR RATIO	0.923	0.926	0.003	0.28%
ITR (H)*	3.90	3.84	-0.06	-1.47%
ITR	3.59	3.57	-0.03	-0.70%
EMUL ITR	5.96	5.98	0.02	0.35%
ITRR (H) *	1.000	0.985	-0.015	-1.47%
ITRR	1.000	0.993	-0.007	-0.70%
Proc. Usage				
PBT/CMD (H) *	256.298	260.122	3.825	1.49%
PBT/CMD	257.129	259.659	2.529	0.98%
CP/CMD (H) *	114.814	118.186	3.373	2.94%
CP/CMD	102.069	104.710	2.641	2.59%
EMUL/CMD (H) *	141.484	141.936	0.452	0.32%
EMUL/CMD	155.061	154.948	-0.112	-0.07%

Table 55 (Page 2 of 2). CMS-intensive migration from the GA level of VM/ESA Release 1.0 (370 Feature) on the 9221-120

Release Level Run ID	VM/ESA 1.0 GA H17E0106	VM/ESA 1.5 GA H17E0113	Difference	%Difference
Environment				
REAL STORAGE	**16MB**	**16MB**		
EXP. STORAGE	**0MB**	**0MB**		
USERS	**100**	**100**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
PROCESSORS	**1**	**1**		
Processor Util.				
TOTAL (H) *	86.43	87.25	0.82	0.94%
TOTAL	86.71	87.09	0.38	0.44%
TOTAL EMUL (H) *	47.71	47.61	-0.11	-0.22%
TOTAL EMUL	52.29	51.97	-0.32	-0.61%
TVR(H) *	1.81	1.83	0.02	1.17%
TVR	1.66	1.68	0.02	1.20%
Storage				
NUCLEUS SIZE	508KB	515KB	7KB	1.38%
TRACE TABLE	256KB	256KB	0KB	0.00%
WKSET	73	74	2	2.17%
PGBLPGS	3405	3403	-2	-0.06%
PGBLPGS/USER	34.1	34.0	0.0	-0.06%
FREEPGS	500	500	0	0.05%
SHRPGS	556	532	-24	-4.24%
Paging				
READS/SEC	21	21	1	2.76%
WRITES/SEC	13	14	1	4.58%
PAGE/CMD	9.999	10.402	0.403	4.03%
PAGE IO RATE	32.508	33.655	1.147	3.53%
PAGE IO/CMD	9.640	10.034	0.394	4.09%
Queues				
DISPATCH LIST	3.88	3.70	-0.18	-4.64%
ELIGIBLE LIST	0.02	0.02	0.00	0.00%
I/O				
VIO RATE	57	56	0	-0.27%
VIO/CMD	16.775	16.822	0.046	0.28%
RIO RATE	61	63	2	2.72%
RIO/CMD	18.047	18.638	0.592	3.28%
MDSK/CMD	5.634	5.665	0.031	0.54%
PRIVOPs				
PRIVOP/CMD (R)	10.826	10.742	-0.084	-0.78%
DIAG/CMD (R)	18.397	18.582	0.185	1.01%
VTAM Machines				
WKSET	463.5	464	0.5	0.11%
TOT CPU/CMD	45.7989	45.9647	0.1658	0.36%
CP CPU/CMD	18.5337	18.6344	0.1007	0.54%
VIRT CPU/CMD	27.2652	27.3304	0.0652	0.24%
DIAG 98/CMD	3.361	3.338	-0.023	-0.69%

Note: T=TPNS, H=Hardware Monitor, R=RTM, Unmarked=VMMAP

VM/ESA 1.5 (370 Feature)

The following table compares VM/ESA Release 1.5 (370 Feature) to VM/ESA Release 1.0 (370 Feature) at the current service level. Average response time improved by 1.4% and the internal throughput rate increased by 3.2%.

Table 56 (Page 1 of 2). CMS-intensive migration from the PUT 9305 level of VM/ESA Release 1.0 (370 Feature) on the 9221-120

Release Level Run ID	VM/ESA 1.0 Put 9305 H17E0108	VM/ESA 1.5 GA H17E0113	Difference	%Difference
Environment				
REAL STORAGE	**16MB**	**16MB**		
EXP. STORAGE	**0MB**	**0MB**		
USERS	**100**	**100**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
PROCESSORS	**1**	**1**		
Response Time				
TRIV INT	0.320	0.330	0.010	3.13%
NONTRIV INT	4.469	4.209	-0.260	-5.82%
TOT INT	1.150	1.100	-0.050	-4.35%
TOT INT ADJ	1.069	1.019	-0.050	-4.67%
AVG FIRST (T) *	1.076	1.104	0.028	2.60%
AVG LAST (T) *	1.766	1.741	-0.025	-1.42%
Throughput				
AVG THINK (T) *	28.31	28.26	-0.05	-0.18%
ETR	3.10	3.11	0.00	0.08%
ETR (T) *	3.34	3.35	0.01	0.41%
ETR RATIO	0.929	0.926	-0.003	-0.34%
ITR (H) *	3.72	3.84	0.12	3.24%
ITR	3.46	3.57	0.10	2.96%
EMUL ITR	5.70	5.98	0.27	4.75%
ITRR (H) *	1.000	1.032	0.032	3.24%
ITRR	1.000	1.030	0.030	2.96%
Proc. Usage				
PBT/CMD (H) *	268.559	260.122	-8.437	-3.14%
PBT/CMD	268.250	259.659	-8.591	-3.20%
CP/CMD (H) *	119.181	118.186	-0.995	-0.83%
CP/CMD	105.384	104.710	-0.674	-0.64%
EMUL/CMD (H) *	149.378	141.936	-7.442	-4.98%
EMUL/CMD	162.866	154.948	-7.917	-4.86%
Processor Util.				
TOTAL (H) *	89.70	87.25	-2.46	-2.74%
TOTAL	89.60	87.09	-2.51	-2.80%
TOTAL EMUL (H) *	49.89	47.61	-2.29	-4.59%
TOTAL EMUL	54.40	51.97	-2.43	-4.47%
TVR(H)	1.80	1.83	0.03	1.94%
TVR	1.65	1.68	0.03	1.82%
Storage				
NUCLEUS SIZE	515KB	515KB	0KB	0.00%
TRACE TABLE	256KB	256KB	0KB	0.00%
WKSET	74	74	0	-0.05%
PGBLPGS	3403	3403	0	0.00%
PGBLPGS/USER	34.0	34.0	0.0	0.00%
FREEPGS	500	500	0	-0.01%
SHRPGS	564	532	-32	-5.66%

Table 56 (Page 2 of 2). CMS-intensive migration from the PUT 9305 level of VM/ESA Release 1.0 (370 Feature) on the 9221-120

Release Level Run ID	VM/ESA 1.0 Put 9305 H17E0108	VM/ESA 1.5 GA H17E0113	Difference	%Difference
Environment				
REAL STORAGE	**16MB**	**16MB**		
EXP. STORAGE	**0MB**	**0MB**		
USERS	**100**	**100**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
PROCESSORS	**1**	**1**		
Paging				
READS/SEC	21	21	0	-0.33%
WRITES/SEC	14	14	0	0.00%
PAGE/CMD	10.467	10.402	-0.064	-0.61%
PAGE IO RATE	33.781	33.655	-0.126	-0.37%
PAGE IO/CMD	10.114	10.034	-0.079	-0.79%
Queues				
DISPATCH LIST	3.92	3.70	-0.22	-5.61%
ELIGIBLE LIST	0.02	0.02	0.00	0.00%
I/O				
VIO RATE	56	56	0	0.04%
VIO/CMD	16.885	16.822	-0.064	-0.38%
RIO RATE	63	63	0	-0.35%
RIO/CMD	18.781	18.638	-0.142	-0.76%
MDSK/CMD	5.688	5.665	-0.023	-0.41%
PRIVOPs				
PRIVOP/CMD (R)	11.111	10.742	-0.369	-3.32%
DIAG/CMD (R)	18.680	18.582	-0.098	-0.52%
VTAM Machines				
WKSET	470	464	-6	-1.28%
TOT CPU/CMD	46.4048	45.9647	-0.4401	-0.95%
CP CPU/CMD	18.7116	18.6344	-0.0772	-0.41%
VIRT CPU/CMD	27.6932	27.3304	-0.3628	-1.31%
DIAG 98/CMD	3.393	3.338	-0.055	-1.63%

Note: T=TPNS, H=Hardware Monitor, R=RTM, Unmarked=VMMAP

VMSES/E Performance

A number of requirements were addressed in VM/ESA Release 1.5 (370 Feature) Refresh for Serviceability Enhancements:

- Reset the service stream because many of the service disks have reached capacity.
- Replace the VMSES service tool with VMSES/E since the VMSES/E service tool is more desirable in usability, capability, and quality. The VMSES/E service tool is available on the 370 Feature for use by LPs via the VM/ESA VMSES/E (370 Feature) for Licensed Programs. Note that the base VM/ESA system cannot use this support.
- Improve the preventive service process. The Recommended Service Upgrade (RSU) preventive service process is superior to the Program Update Tape (PUT) process but can be only used for VMSES/E supported products.
- Improve the install process.

It is expected that by replacing VMSES with VMSES/E, there will be both a disk space improvement and performance response time improvement. There will also be an improvement because a number of manual tasks have been automated. This saves on the amount of time reading a manual and executing the tasks desired. VMSES/E has automated many of these steps.

Three primary VMSES/E tools that help with the servicing of products were measured to quantify the effects of the new function, performance enhancements, and code quality improvements:

- 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.

On the 9221-120 configuration, the sum of the total elapsed time for the VMSES/E in VM/ESA Release 1.5 (370 Feature) receive, apply, and build commands improved 26% when compared to VMSES in VM/ESA Release 1.0 (370 Feature). Each individual command improved or degraded for the following reasons:

- After VMSES in VM/ESA Release 1.0 (370 Feature), the receive function was changed such that it performed more tape processing. The VMSES in VM/ESA Release 1.0 (370 Feature) receive function was limited in how it received service. The VMSES/E in VM/ESA Release 1.5 (370 Feature) receive function had processing/functional enhancements to allow for selective choices. Plus, the changes to the VMSES/E in VM/ESA Release 1.5 (370 Feature) receive processing helped the performance of the apply function by reducing the amount of processing the apply function had to perform.
- There is a big difference between VMSES in VM/ESA Release 1.0 (370 Feature) apply and VMSES/E in VM/ESA Release 1.5 (370 Feature) apply. VMFAPPLY experienced a major algorithm re-write, and the performance improvement of apply on VMSES/E in VM/ESA Release 1.5 (370 Feature) was

phenomenal -- the larger the service applied, the greater the benefit. This more than offset the increased costs due to receive or build.

- The build function after VMSES in VM/ESA Release 1.0 (370 Feature) has gone through a lot of change too. Some of areas of build were improved (for example, build processing was improved by the elimination of the validation associated with the CHKREQ tag). But the build function also has taken on more automation, and overall response time degraded somewhat. The increased automation of the build process reduced the complexity of the manual steps that build required before VMSES/E in VM/ESA Release 1.5 (370 Feature).

Also, VMFMRDSK and VMFBLD with the STATUS option were new after VMSES in VM/ESA Release 1.0 (370 Feature). VMFMRDSK was much faster than the old commands used for merging disks (note that the cost of manually investigating pre-VMFMRDSK function was not taken into account). Both of these new functions save manually researching what needs to be rebuilt.

The measurements described in the following sections are provided to demonstrate the performance of these changes on VMSES/E in VM/ESA Release 1.5 (370 Feature).

Hardware Configuration

Processor model:	9221-120
Processors used:	1
Storage	
Real:	16MB
Expanded:	0 MB
Tape:	3480 (one service tape for the receive command)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes -* *PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
9336-2			3	2	3	5		3

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*RELSHARE*	*RESERVED*	*Other Options*
MAINT	1	MAINT	16MB			

Measurement Discussion: All measurements were performed on a dedicated, first-level system with only one active user logged on (the MAINT user ID). The objective of these measurements was to show that the functional enhancements of VMSES/E in VM/ESA Release 1.5 (370 Feature) did not degrade performance when compared to VMSES in VM/ESA Release 1.0 (370 Feature) 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 for both VMSES and VMSES/E. The purpose of initializing SIM was to remove the one-time costs associated with setting up SIM.

Once SIM was initialized, a Corrective (COR) service tape containing CP service was loaded onto the system. The service workload used for these measurements was created so that even though the format of the service tapes

differed between VMSES/E in VM/ESA Release 1.5 (370 Feature) and VMSES in VM/ESA Release 1.0 (370 Feature), the same exact service was used for both VMSES in VM/ESA Release 1.0 (370 Feature) and VMSES/E in VM/ESA Release 1.5 (370 Feature), respectively. Hence, the receive, apply and build function of both releases worked on exactly the same service.

The CP service from the COR tape was received. VMFREC was used to receive a total of 1318 CP parts from three tape files. Next, the apply function (VMFAPPLY) was used to process 181 PTFs. Finally, 20 build lists were processed after running the VMFBLD command with the SERVICED option. Note that the build function (VMFBLD) with the STATUS option did not exist in VMSES 1.0. The VMFBLD with the STATUS option was run for VMSES/E in VM/ESA Release 1.5 (370 Feature) only and identified 124 build requirements.

The methodology described in this section applies to both VMSES/E in VM/ESA Release 1.5 (370 Feature) and VMSES in VM/ESA Release 1.0 (370 Feature). Performance data were collected before and after each command execution to determine total response time. The performance data were generated by the CP QUERY TIME 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 running 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 Scenario Time (seconds): the summation of the Total Time for all VMSES and VMSES/E commands that are defined below.

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, Apply, and Build

VMSES in VM/ESA Release 1.0 (370 Feature) Commands:

```
VMFREC   VM370 CP (COR
VMFAPPLY VM370 CP (COR
VMFMAC   DMKGPI DMKVM
COPY     DMKGPI MACLIB A = = O (OLDDATE REPLACE
ERASE    DMKGPI MACLIB A
VMFMAC   DMKPSI DMKVM
COPY     DMKPSI MACLIB A = = O (OLDDATE REPLACE
ERASE    DMKPSI MACLIB A
VMFMAC   DMKOM  DMKVM
COPY     DMKOM  MACLIB A = = O (OLDDATE REPLACE
ERASE    DMKOM  MACLIB A
VMFBLD   VM370 CP DMKBLINF
VMFBLD   VM370 CP CPMLOAD
VMFBLD   VM370 CP DMKBLDIR
VMFBLD   VM370 CP DMKDDRSA
VMFBLD   VM370 CP CPLOAD (NOIPL
VMFBLD   VM370 CP DMKBLLDR (NOIPL
```

VMSES/E in VM/ESA Release 1.5 (370 Feature) Commands:

```
VMFREC   PPF VM370 CP
VMFAPPLY PPF VM370 CP
VMFBLD   PPF VM370 CP
```

Scenario Details:

1318 parts received from 3 tape files.
180 PTFs applied after receiving parts from COR tape.
20 build lists processed; 124 objects built.

Table 57. Service measurement data: VMSES in VM/ESA Release 1.0 (370 Feature) and VMSES/E in VM/ESA Release 1.5 (370 Feature) on the 9221-120

Release	VMSES 1.0	VMSES/E 1.5	Difference	%Difference
Total Time (Q) for:				
Receive	1220	1523	303	25%
Apply	4912	2167	-2745	-56%
Build	1158	1698	540	47%
Total Scenario Time:	7290	5388	-1902	-26%

Note: Q=CP QUERY TIME

Effect of IABIAS on Response Times

Workload: FS8F0R

Hardware Configuration

Processor model: 9221-170
Processors used: 1
Storage
  Real: 64MB (default MDC)
  Expanded: 0MB
Tape: 3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*User*	*Server*	*System*
3390-2	3990-2	1	16	4	6			
3390-2	3990-3	1		2	2	8 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*
3088-08	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	350	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	300	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	300	Users	3MB/XC	100		

Measurement Discussion: Overall system responsiveness is improved if short transactions are given priority over long-running transactions. In VM/ESA, once a transaction is placed in the dispatch list, the interactive bias mechanism is the way in which the CP scheduler gives preference to short transactions.

The amount of interactive bias can be tuned by using the SET SRM IABIAS command. The amount of interactive bias is determined by two parameters: intensity and duration. Intensity is a number from 0 to 100 that determines where in the dispatch list a new transaction is inserted for its first minor time slice. The larger the value, the higher the initial priority. The duration setting is

the number of minor time slices that the interactive bias is in effect. A transaction receives the full intensity interactive bias only for its first minor time slice. The degree of bias fades out during subsequent minor time slices and disappears after "duration" minor time slices have been consumed by that transaction.

For more information on the SET SRM IABIAS command, see *VM/ESA CP Command and Utility Reference* and *VM/ESA Performance*.

Two VM/ESA Release 2.2 measurements were obtained using the FS8F CMS-interactive workload. The first measurement used the default IABIAS settings (intensity of 90, duration of 2 minor time slices), while the second measurement used stronger settings (intensity of 100, duration of 4). Relative to the default IABIAS settings, the "IABIAS 100 4" results showed a 27% improvement in trivial response time (TRIV INT) and a 13% improvement in average external response time (AVG LAST (T)). These results suggest that trivial response times and overall system performance may be improved by using the SET SRM IABIAS command to use a larger interactive bias than the system default.

These measurements were obtained on a 9221-170 processor running at about 85% processor utilization. The processor was the only constraining resource. Changes to the IABIAS settings would typically have smaller effects on faster processors, processors running at lower utilizations, and environments where response times are more determined by resources other than the processor.

Table 58 (Page 1 of 3). The effect of IABIAS on response times - 9221-170

IABIAS Settings **Release** **Run ID**	**90 2** **VM/ESA 2.2** **H17E0300**	**100 4** **VM/ESA 2.2** **H17E0302**	**Difference**	**%Difference**
Environment				
Real Storage	**64MB**	**64MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**300**	**300**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**1**	**1**		
Response Time				
TRIV INT	0.226	0.166	-0.060	-26.55%
NONTRIV INT	0.957	0.846	-0.111	-11.60%
TOT INT	0.735	0.622	-0.113	-15.37%
TOT INT ADJ	0.629	0.535	-0.094	-14.87%
AVG FIRST (T)	0.385	0.333	-0.052	-13.51%
AVG LAST (T)	0.670	0.586	-0.084	-12.54%
Throughput				
AVG THINK (T)	28.04	28.07	0.03	0.11%
ETR	8.93	8.98	0.05	0.56%
ETR (T)	10.44	10.43	0.00	-0.03%
ETR RATIO	0.856	0.861	0.005	0.59%
ITR (H)	11.82	11.95	0.13	1.13%
ITR	10.12	10.30	0.18	1.79%
EMUL ITR	15.16	15.38	0.22	1.47%
ITRR (H)	1.000	1.011	0.011	1.13%
ITRR	1.000	1.018	0.018	1.79%

Effect of IABIAS

Table 58 (Page 2 of 3). The effect of IABIAS on response times - 9221-170

IABIAS Settings **Release** **Run ID**	**90 2** **VM/ESA 2.2** **H17E0300**	**100 4** **VM/ESA 2.2** **H17E0302**	**Difference**	**%Difference**
Environment				
Real Storage	**64MB**	**64MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**300**	**300**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**1**	**1**		
Proc. Usage				
PBT/CMD (H)	84.618	83.675	-0.943	-1.11%
PBT/CMD	84.316	83.386	-0.929	-1.10%
CP/CMD (H)	33.714	33.194	-0.521	-1.54%
CP/CMD	27.786	26.837	-0.949	-3.41%
EMUL/CMD (H)	50.904	50.482	-0.422	-0.83%
EMUL/CMD	56.530	56.549	0.020	0.03%
Processor Util.				
TOTAL (H)	88.32	87.30	-1.01	-1.15%
TOTAL	88.00	87.00	-1.00	-1.14%
TOTAL EMUL (H)	53.13	52.67	-0.46	-0.86%
TOTAL EMUL	59.00	59.00	0.00	0.00%
TVR(H)	1.66	1.66	0.00	-0.29%
TVR	1.49	1.47	-0.02	-1.14%
Storage				
NUCLEUS SIZE (V)	2536KB	2536KB	0KB	0.00%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
WKSET (V)	86	86	0	0.00%
PGBLPGS	13506	13514	8	0.06%
PGBLPGS/USER	45.0	45.0	0.0	0.06%
FREEPGS	943	935	-8	-0.85%
FREE UTIL	0.89	0.90	0.01	0.98%
SHRPGS	1089	1086	-3	-0.28%
Paging				
READS/SEC	83	83	0	0.00%
WRITES/SEC	68	68	0	0.00%
PAGE/CMD	14.468	14.473	0.005	0.03%
PAGE IO RATE (V)	25.800	25.700	-0.100	-0.39%
PAGE IO/CMD (V)	2.472	2.463	-0.009	-0.35%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	8.623	8.530	-0.093	-1.08%
Queues				
DISPATCH LIST	11.22	10.60	-0.62	-5.49%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	117	117	0	0.00%
VIO/CMD	11.210	11.214	0.004	0.03%
RIO RATE (V)	70	69	-1	-1.43%
RIO/CMD (V)	6.707	6.613	-0.094	-1.39%
NONPAGE RIO/CMD (V)	4.235	4.150	-0.085	-2.00%
DASD RESP TIME (V)	23.800	24.000	0.200	0.84%
MDC READS (I/Os)	29	29	0	0.00%
MDC WRITES (I/Os)	1.41	1.47	0.06	4.26%
MDC AVOID	26	27	1	3.85%
MDC HIT RATIO	0.88	0.89	0.01	1.14%

Table 58 (Page 3 of 3). The effect of IABIAS on response times - 9221-170

IABIAS Settings **Release** **Run ID**	**90 2** **VM/ESA 2.2** **H17E0300**	**100 4** **VM/ESA 2.2** **H17E0302**	**Difference**	**%Difference**
Environment				
Real Storage	**64MB**	**64MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**300**	**300**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**1**	**1**		
PRIVOPs				
PRIVOP/CMD	14.506	14.411	-0.095	-0.65%
DIAG/CMD	32.499	32.464	-0.035	-0.11%
DIAG 04/CMD	6.225	6.171	-0.054	-0.87%
DIAG 08/CMD	0.734	0.738	0.004	0.53%
DIAG 0C/CMD	1.127	1.127	0.000	-0.02%
DIAG 14/CMD	0.024	0.025	0.000	0.80%
DIAG 58/CMD	1.252	1.252	0.000	0.02%
DIAG 98/CMD	2.283	2.258	-0.026	-1.12%
DIAG A4/CMD	3.603	3.601	-0.002	-0.05%
DIAG A8/CMD	2.825	2.821	-0.004	-0.15%
DIAG 214/CMD	13.278	13.318	0.040	0.30%
SIE/CMD	63.141	62.396	-0.745	-1.18%
SIE INTCPT/CMD	44.830	43.677	-1.153	-2.57%
FREE TOTL/CMD	58.542	58.371	-0.171	-0.29%
VTAM Machines				
WKSET (V)	286	285	-1	-0.35%
TOT CPU/CMD (V)	17.5392	16.8530	-0.6862	-3.91%
CP CPU/CMD (V)	6.8666	6.6294	-0.2372	-3.45%
VIRT CPU/CMD (V)	10.6725	10.2236	-0.4489	-4.21%
DIAG 98/CMD (V)	2.283	2.258	-0.025	-1.12%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

VM/VTAM 3.4.1

Workload: FS8F0R

Hardware Configuration

Processor model: 9121-480
Processors used: 2
Storage
  Real: 224MB
  Expanded: 32MB (all for MDC)
Tape: 3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2	16 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*
3088-08	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XC	10000	512	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1860	Users	3MB/XC	100		

Measurement Discussion: A pair of measurements was obtained on a 9121-480 with VM/ESA Release 2.1 to observe the performance effects of migrating from VTAM 3.4.0 to VTAM 3.4.1. The results showed equivalent response times within measurement variability and a 0.5% decrease in internal throughput (ITR (H)).

The ITR decrease was due to an increase in emulation time (VIRT CPU/CMD(V)) in the VTAM/VSCS virtual machine. This increase is attributed to the fact that, in VTAM 3.4.1, five of the VTAM internal trace options (PIU, API, NRM, SSCP, and MSG) are now forced on. This is expected to result in a substantial VTAM serviceability improvement.

No trace options were selected for the 3.4.0 base run. Customers who are running VTAM 3.4.0 with some of these options already in effect will see a smaller performance impact when migrating to VTAM 3.4.1.

Table 59 (Page 1 of 2). Migration from VM/VTAM 3.4.0 to VM/VTAM 3.4.1

VTAM Release **VM Release** **Run ID**	**VTAM 3.4.0** **VM/ESA 2.1** **L26E186N**	**VTAM 3.4.1** **VM/ESA 2.1** **L26E186I**	**Difference**	**%Difference**
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.123	0.123	0.000	0.00%
NONTRIV INT	0.409	0.428	0.019	4.65%
TOT INT	0.313	0.325	0.012	3.83%
TOT INT ADJ	0.276	0.287	0.011	3.95%
AVG FIRST (T)	0.235	0.240	0.005	2.34%
AVG LAST (T)	0.341	0.354	0.013	3.82%
Throughput				
AVG THINK (T)	26.23	26.16	-0.07	-0.27%
ETR	57.66	57.71	0.05	0.09%
ETR (T)	65.42	65.40	-0.02	-0.03%
ETR RATIO	0.881	0.882	0.001	0.12%
ITR (H)	73.66	73.27	-0.39	-0.52%
ITR	32.46	32.34	-0.13	-0.39%
EMUL ITR	48.30	48.01	-0.30	-0.61%
ITRR (H)	1.000	0.995	-0.005	-0.52%
ITRR	1.000	0.996	-0.004	-0.39%
Proc. Usage				
PBT/CMD (H)	27.153	27.296	0.143	0.53%
PBT/CMD	27.208	27.216	0.008	0.03%
CP/CMD (H)	9.480	9.477	-0.003	-0.03%
CP/CMD	9.018	8.868	-0.150	-1.67%
EMUL/CMD (H)	17.673	17.819	0.146	0.83%
EMUL/CMD	18.190	18.348	0.158	0.87%
Processor Util.				
TOTAL (H)	177.64	178.52	0.88	0.50%
TOTAL	178.00	178.00	0.00	0.00%
UTIL/PROC (H)	88.82	89.26	0.44	0.50%
UTIL/PROC	89.00	89.00	0.00	0.00%
TOTAL EMUL (H)	115.62	116.54	0.92	0.80%
TOTAL EMUL	119.00	120.00	1.00	0.84%
MASTER TOTAL (H)	88.54	89.03	0.49	0.55%
MASTER TOTAL	89.00	89.00	0.00	0.00%
MASTER EMUL (H)	51.32	51.86	0.54	1.05%
MASTER EMUL	53.00	54.00	1.00	1.89%
TVR(H)	1.54	1.53	0.00	-0.30%
TVR	1.50	1.48	-0.01	-0.83%
Storage				
NUCLEUS SIZE (V)	2364KB	2364KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	82	82	0	0.00%
PGBLPGS	47494	47502	8	0.02%
PGBLPGS/USER	25.5	25.5	0.0	0.02%
FREEPGS	5103	5097	-6	-0.12%
FREE UTIL	0.95	0.95	0.00	0.12%
SHRPGS	1192	1230	38	3.19%

Table 59 (Page 2 of 2). Migration from VM/VTAM 3.4.0 to VM/VTAM 3.4.1

VTAM Release VM Release Run ID	VTAM 3.4.0 VM/ESA 2.1 L26E186N	VTAM 3.4.1 VM/ESA 2.1 L26E186I	Difference	%Difference
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	613	612	-1	-0.16%
WRITES/SEC	432	432	0	0.00%
PAGE/CMD	15.973	15.963	-0.011	-0.07%
PAGE IO RATE (V)	172.000	171.800	-0.200	-0.12%
PAGE IO/CMD (V)	2.629	2.627	-0.002	-0.09%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	6.680	6.666	-0.013	-0.20%
Queues				
DISPATCH LIST	35.53	33.43	-2.10	-5.90%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	675	671	-4	-0.59%
VIO/CMD	10.318	10.259	-0.058	-0.56%
RIO RATE (V)	374	368	-6	-1.60%
RIO/CMD (V)	5.717	5.627	-0.090	-1.58%
NONPAGE RIO/CMD (V)	3.088	3.000	-0.088	-2.84%
DASD RESP TIME (V)	18.900	19.100	0.200	1.06%
MDC READS (blks)	366	364	-2	-0.55%
MDC WRITES (blks)	139	139	0	0.00%
MDC MODS	106	106	0	0.00%
MDC AVOID	174	173	-1	-0.57%
MDC HIT RATIO	0.91	0.91	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.956	13.967	0.011	0.08%
DIAG/CMD	27.999	27.810	-0.189	-0.68%
DIAG 04/CMD	2.552	2.496	-0.056	-2.21%
DIAG 08/CMD	0.739	0.738	-0.001	-0.14%
DIAG 0C/CMD	1.322	1.319	-0.003	-0.26%
DIAG 14/CMD	0.025	0.025	0.000	0.14%
DIAG 58/CMD	1.248	1.247	-0.001	-0.04%
DIAG 98/CMD	1.243	1.209	-0.034	-2.71%
DIAG A4/CMD	3.660	3.646	-0.013	-0.36%
DIAG A8/CMD	2.837	2.831	-0.006	-0.21%
DIAG 214/CMD	13.225	13.159	-0.066	-0.50%
SIE/CMD	51.787	51.649	-0.138	-0.27%
SIE INTCPT/CMD	34.697	34.605	-0.092	-0.27%
FREE TOTL/CMD	67.256	67.092	-0.164	-0.24%
VTAM Machines				
WKSET (V)	518	525	7	1.35%
TOT CPU/CMD (V)	3.7619	3.9669	0.2050	5.45%
CP CPU/CMD (V)	1.5455	1.5630	0.0175	1.13%
VIRT CPU/CMD (V)	2.2164	2.4039	0.1875	8.46%
DIAG 98/CMD (V)	1.243	1.208	-0.035	-2.79%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

3745 Comparison to CTCA

This section presents and discusses the results of a comparable pair of measurements on VM/ESA Release 2.1 where one measurement was with CTCA connectivity and the other used a 3745.[16] This information is useful for getting a better understanding of how the CMS-intensive performance results shown in this report (which used CTCA) relate to environments that have network connectivity through 3745s or similar communication control units.

Workload: FS7FMAXR

Hardware Configuration

Processor model:	9121-480
Processors used:	2
Storage	
Real:	224MB
Expanded:	32MB (all for MDC)
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	4	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2		16 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*
3088-08	1	NA	4.5MB
or			
3745-410[17]	1	42	56Kb

[16] These results were first documented in *VM/ESA Release 2.1 Performance Report* and are reprinted here for convenience.

[17] The 3745-410 control unit was 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 3000 users.

Software Configuration

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

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
CRRSERV1	1	SFS	16MB/XC	100		
ROSERV1	1	SFS	32MB/XC	100		QUICKDSP ON
RWSERVn	2	SFS	64MB/XC	1500	1200 (CTCA) 1300 (3745)	QUICKDSP ON
SMART	1	RTM	16MB/370	3%	300	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	450 (CTCA) 600 (3745)	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Uxxxx	1620	Users	3MB/XC	100		

Measurement Discussion: The number of users was adjusted so that the two measurements were done at approximately the same processor utilization (88%).

The number of pages reserved for the VTAM machine was increased from 450 in the CTCA case to 600 in the 3745 case in order to accommodate the larger VTAM working set size that occurs in the 3745 case. The number of pages that was reserved for the SFS server virtual machines was slightly different (1200 and 1300) for the two measurements but the measurement data indicate that this had no significant effect on the results.

As expected, the two cases showed somewhat different performance characteristics. Relative to the CTCA measurement, the 3745 measurement had a 30% higher (0.136 seconds) external response time (AVG LAST (T)), while the internal throughput rate (ITR (H)) improved by 1.8%.

The response time increase is probably due to the reduced bandwidth of the 3745 (56Kb) relative to the CTCA (4.5MB). However, the details of how VTAM and the 3745 were configured may be a factor as well.

Some of the VTAM working set increase is due to the 3.9% increase in the number of users that were run. However, most of the increase comes from changing from CTCA to 3745 communications. One contributing factor is that message buffers are held longer in the 3745 case because of the lower bandwidth.

The ITR increase is entirely due to a decrease in VTAM processor usage. This is apparent from the results because the decrease in TOT CPU/CMD (V) for the VTAM machine (-0.59 ms) is very similar to the system-wide decrease in total processor usage (-0.56 ms, PBT/CMD (H)). There are two main reasons for this improvement:

1. Some of the processing that was done by VTAM in the CTCA case is instead handled by the 3745.

2. There is a 72% decrease in DIAGNOSE code X¢98¢ (DIAG 98/CMD (V)), which VTAM uses to do network I/O. This decrease occurs as a consequence of the longer response times, which allow more time for multiple messages to accumulate and be transmitted as a single DIAGNOSE code X¢98¢ request.

The large changes in trivial response time (TRIV INT) and total internal response time (TOT INT) are mostly due to a large difference in what CP determines to be a transaction. This is reflected by the 47% increase in the ETR metric, which is the number of CP-determined transactions per second (as reported by RTM).

CP has no way of knowing with certainty when one transaction stops and the next transaction starts, so it does the best it can. Whenever a virtual machine is in voluntary wait for more than 300 milliseconds, CP assumes that a new transaction has started. In the CTCA case, the CTCA is so fast that short interactions that happen to be separated by a short think time (such as 0.1 seconds) are so close together that they are separated by less than this 300 milliseconds. Whenever this happens, two real user/system interactions get combined into one CP transaction. In the 3745 case, these same interactions appear to CP as separate transactions.

The true change in internal response time is represented by TOT INT ADJ, which normalizes TOT INT to ETR (T), the TPNS-based transaction rate. ETR (T) counts the number of actual user/system interactions and is consistent in meaning from one measurement to another.

Table 60 (Page 1 of 3). 3745 comparison to CTCA / FS7FMAXR on the 9121-480

Communications **Release** **Run ID**	**CTCA** **VM/ESA 2.1** **L26S1563**	**3745** **VM/ESA 2.1** **L26S1620**	**Difference**	**%Difference**
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1560**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.110	0.051	-0.059	-53.64%
NONTRIV INT	0.421	0.444	0.023	5.46%
TOT INT	0.319	0.246	-0.073	-22.88%
TOT INT ADJ	0.280	0.312	0.032	11.43%
AVG FIRST (T)	0.228	0.274	0.046	20.18%
AVG LAST (T)	0.346	0.482	0.136	39.51%
Throughput				
AVG THINK (T)	26.25	26.49	0.24	0.91%
ETR	48.21	70.93	22.72	47.13%
ETR (T)	54.92	55.92	1.00	1.82%
ETR RATIO	0.878	1.268	0.391	44.50%
ITR (H)	62.58	63.69	1.11	1.78%
ITR	27.49	40.41	12.93	47.03%
EMUL ITR	41.17	60.32	19.15	46.53%
ITRR (H)	1.000	1.018	0.018	1.78%
ITRR	1.000	1.470	0.470	47.03%

3745 Comparison to CTCA

Table 60 (Page 2 of 3). 3745 comparison to CTCA / FS7FMAXR on the 9121-480

Communications	**CTCA**	**3745**		
Release	**VM/ESA 2.1**	**VM/ESA 2.1**	**Difference**	**%Difference**
Run ID	**L26S1563**	**L26S1620**		
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1560**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Proc. Usage				
PBT/CMD (H)	31.961	31.403	-0.558	-1.74%
PBT/CMD	31.865	31.474	-0.391	-1.23%
CP/CMD (H)	11.251	10.972	-0.279	-2.48%
CP/CMD	10.561	10.372	-0.189	-1.79%
EMUL/CMD (H)	20.710	20.431	-0.278	-1.34%
EMUL/CMD	21.304	21.102	-0.202	-0.95%
Processor Util.				
TOTAL (H)	175.52	175.60	0.08	0.04%
TOTAL	175.00	176.00	1.00	0.57%
UTIL/PROC (H)	87.76	87.80	0.04	0.04%
UTIL/PROC	87.50	88.00	0.50	0.57%
TOTAL EMUL (H)	113.74	114.25	0.51	0.45%
TOTAL EMUL	117.00	118.00	1.00	0.85%
MASTER TOTAL (H)	87.59	87.85	0.26	0.30%
MASTER TOTAL	88.00	88.00	0.00	0.00%
MASTER EMUL (H)	52.00	52.32	0.32	0.61%
MASTER EMUL	54.00	54.00	0.00	0.00%
TVR(H)	1.54	1.54	-0.01	-0.41%
TVR	1.50	1.49	0.00	-0.28%
Storage				
NUCLEUS SIZE (V)	2364KB	2364KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	79	79	0	0.00%
PGBLPGS	48821	48537	-284	-0.58%
PGBLPGS/USER	31.3	30.0	-1.3	-4.26%
FREEPGS	4262	4495	233	5.47%
FREE UTIL	0.91	0.90	-0.01	-1.60%
SHRPGS	1259	1239	-20	-1.59%
Paging				
READS/SEC	516	522	6	1.16%
WRITES/SEC	342	347	5	1.46%
PAGE/CMD	15.623	15.540	-0.083	-0.53%
PAGE IO RATE (V)	126.000	127.200	1.200	0.95%
PAGE IO/CMD (V)	2.294	2.275	-0.020	-0.85%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	6.701	6.778	0.077	1.15%
Queues				
DISPATCH LIST	32.35	39.51	7.16	22.13%
ELIGIBLE LIST	0.00	0.00	0.00	na

Table 60 (Page 3 of 3). 3745 comparison to CTCA / FS7FMAXR on the 9121-480

Communications	**CTCA**	**3745**		
Release	**VM/ESA 2.1**	**VM/ESA 2.1**	**Difference**	**%Difference**
Run ID	**L26S1563**	**L26S1620**		
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1560**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
I/O				
VIO RATE	528	484	-44	-8.33%
VIO/CMD	9.614	8.655	-0.959	-9.97%
RIO RATE (V)	347	297	-50	-14.41%
RIO/CMD (V)	6.318	5.311	-1.007	-15.94%
MDC READS	261	268	7	2.68%
MDC WRITES	123	129	6	4.88%
MDC MODS	77	80	3	3.90%
MDC HIT RATIO	0.84	0.84	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	24.777	24.847	0.070	0.28%
DIAG/CMD	24.411	23.371	-1.039	-4.26%
DIAG 08/CMD	0.759	0.760	0.001	0.13%
DIAG 14/CMD	0.024	0.025	0.000	0.51%
DIAG 58/CMD	1.251	1.250	-0.001	-0.06%
DIAG 98/CMD	1.478	0.418	-1.060	-71.72%
DIAG A4/CMD	2.095	2.101	0.005	0.25%
DIAG A8/CMD	1.631	1.648	0.017	1.02%
DIAG 214/CMD	12.240	12.213	-0.027	-0.22%
SIE/CMD	62.584	58.031	-4.553	-7.27%
SIE INTCPT/CMD	45.060	41.202	-3.858	-8.56%
FREE TOTL/CMD	72.416	71.515	-0.902	-1.25%
VTAM Machines				
WKSET (V)	461	631	170	36.88%
TOT CPU/CMD (V)	4.1172	3.5269	-0.5903	-14.34%
CP CPU/CMD (V)	1.7399	1.4903	-0.2496	-14.35%
VIRT CPU/CMD (V)	2.3773	2.0367	-0.3406	-14.33%
DIAG 98/CMD (V)	1.477	0.418	-1.059	-71.71%
SFS Servers				
WKSET (V)	2997	3204	207	6.91%
TOT CPU/CMD (V)	3.6620	3.7256	0.0636	1.74%
CP CPU/CMD (V)	1.7096	1.7486	0.0390	2.28%
VIRT CPU/CMD (V)	1.9524	1.9771	0.0247	1.27%
FP REQ/CMD(Q)	1.263	1.269	0.006	0.48%
IO/CMD (Q)	1.821	1.891	0.070	3.84%
IO TIME/CMD (Q)	0.022	0.024	0.002	9.09%
SFS TIME/CMD (Q)	0.029	0.032	0.003	10.34%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Query Filepool Counters, Unmarked=RTM

FS7F to FS8F Workload Comparison

A number of incremental improvements were made to the FS7F CMS-intensive workload:

- larger NAMES file
- one T-disk was changed to a virtual disk in storage
- use of CMS pipelines
- HASM replaced by HLASM
- increased spool file activity
- more representative use of the TELL command

The updated workload was renamed FS8F. This section provides measurement results on the 9121-480 that bridge from the FS7F workload to the new FS8F workload.

See Appendix A, “Workloads” on page 208 for workload descriptions.

Workloads: FS7F0R and FS8F0R

Hardware Configuration

Processor model:	9121-480
Processors used:	2
Storage	
Real:	224MB
Expanded:	32MB (all for MDC)
Tape:	3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-3	2						2 R
3390-2	3990-2	4	16	5	6			
3390-2	3990-3	4		2	2	16 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*
3088-08	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*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	512	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1860	Users	3MB/XC	100		

Measurement Discussion: The performance of the two workloads is quite similar. Relative to FS7F, the FS8F workload showed a 7% decrease in external response time (AVG LAST(T)) and a 1.5% increase in internal throughput (ITR(H). The use of virtual disks in storage was one of the more significant changes, and accounts for much of the 9.7% decrease in non-paging real I/Os (NONPAGE RIO/CMD(V)).

Table 61 (Page 1 of 3). Comparison of the FS7F and FS8F CMS Workloads

Release Run ID	**VM/ESA 2.1 L26E186G**	**VM/ESA 2.1 L26E186I**	**Difference**	**%Difference**
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.126	0.123	-0.003	-2.38%
NONTRIV INT	0.458	0.428	-0.030	-6.55%
TOT INT	0.347	0.325	-0.022	-6.34%
TOT INT ADJ	0.307	0.287	-0.020	-6.61%
AVG FIRST (T)	0.250	0.240	-0.010	-3.80%
AVG LAST (T)	0.381	0.354	-0.028	-7.34%
Throughput				
AVG THINK (T)	26.34	26.16	-0.17	-0.66%
ETR	57.67	57.71	0.04	0.07%
ETR (T)	65.17	65.40	0.24	0.36%
ETR RATIO	0.885	0.882	-0.003	-0.29%
ITR (H)	72.21	73.27	1.06	1.47%
ITR	31.97	32.34	0.36	1.14%
EMUL ITR	46.97	48.01	1.03	2.20%
ITRR (H)	1.000	1.015	0.015	1.47%
ITRR	1.000	1.011	0.011	1.14%
Proc. Usage				
PBT/CMD (H)	27.696	27.296	-0.400	-1.45%
PBT/CMD	27.621	27.216	-0.406	-1.47%
CP/CMD (H)	9.419	9.477	0.058	0.62%
CP/CMD	8.747	8.868	0.121	1.39%
EMUL/CMD (H)	18.278	17.819	-0.459	-2.51%
EMUL/CMD	18.875	18.348	-0.527	-2.79%

FS7F to FS8F Workload Comparison

Table 61 (Page 2 of 3). Comparison of the FS7F and FS8F CMS Workloads

Release Run ID	VM/ESA 2.1 L26E186G	VM/ESA 2.1 L26E186I	Difference	%Difference
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Processor Util.				
TOTAL (H)	180.49	178.52	-1.96	-1.09%
TOTAL	180.00	178.00	-2.00	-1.11%
UTIL/PROC (H)	90.24	89.26	-0.98	-1.09%
UTIL/PROC	90.00	89.00	-1.00	-1.11%
TOTAL EMUL (H)	119.11	116.54	-2.57	-2.16%
TOTAL EMUL	123.00	120.00	-3.00	-2.44%
MASTER TOTAL (H)	90.07	89.03	-1.04	-1.15%
MASTER TOTAL	90.00	89.00	-1.00	-1.11%
MASTER EMUL (H)	53.61	51.86	-1.75	-3.27%
MASTER EMUL	55.00	54.00	-1.00	-1.82%
TVR(H)	1.52	1.53	0.02	1.09%
TVR	1.46	1.48	0.02	1.36%
Storage				
NUCLEUS SIZE (V)	2364KB	2364KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	82	82	0	0.00%
PGBLPGS	47724	47502	-222	-0.47%
PGBLPGS/USER	25.7	25.5	-0.1	-0.47%
FREEPGS	4985	5097	112	2.25%
FREE UTIL	0.98	0.95	-0.02	-2.20%
SHRPGS	1180	1230	50	4.24%
Paging				
READS/SEC	628	612	-16	-2.55%
WRITES/SEC	425	432	7	1.65%
PAGE/CMD	16.159	15.963	-0.196	-1.21%
PAGE IO RATE (V)	173.600	171.800	-1.800	-1.04%
PAGE IO/CMD (V)	2.664	2.627	-0.037	-1.39%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	6.552	6.666	0.114	1.74%
Queues				
DISPATCH LIST	39.21	33.43	-5.78	-14.75%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	624	671	47	7.53%
VIO/CMD	9.575	10.259	0.684	7.14%
RIO RATE (V)	390	368	-22	-5.64%
RIO/CMD (V)	5.985	5.627	-0.358	-5.98%
NONPAGE RIO/CMD (V)	3.321	3.000	-0.321	-9.66%
DASD RESP TIME (V)	19.900	19.100	-0.800	-4.02%
MDC READS (blks)	361	364	3	0.83%
MDC WRITES (blks)	178	139	-39	-21.91%
MDC MODS	144	106	-38	-26.39%
MDC AVOID	180	173	-7	-3.89%
MDC HIT RATIO	0.91	0.91	0.00	0.00%

Table 61 (Page 3 of 3). Comparison of the FS7F and FS8F CMS Workloads

Release Run ID	VM/ESA 2.1 L26E186G	VM/ESA 2.1 L26E186I	Difference	%Difference
Environment				
Real Storage	**224MB**	**224MB**		
Exp. Storage	**32MB**	**32MB**		
Users	**1860**	**1860**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
PRIVOPs				
PRIVOP/CMD	14.113	13.967	-0.146	-1.04%
DIAG/CMD	27.035	27.810	0.775	2.87%
DIAG 04/CMD	2.480	2.496	0.016	0.64%
DIAG 08/CMD	0.760	0.738	-0.022	-2.91%
DIAG 0C/CMD	1.341	1.319	-0.022	-1.67%
DIAG 14/CMD	0.025	0.025	0.000	0.00%
DIAG 58/CMD	1.249	1.247	-0.002	-0.13%
DIAG 98/CMD	1.208	1.209	0.001	0.10%
DIAG A4/CMD	3.919	3.646	-0.273	-6.96%
DIAG A8/CMD	1.885	2.831	0.946	50.16%
DIAG 214/CMD	12.958	13.159	0.201	1.55%
SIE/CMD	51.453	51.649	0.196	0.38%
SIE INTCPT/CMD	34.473	34.605	0.132	0.38%
FREE TOTL/CMD	67.826	67.092	-0.734	-1.08%
VTAM Machines				
WKSET (V)	522	525	3	0.57%
TOT CPU/CMD (V)	4.0068	3.9669	-0.0399	-1.00%
CP CPU/CMD (V)	1.5772	1.5630	-0.0142	-0.90%
VIRT CPU/CMD (V)	2.4297	2.4039	-0.0258	-1.06%
DIAG 98/CMD (V)	1.207	1.208	0.001	0.11%

Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM

Appendix A. Workloads

The workloads that were used to evaluate VM/ESA Release 2.2 are described in this appendix.

CMS-Intensive (FS7F and FS8F)

The FS8F workload is similar to FS7F (used in previous VM/ESA releases) except for the following enhancements that were added to exploit new function and to be more representative of what customers are doing.

- larger NAMES file
- one T-disk was changed to a virtual disk in storage
- use of CMS pipelines
- HASM replaced by HLASM
- increased spool file activity
- more representative use of the TELL command

Most measurements in this report were done with the FS8F workload but a small subset were done with the FS7F workload. For a description of the FS7F workload, refer to the *VM/ESA Release 2.1 Performance Report*.

FS8F Workload Description

FS8F simulates a CMS user environment, with variations simulating a minidisk environment, an SFS environment, or some combination of the two. Table 62 shows the search-order characteristics of the two environments used for measurements discussed in this document.

Table 62. FS8F workload characteristics

Filemode	ACCESS	Number of Files	FS8F0R	FS8FMAXR
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	*m*	minidisk	minidisk
Y	R/O	*n*	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.

The measurement environments have the following characteristics in common:

- The HELP disk has the FSTs saved in a shared segment.

- The CMSINST, VMLIB, and VMMTLIB shared segments are used.
- A Bactrian-distribution think time averaging 30 seconds is used. (See “Glossary of Performance Terms” on page 228 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 workload. TPNS runs in a separate processor and simulates LU2 terminals. User traffic travels between the processors through 3088 multisystem channel communication units.

FS8F Variations

Two FS8F workload variants were used for measurements, one for minidisk-based CMS users, and the other for SFS-based CMS users.

FS8F0R 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.

FS8FMAXR Workload: All file modes, except S and Y (which SFS does not support), the HELP minidisk, and T-disks that are created by the workload, 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.

FS8F Licensed Programs

The following licensed programs were used in the FS8F measurements described in this document:

- VS COBOL II Compiler and Library V1R4M0
- Document Composition Facility V1R4M0
- VS FORTRAN Compiler/Library/Debug V2R5M0
- Assembler H V2R1M0 or IBM High Level AssemblerV1R1M0
- OS PL/I V2R3M0 Compiler & Library
- C & PL/I Common Library V1R2M0
- VTAM V3R4M1
- NCP V5R4M0

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 15 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.

FS8F Script Description

FS8F 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 15 minute interval. After the selected interval, each user starts running a general workload script. The scripts are summarized in Table 63.

Table 63. FS8F workload script summary

Script Name	% Used	Script Description
LOGESA	*	Logon and Initialization
WAIT	*	Wait state
CMSSTRT	*	Stagger start of user activity
ASM617F	5	Assemble (HLASM) and Run
ASM627F	5	Assemble and Run
XED117F	5	Edit a VS BASIC Program
XED127F	10	Edit a VS BASIC Program
XED137F	10	Edit a COBOL Program
XED147F	10	Edit a COBOL Program
COB217F	5	COBOL Compile
COB417F	5	Run a COBOL Program
FOR217F	5	VS FORTRAN Compile
FOR417F	5	FORTRAN Run
PRD517F	5	Productivity Aids Session
DCF517F	5	Edit and Script a File
PLI317F	5	PL/I Optimizer Session
PLI717F	5	PL/I Optimizer Session
WND517F	8	Run Windows with IPL CMS
WND517FL	2	Run Windows with LOGON/LOGOFF
HLP517F	5	Use HELP

Note: Scripts with an asterisk (*) in the "% Used" column are run only once each for each user during initialization.

The following are descriptions of each script used in the FS8F workload.

LOGESA: Initialization Script

LOGON userid
SET AUTOREAD ON
IF FS8F0R workload
THEN
  Erase extraneous files from 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
Clear the screen
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 15 minutes, the start for each user to prevent all users from starting scripts at the same time.

ASM617F: Assemble (HLASM) 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 HLASM (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

ASM627F: Assemble (F-Assembler) and Run

QUERY reader and printer
Clear the screen
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

XED117F: Edit a VS BASIC Program

XEDIT the program
Get into input mode
Enter 29 input lines
Quit without saving file (QQUIT)

XED127F: 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

XED137F: 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

XED147F: 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

COB217F: Compile a COBOL Program

Set ready message short
Clear the screen
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

COB417F: Run a COBOL Program

Define temporary disk space for 2 disks using an EXEC
Clear the screen
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

FOR217F: Compile 6 VS FORTRAN Programs

NUCXDROP NAMEFIND using an EXEC
Clear the screen
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 command
Clear the screen
Compile a FORTRAN program
Issue INDICATE commands
Compile another FORTRAN program
Issue INDICATE commands
Compile another FORTRAN program
Clear the screen
Issue INDICATE command
Erase extraneous files from A-disk
PURGE the printer

FOR417F: Run 2 FORTRAN Programs

SPOOL PRT CLASS D
Clear the screen
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 do not issue spool printer
List and erase output files
Reset GLOBALs and clear FILEDEFs

PRD517F: Productivity Aids Session

Run an EXEC to set up names file for user
Clear the screen
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

DCF517F: 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

PLI317F: 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
COPY names file and RENAME it
TELL a group of users one pass of script run
ERASE names file
PURGE the printer

PLI717F: 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 another USERID one pass of script run
PURGE the printer

WND517F: 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

WND517FL: 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

HLP517F: 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

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 PACEX*c* workload, in order to scale the load to the environment being tested. The most commonly-used multiples are PACEX4, PACEX6, and PACEX8.

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 PACEX*c* 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 PACEX8; 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.

PACE Workload

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 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 for this report are taken at the 70% processor utilization point.

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%). 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 1.3.2, CICS/VSE 2.2.0, ACF/VTAM* 3.4.0, and VSE/VSAM 2.2.0. 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 executes a combination of COBOL and 24-bit assembler applications to produce an 80% read and 20% write mixture. Each application uses several transactions that employ differing sets of CICS functions. The following table indicates the number of transactions for each application and the frequency of specific CICS functions within each:

Table 64. CICS/VSE Transaction Characteristics

TRANSACTION TYPE	VSAM CALLS	READ	READ NEXT	ADD	UPDATE	DELETE	TRANS DATA	TEMP STOR	% MIX
Banking	3	2			1	1			8
	10	8		2		2			8
Hotel	2	1			1				3
Reservations	2	1		1					3
Inventory Control	0							1	3
	17	1	16						6
	14	5	14	3	2				8
Order Entry	3	1					1		5
	3	1	2	1	1		2		5
	9	9					2		5
	22	9		4	9		1		5
Product	18	8	10						10
Specification	34	2	32						9
Stock Control	18	1	8		9				5
	9			9					3
	3	1	1			1			10
Teller System	0								4

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.

The ITR is calculated as

$$\frac{\text{transactions}}{\text{processor busy seconds}}$$

Connectivity (CONN1)

The CONN1 workload exercises APPC/VM data transfers. CONN1 consists of two parts: the client part and the server part. Each part runs in a separate virtual machine. The basic scenario is as follows:

- Initialization Process
  1. The server application is started. It identifies itself to CP as a resource and waits for APPC/VM interrupts.
  2. The client application is started. It issues a CONNECT to the server.
  3. The server issues an APPC/VM RECEIVE for that client. The server is ready to receive requests.
  4. The client application prompts the user for the number of APPC/VM iterations (i) to perform and the number of bytes (n) that the server should transfer back to the client for each iteration.
- Send/Receive Process

  The following steps are performed (i) times:
  1. The client application issues an APPC/VM SENDDATA with RECEIVE=YES WAIT=YES, supplying an ANSBUF area for the server's reply.
  2. The server's RECEIVE (or SENDDATA RECEIVE=YES) completes. It now has a request from a client for (n) bytes of information.
  3. The server issues an APPC/VM SENDDATA RECEIVE=YES WAIT=NO in order to send (n) bytes back to that client and wait for the next request.
  4. The client resumes execution after the APPC/VM SENDDATA that it had issued.

  The data transfers described above are from memory to memory. There is no DASD I/O after the applications are loaded.

  CONN1 uses the CMSIUCV interface to set up APPC connections.

Appendix B. 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 in the DMSDAC and DMSSAC logical segments.
CMSPIPES	Contain CMSPIPES code in the PIPES logical segment.
CMSINST	Contains the execs-in-storage segment.
CMSVMLIB	Contains the following logical segments: • VMLIB contains the CSL code. • VMMTLIB contains the CMS multitasking code. • CMSQRYL and CMSQRYH contain the code for some CMS QUERY and SET commands. This code would otherwise be read from the S-disk when these commands are used.
HELP	Contains FSTs for the HELP disk.
GOODSEG	Contains FSTs for the C-disk. The C-disk is in the CMS search order used by the minidisk version of the FS8F workload.
FORTRAN	This segment space has two members: DSSVFORT for the FORTRAN compiler and FTNLIB20 for the library composite modules.
DSMSEG4B	Contains DCF (Document Composition Facility) code.
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 MAINT U3 OPERATOR MARK
NOBACKUP
FULLDUMP
FILEPOOLID fp_name
NOFORMAT
ACCOUNT
CATBUFFERS 415
MSGS
SAVESEGID CMSFILES
USERS  nnnn
```

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 or the *VM/ESA Performance* 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 MAINT U3 OPERATOR MARK
NOBACKUP
FULLDUMP
FILEPOOLID fp_name
NOFORMAT
ACCOUNT
MSGS
SAVESEGID CMSFILES
CRR
LUNAME lu_name
```

For more information on CRR and CRR tuning parameters, see the *SFS and CRR Planning, Administration, and Operation* manual or the *VM/ESA Performance* manual.

Appendix C. Master Table of Contents

This appendix provides a high-level table of contents that covers all of the performance measurement results that are published in the VM/ESA performance reports. This information is provided in two tables. Table 65 covers all performance measurement results except for migration results, which are covered by Table 28 on page 78. Both of these tables refer to the performance reports using the following notation:

10 *VM/ESA Release 1.0 Performance Report*

11 *VM/ESA Release 1.1 Performance Report*

20 *VM/ESA Release 2.0 Performance Report*

21 *VM/ESA Release 2.1 Performance Report*

22 *VM/ESA Release 2.2 Performance Report* (this document)

See "Referenced Publications" on page ix for more information on these reports.

Table 65 (Page 1 of 2). Sources of VM performance measurement results

Subject	Report(s)
Migration	see page 78
New Functions	
Coordinated Resource Recovery	10
VM Data Spaces (Use by SFS)	11
3990-3 DASD Fast Write Support	11
CMS Pipelines	11
Inter-System Facility for Communications (ISFC)	11 22
ECKD* Support	11
FBA DASD Support	20
CP Configurability	20
DIAGNOSE Code X'250'	20
Extended CMS File System Interfaces	20
Virtual Disk in Storage	21 22
Load Wait State PSW Improvements	21
REXX SAA* Level 2 Architecture	21
Minidisk Cache Enhancements	22
Share Capping and Proportional Distribution	22
SPXTAPE Command	22

Table 65 (Page 2 of 2). Sources of VM performance measurement results

Subject	Report(s)
Special Environments	
Capacity of a Single VTAM/VSCS Virtual Machine	10
APPC/VM	10
APPC/VM VTAM Support (AVS)	10
Effect of Virtual Machine Mode (370, XA, XC)	10 11
Minidisk to SFS	10 11 20
Effect of Real/Expanded Storage Size	11
Effect of Virtual Machine Size	11
LPAR Performance	20
RACF* 1.9	20
VSE Guests using Shared DASD	20
VMSES/E	20 21 22
VSE/ESA Guest Performance (Mode Variations)	21
3745 Comparison to CTCA	21 22
Tuning Studies	
Recommended 9221 Tuning	11
GCS IPOLL Option	11
Using Expanded Storage for MDC on a 9121	11
SET RESERVE	11
OfficeVision MSGFLAGS Setting	11
CMS File Cache for SFS	20
I/O Assist for Guests	20
Adjusting the Minidisk File Cache Size	21
VM/ESA REXX Performance Guidelines	21
Minidisk Cache Tuning: Restricting the Arbiter	21 22
Effect of IABIAS on Response Time	22
Processor Capacity	
3090-600J	10
9121-480	11
9021-720	11
9021-900	20
9121-742	21
9021-941	22
Additional Studies	
Greater than 16M of Real Storage (370 Feature)	10
CMS Instruction Trace Data	10
Measurement Variability	20 21
High Level Assembler Evaluation	21
CP Monitor Overhead	21
Processor-Constrained Environment	22
Comparison of VM/VTAM 3.4.0 to 3.4.1	22
FS7F to FS8F Workload Comparison	22

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/ESA Realtime Monitor.

(C). Denotes data obtained from the VSE console timestamps or from the CICSPARS reports (CICS transaction performance data).

(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.

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.

Absolute Share. An ABSOLUTE share allocates to a virtual machine an absolute percentage of all the available system resources.

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:

$$\frac{1}{1000} \times \sum_{f \in \text{filepools}} \frac{Agent\_Holding\_Time_f}{\text{SFSTIME}_f}$$

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:

$$\frac{1}{1000} \times \sum_{f \in \text{filepools}} \frac{Filpool\_Request\_Service\_Time_f}{\text{SFSTIME}_f}$$

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:

$$\frac{\text{Agents In Call}}{\sum_{f \in \text{filepools}} \frac{Total\_Filepool\_Requests_f}{\text{SFSTIME}_f}}$$

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:

$$\frac{1}{\text{ETR (T)}} \times \sum_{t \in \text{TPNS machines}} \frac{First\_Response_t \times Total\_Transmits_t}{TPNS\_Time_t}$$

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:

$$\frac{1}{\text{ETR (T)}} \times \sum_{t \in \text{TPNS machines}} \frac{Last\_Response_t \times Total\_Transmits_t}{TPNS\_Time_t}$$

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:

$$\frac{\sum_{f \in \text{filepools}} \frac{Lock\_Wait\_Time_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{Total\_Lock\_Conflicts_f}{\text{SFSTIME}_f}}$$

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:

$$\frac{\sum_{f \in \text{filepools}} \frac{Agent\_Holding\_Time_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{Begin\_LUWs_f}{\text{SFSTIME}_f}}$$

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:

$$\frac{1}{\text{ETR (C)}} \times \sum_{t \in \text{CICSPARS files}} \frac{Last\_Response_t \times Total\_Transmits_t}{CICS\_Time_t}$$

Last_Response is taken from the AVG TASK RESPONSE TIME line and *Total_Transmits* is from the TOTAL TASKS SELECTED line the CICSPARS 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:

$$\frac{1}{\text{ETR (T)}} \times \sum_{t \in \text{TPNS machines}} \frac{Think\_Time_t \times Total\_Transmits_t}{TPNS\_Time_t}$$

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:

$$\frac{\sum_{f \in \text{filepools}} \frac{Total\_BIO\_Request\_Time_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{Total\_BIO\_Requests_f}{\text{SFSTIME}_f}}$$

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:

$$\frac{\sum_{f \in \text{filepools}} \frac{Total\_DASD\_Block\_Transfers_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{Total\_BIO\_Requests_f}{\text{SFSTIME}_f}}$$

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:

$$\frac{\sum_{f \in \text{filepools}} \frac{Total\_DASD\_Block\_Transfers_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{Total\_IO\_Requests_f}{\text{SFSTIME}_f}}$$

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:

$$\frac{1}{1000} \times \frac{\sum_{f \in \text{filepools}} Checkpoint\_Time_f}{\sum_{f \in \text{filepools}} Checkpoints\_Taken_f}$$

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:

$$\frac{1}{10} \times \sum_{f \in \text{filepools}} \frac{Checkpoint\_Time_f}{\text{SFSTIME}_f}$$

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:

$$\sum_{f \in \text{filepools}} Checkpoints\_Taken_f$$

Checkpoints_Taken is from the QUERY FILEPOOL STATUS command.

CICSPARS. CICS Performance Analysis Reporting System, a licensed program that provides CICS response time and transaction information.

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 FS7F, FS8F, and VSECICS 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 the FS7F, FS8F, and VSECICS workloads:

$$10 \times \frac{\text{(TOTAL-TOTAL EMUL)}}{\text{ETR (T)}}$$

For the PACE workload:

PBT/CMD-EMUL/CMD

CP/CMD (H). See CP/CMD. This is the hardware based measure. This is calculated by:

For 9221 processors:

For the FS7F, FS8F, and VSECICS workloads:

$$\frac{CP\_CPU\_PCT \times \text{TOTAL (H)}}{10 \times \text{ETR (T)}}$$

For the PACE workload:

$$6000 \times \frac{CP\_CPU\_PCT \times \text{TOTAL (H)}}{\text{ETR (H)}}$$

CP_CPU_PCT is taken from the Host CPU Busy line in the CPU Busy/MIPs section of the RE0 report.

For all workloads running on 9121 and 9021 processors:

PBT/CMD (H)-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:

$$\left(\frac{1}{V\_Time \times \text{ETR (T)}}\right) \times \sum_{s \in \text{server class}} (TCPU_s\text{-}VCPU_s)$$

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.

CPU PCT BUSY (V). CPU Percent Busy. The percentage of total available processor time used by the designated virtual machine. Total available processor time is the sum of online time for all processors and represents total processor capacity (not processor usage).

This is the from the CPU Pct field in the VMPRF USER_RESOURCE_USER report.

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:

$$\frac{\text{DASD IO TOTAL (V)}}{V\_Time}$$

V_Time is taken from the time stamps at the beginning of the VMPRF DASD Activity Ordered by Activity 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:

$$\sum_{d \in \text{VSE Guest DASD}} Total_d$$

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.

DASD RESP TIME (V). Average DASD response time in milliseconds. This includes DASD service time plus

(except for page and spool volumes) any time the I/O request is queued in the host until the requested device becomes available.

This is taken from the DASD Resp Time field in the VMPRF SYSTEM_SUMMARY_BY_TIME report.

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:

$$\sum_{f \in \text{filepools}} Deadlocks_f$$

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 04/CMD. The number of DIAGNOSE code X¢04¢ instructions used per command. DIAGNOSE code X¢04¢ is the privilege class C and E CP function call to examine real storage. This is a product-sensitive programming interface. This is calculated by:

$$\frac{DIAG\_04}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_04 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 08/CMD. The number of DIAGNOSE code X¢08¢ instructions used per command. DIAGNOSE code X¢08¢ is the CP function call to issue CP commands from an application. This is calculated by:

$$\frac{DIAG\_08}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_08 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 0C/CMD. The number of DIAGNOSE code X¢0C¢ instructions used per command. DIAGNOSE code X¢0C¢ is the CP function call to obtain the time of day, virtual CPU time used by the virtual machine, and total CPU time used by the virtual machine. This is calculated by:

$$\frac{DIAG\_0C}{\text{RTM\_Time} \times \text{ETR (T)}}$$

DIAG_0C is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 10/CMD. The number of DIAGNOSE code X¢10¢ instructions used per command. DIAGNOSE code X¢10¢ is the CP function call to release pages of virtual storage. This is calculated by:

$$\frac{DIAG\_10}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_10 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 14/CMD. The number of DIAGNOSE code X¢14¢ instructions used per command. DIAGNOSE code X¢14¢ is the CP function call to perform virtual spool I/O. This is calculated by:

$$\frac{DIAG\_14}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_14 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 58/CMD. The number of DIAGNOSE code X¢58¢ instructions used per command. DIAGNOSE code X¢58¢ is the CP function call that enables a virtual machine to communicate with 3270 virtual consoles. This is calculated by:

$$\frac{DIAG\_58}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_58 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 98/CMD. The number of DIAGNOSE code X¢98¢ instructions used per command. This allows a specified virtual machine to lock and unlock virtual pages and to run its own channel program. This is calculated by:

$$\frac{DIAG\_98}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_98 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 98/CMD (V) VTAM Servers. See DIAG 98/CMD for a description of this instruction. This represents the sum of all DIAGNOSE code X¢98¢ instructions per command for all VTAM and VSCS servers. This is calculated by:

$$\frac{DIAG\_98\_VTAM + DIAG\_98\_VSCS}{\text{ETR (T)}}$$

*DIAG*_98_*VTAM* and *DIAG*_98_*VSCS* are taken from the VMPRF Virtual Machine Communication by User Class report for the VTAM and VSCS server classes respectively.

DIAG A4/CMD. The number of DIAGNOSE code X¢A4¢ instructions used per command. DIAGNOSE code X¢A4¢ is the CP function call that supports synchronous I/O to supported DASD. This is calculated by:

$$\frac{DIAG\_A4}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG_A*4 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG A8/CMD. The number of DIAGNOSE code X¢A8¢ instructions used per command. DIAGNOSE code X¢A8¢ is the CP function call that supports synchronous general I/O to fully supported devices. This is calculated by:

$$\frac{DIAG\_A8}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG_A*8 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 214/CMD. The number of DIAGNOSE code X¢214¢ instructions used per command. DIAGNOSE code X¢214¢ is used by the Pending Page Release function. This is calculated by:

$$\frac{DIAG\_214}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG_*214 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG/CMD. The total number of DIAGNOSE instructions used per command or job. This is calculated by:

For the FS7F, FS8F, and VSECICS workloads:

$$\frac{1}{(\text{ETR (T)} \times RTM\_Time)} \times \sum_{x \in \text{ DIAGNOSE}} TOTALCNT_x$$

For the PACE workload:

$$\frac{60}{(\text{ETR (H)} \times RTM\_Time)} \times \sum_{x \in \text{ DIAGNOSE}} TOTALCNT_x$$

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.

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).

$$\frac{1}{Num\_Entries} \times \sum_{t \in \text{ SCLOG entries}} Q0_t + Q0L_t + Q1_t + Q1L_t + Q2_t + Q2L_t + Q3_t + Q3L_t$$

$Q0_t$, $Q0L_t$ .. are from the Q0CT, Q0L ... columns in the RTM SCLOG screen. *Num_Entries* is the total number of entries in the RTM SCLOG screen.

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 (C). The total time, in seconds, required to execute the PACE batch workload.

This is calculated using the timestamps that appear on the console of the VSE/ESA guest virtual machine. The time the first job started is subtracted from the time the last job ended.

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).

$$\frac{1}{Num\_Entries} \times \sum_{t \in \text{ SCLOG entries}} E0_t + E0L_t + E1_t + E1L_t + E2_t + E2L_t + E3_t + E3L_t$$

$E0_t$, $E0L_t$ .. are from the E0CT, E0L ... columns in the RTM SCLOG screen. *Num_Entries* is the total number of entries in the RTM SCLOG screen.

EMUL ITR. Emulation Internal Throughput Rate. The average number of transactions completed per second of emulation time.

This is from the EM_ITR field under TOTALITR of the RTM TRANSACT screen.

EMUL/CMD. For the FS7F, FS8F, and VSECICS 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 FS7F, FS8F, and VSECICS workloads, this is calculated by:

$$10 \times \frac{\text{TOTAL EMUL}}{\text{ETR (T)}}$$

For the PACE workload, this is calculated by:

$$6000 \times \frac{\text{TOTAL EMUL}}{\text{ETR (H)}}$$

EMUL/CMD (H). See EMUL/CMD. This is the hardware based measurement.

For the FS7F, FS8F, and VSECICS workloads, this is calculated by:

$$10 \times \frac{\text{TOTAL EMUL (H)}}{\text{ETR (T)}}$$

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.

This is found in the NSEC column for ALL_TRANS for the total RTM interval time on the RTM Transaction screen.

ETR (C). See ETR. The external throughput rate for the VSE guest measurements.
For the PACE workloads, it is calculated by:

$$60 \times \frac{Jobs}{\text{Elapsed Time (C)}}$$

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.

For the VSECICS workload, it is calculated by:

$$\frac{1}{CICS\_Time} \times \sum_{t \in CICSPARSfiles} Total\_Transmits_t$$

Total_Transmits is from the TOTAL TASKS SELECTED line in the CICSPARS reports. *CICS_Time* is the run interval time, which is 900 seconds for all measurements.

ETR (T). See ETR. TPNS-based calculation of ETR. It is calculated by:

$$\sum_{t \in \text{TPNS machines}} \frac{Total\_Transmits_t}{TPNS\_Time_t}$$

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 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. This is calculated by:

For the FS7F, FS8F, and VSECICS workloads:

$$\frac{Fast\_Clear\_Sec}{\text{ETR (T)}}$$

For the PACE workload:

$$60 \times \frac{Fast\_Clear\_Sec}{\text{ETR (H)}}$$

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:

$$\sum_{f \in filepools} \frac{Total\_Filepool\_Requests_f}{\text{SFSTIME}_f}$$

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. This is calculated by:

For the FS7F, FS8F, and VSECICS workloads:

$$\frac{Free\_Total\_Sec}{\text{ETR (T)}}$$

For the PACE workload:

$$60 \times \frac{Free\_Total\_Sec}{\text{ETR (H)}}$$

Free_Total_Sec is taken from the NSEC column for the total RTM time interval on the RTM SYSTEM screen.

FREE UTIL. The proportion of the amount of available free storage actually used. This is calculated by:

$$\frac{Free\_Size}{\text{FREEPGS} \times 4096}$$

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).

This is found in the FPGS column for the total RTM time interval (<-..) on the RTM SYSTEM screen.

FST (File Status Table). CMS control block that contains information about a file belonging to a minidisk or SFS directory.

GB. Gigabytes. 1024 megabytes.

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:

$$60 \times \frac{\text{GUESTWT/SEC}}{\text{ETR (H)}}$$

GUESTWT/SEC. The number of entries into guest enabled wait state per second.

This field is taken from the NSEC column for the RTM total count since last reset, for the GUESTWT field in the RTM SYSTEM screen.

Hardware Instrumentation. See Processor Instrumentation

HT5. One of the CMS-intensive workloads used in the Large Systems Performance Reference (LSPR) to evaluate relative processor performance.

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:

$$\frac{1}{(1000 \times \text{ETR (T)})} \times \sum_{(f \in \textit{filepools})} \frac{\textit{Total\_BIO\_Request\_Time}_f}{\text{SFSTIME}_f}$$

Total_BIO_Request_Time is from the QUERY FILEPOOL STATUS command.

IO/CMD (Q). SFS file I/Os per command. This is calculated by:

$$\frac{1}{\text{ETR (T)}} \times \sum_{f \in \textit{filepools}} \frac{\textit{Total\_IO\_Requests}_f}{\text{SFSTIME}_f}$$

Total_IO_Requests is from the QUERY FILEPOOL STATUS command.

ISFC. Inter-System Facility for Communications

ITR. Internal Throughput Rate. This is the number of units of work accomplished per unit of processor busy time in an nonconstrained environment. For the FS7F, FS8F, and VSECICS 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 the FS7F, FS8F, and VSECICS workloads, this is found from the TOTALITR for SYS_ITR on the RTM TRANSACT screen.

For the PACE workload:

$$100 \times \frac{\text{ETR (H)}}{\text{UTIL/PROC}}$$

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 FS7F, FS8F, and VSECICS 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 FS7F, FS8F, and VSECICS workloads:

$$100 \times \frac{\text{ETR (T)}}{\text{TOTAL (H)}}$$

For the PACE workloads:

$$6000 \times \frac{\textit{Jobs}}{\text{Elapsed time (H)} \times \text{TOTAL (H)}}$$

Jobs is the number of jobs run in the workload. The values of *Jobs* are 7, 28, 42 and 56 for the PACE1, PACEX4, PACEX6 and PACEX8 workloads respectively.

ITRR. Internal Throughput Rate Ratio. This is the RTM based ITR normalized to a specific run. This is calculated by:

$$\frac{\text{ITR}}{\text{ITR}_1}$$

ITR_1 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:

$$\frac{\text{ITR (H)}}{\text{ITR (H)}_1}$$

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:

$$\sum_{f \in \text{ filepools}} \textit{LUW\_Rollbacks}_f$$

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. this is the same as

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 (<-..).

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.

This is taken from the %CPU column of the GUES-CPn line of the REPORT file for the master processor number as shown by RTM. In RTM, the first processor listed on the CPU screen is the master processor.

MASTER TOTAL. Total utilization of the master processor. For uniprocessor this is the same as TOTAL and is generally not shown.

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 (<-..).

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.

This is taken from the %CPU column of the SYST-CPn line of the REPORT file for the master processor number as shown by RTM. In RTM, the first processor listed on the CPU screen is the master processor.

MB. Megabytes. One megabyte is 1,048,576 bytes.

MDC AVOID. The number of DASD read I/Os per second that were avoided through the use of minidisk caching.

For VM releases prior to VM/ESA Release 2.2, this is taken from the NSEC column for the RTM MDC_IA field for the total RTM time interval on the RTM SYSTEM screen.

For VM/ESA Release 2.2, this is taken from the NSEC column for the RTM VIO_AVOID field for the total RTM time interval on the RTM MDCACHE screen.

MDC HIT RATIO. Minidisk Cache Hit Ratio. For VM releases prior to VM/ESA Release 2.2, the number of blocks found in the minidisk cache for DASD read operations divided by the total number of blocks read that are eligible for minidisk caching.

This is from the MDHR field for the total RTM time interval (<-..) on the RTM SYSTEM screen.

For VM/ESA Release 2.2, the number of I/Os avoided by minidisk caching divided by the total number of virtual DASD read requests (except for page, spool, and virtual disk in storage requests).

This is from the MDHR field for the total RTM time interval (<-..) on the RTM MDCACHE 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 measure only applies to VM releases prior to VM/ESA Release 2.2.

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 (blks). Minidisk Cache Reads. The number of times per second blocks were found in the cache as the result of a read operation. This measure only applies to VM releases prior to VM/ESA Release 2.2.

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 READS (I/Os). Minidisk Cache Reads. The total number of virtual read I/Os per second that read data from the minidisk cache. This measure does not apply to VM releases prior to VM/ESA Release 2.2.

This is taken from the NSEC column for the RTM MDC_READS field for the total RTM time interval on the RTM MDCACHE screen.

MDC REAL SIZE (MB). The size, in megabytes, of the minidisk cache in real storage. This measure does not apply to VM releases prior to VM/ESA Release 2.2.

This is the ST_PAGES count on the RTM MDCACHE screen, divided by 256.

MDC WRITES (blks). Minidisk Cache Writes. The number of CMS Blocks moved per second from main storage to expanded storage. This measure only applies to VM releases prior to VM/ESA Release 2.2.

This is taken from the NSEC column for the RTM MDC_PW field for the total RTM time interval on the RTM SYSTEM screen.

MDC WRITES (I/Os). Minidisk Cache Writes. The total number of virtual write I/Os per second that write data into the minidisk cache. This measure does not apply to VM releases prior to VM/ESA Release 2.2.

This is taken from the NSEC column for the RTM MDC_WRITS field for the total RTM time interval on the RTM MDCACHE screen.

MDC XSTOR SIZE (MB). The size, in megabytes, of the minidisk cache in expanded storage.

For VM releases prior to VM/ESA Release 2.2, this is MDNE for the total RTM time interval (<-..) on the RTM SYSTEM screen, divided by 256.

For VM/ESA Release 2.2, this is the XST_PAGES count on the RTM MDCACHE screen, divided by 256.

Millisecond. One one-thousandth of a second.

Minidisk Caching. Refers to a CP facility that uses a portion of storage as a read cache of DASD blocks. It is used to help eliminate I/O bottlenecks and improve system response time by reducing the number of DASD read I/Os. Prior to VM/ESA Release 2.2, the minidisk cache could only reside in expanded storage and only applied to 4KB-formatted CMS minidisks accessed via diagnose or *BLOCKIO interfaces. Minidisk caching was redesigned in VM/ESA Release 2.2 to remove these restrictions. With VM/ESA Release 2.2, the minidisk cache can reside in real and/or expanded storage and the minidisk can be in any format. In addition to the diagnose and *BLOCKIO interfaces, minidisk caching now also applies to DASD accesses that are done using SSCH, SIO, or SIOF.

Minidisk File Cache. A buffer used by CMS when a file is read or written to sequentially. When a file is read sequentially, CMS reads ahead as many blocks as will fit into the cache. When a file is written sequentially, completed blocks are accumulated until the cache is filled and then are written out together.

MPG. Multiple preferred guests is a facility on a processor that has the Processor Resource/Systems Manager* (PR/SM*) feature installed. This facility supports up to 6 preferred virtual machines. One can be V=R, the others are V=F.

ms. Millisecond.

Native. Refers to the case where an operating system is run directly on the hardware as opposed to being run as a guest on VM.

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.

NONPAGE RIO/CMD (V). The number of real SSCH and RSCH instructions issued per command for purposes other than paging. This is calculated by:

$$\text{RIO/CMD (V)} - \text{PAGE IO/CMD (V)}$$

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.

This is from TOTALTTM for the RTM NTRIV field on the RTM TRANSACT screen.

NUCLEUS SIZE (V). The resident CP nucleus size in kilobytes.

This is from the <K bytes> column on the Total Resident Nucleus line in the VMPRF System Configuration Report.

PAGE/CMD. The number of pages moved between real storage and DASD per command or job. This is calculated by:

For the FS7F, FS8F, and VSECICS workloads:

$$\frac{\text{READS/SEC + WRITES/SEC}}{\text{ETR (T))}}$$

For the PACE workload:

$$60 \times \frac{\text{READS/SEC + WRITES/SEC}}{\text{ETR (H)}}$$

PAGE IO RATE (V). The number of real SSCH or RSCH instructions issued on behalf of system paging.

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.

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 FS7F, FS8F, and VSECICS 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 FS7F, FS8F, and VSECICS workloads:

$$10 \times \frac{\text{TOTAL}}{\text{ETR (T)}}$$

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 FS7F, FS8F, and VSECICS workloads:

$$10 \times \frac{\text{TOTAL (H)}}{\text{ETR (T)}}$$

For the PACE workload:

$$6000 \times \frac{\text{TOTAL (H)}}{\text{ETR (H)}}$$

PD4. 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.

This is from the PPAG field for the total RTM time interval (<-) on the RTM SYSTEM screen.

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 the FS7F, FS8F, and VSECICS workloads:

$$\frac{1}{(\text{ETR (T)}) \times RTM\_Time} \times \sum_{x \in \text{privops}} TOTALCNT_x$$

For the PACE workload:

$$\frac{60}{(\text{ETR (H)} \times RTM\_Time} \times \sum_{x \in \text{privops}} TOTALCNT_x$$

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.

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.

This is from the NC field under CPU statistics on the RTM CPU screen.

PSU. Product Service Upgrade

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.

This is taken from the NSEC column for the PAGREAD field for the total RTM time interval on the RTM SYSTEM screen.

REAL STORAGE. The amount of real storage used for a particular measurement.

Relative Share. A relative share allocates to a virtual machine a portion of the total system resources minus those resources allocated to virtual machines with an ABSOLUTE share. A virtual machine with a RELATIVE share receives access to system resources that is proportional with respect to other virtual machines with RELATIVE shares.

RESERVE. See SET RESERVED

RIO/CMD (V). The number of real SSCH and RSCH instructions issued per command. This is calculated by:

For the FS7F, FS8F, and VSECICS workloads:

$$\frac{\text{RIO RATE (V)}}{\text{ETR (T)}}$$

For the PACE workload:

$$60 \times \frac{\text{RIO RATE (V)}}{\text{ETR (H)}}$$

RIO RATE (V). The number of real SSCH and RSCH instructions issued per second.

This is taken from the I/O Rate column for the overall average on the VMPRF System Performance Summary by Time report; the value reported does not include assisted I/Os.

Rollback Requests (delta). The total number of SFS rollback requests made during a measurement. This is calculated by:

$$\sum_{f \in \text{filepools}} Rollback\_Requests_f$$

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:

$$\sum_{f \in \text{filepools}} Rollbacks\_Due\_to\_Deadlock_f$$

Rollbacks_Due_to_Deadlock is from the QUERY FILEPOOL STATUS command.

RSU. Recommend Service Upgrade

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:

$$\frac{\displaystyle\sum_{f \in \text{filepools}} \frac{Sac\_Calls_f}{\text{SFSTIME}_f}}{\displaystyle\sum_{f \in \text{filepools}} \frac{Total\_Filepool\_Requests_f}{\text{SFSTIME}_f}}$$

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:

$$\frac{1}{\displaystyle\sum_{f \in \text{filepools}} \frac{Checkpoints\_Taken_f}{\text{SFSTIME}_f}}$$

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:

$$\frac{1}{\text{ETR (T)}} \times \sum_{f \in \text{filepools}} \frac{Filepool\_Request\_Service\_Time_f}{\text{SFSTIME}_f}$$

Filepool_Request_Service_Time is from the QUERY FILEPOOL STATUS command.

SHARE. The virtual machine's SHARE setting. The SET SHARE command and the SHARE directory statement allow control of the percentage of system resources a virtual machine receives. These resources include processors, real storage and paging I/O capability. A virtual machine receives its proportion of these resources according to its SHARE setting. See Relative and Absolute Share.

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.

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:

$$\frac{Percent\_Intercept \times \text{SIE/CMD}}{100}$$

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 the FS7F, FS8F, and VSECICS workloads:

$$\frac{SIE\_SEC}{\text{ETR (T)}}$$

For the PACE workload:

$$60 \times \frac{SIE\_SEC}{\text{ETR (H)}}$$

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.

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:

$$\frac{1}{(V\_Time \times \text{ETR (T)})} \times \sum_{s \in \text{server class}} Total\_CPU\_Secs_s$$

Total_CPU_Secs and *V_Time* are from the Resource Utilization by User Class section of the VMPRF reports.

TOT INT. Total Internal Response Time in seconds. Internal response time averaged over all trivial and non-trivial transactions.

This is the value for TOTALTTM for ALL_TRANS on the RTM TRANSACT screen.

TOT INT ADJ. Total internal response time (TOT INT) reported by RTM, 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.

This comes from the %CPU column for all processors for the total RTM interval time (<-..) on the RTM CPU screen.

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.

This comes from the %EM column for all processors for the total RTM interval time (<-..) on the RTM CPU screen.

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.

For 9121 and 9021 processors, this comes from the %CPU column for the GUES-ALL line of the REPORT file times the number of processors.

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.

This is the value of the <K bytes> column on the Trace Table line in the VMPRF System Configuration 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.

This is from TOTALTTM for the TRIV field on the RTM TRANSACT screen.

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.

UTIL/PROC. Per processor utilization. This is calculated by:

$$\frac{\text{TOTAL}}{\text{PROCESSORS}}$$

UTIL/PROC (H). Per 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-ALL line of the REPORT file.

VIO RATE. The total number of all virtual I/O requests per second for all users in the system.

This is from the ISEC field for the total RTM time interval (<-) on the RTM SYSTEM screen.

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 the FS7F, FS8F, and VSECICS workloads:

$$\frac{\text{VIO RATE}}{\text{ETR (T))}}$$

For the 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:

$$\frac{1}{(V\_Time \times \text{ETR (T)})} \times \sum_{s \in \text{ server class}} Virt\_CPU\_Secs_s$$

Virt_CPU_Secs and *V_Time* are from the Resource Utilization by User Class section of the VMPRF reports.

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.

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, ESA, and VMESA.

VSE WKSET (V). The average working set size in pages of the VSE guest machine in the VSECICS or PACE workload.

This is found in the WSS column in the VMPRF Resource Utilization by User Class report for the VSE user class.

VTAMs. The number of virtual machines running VTAM during a measurement interval.

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.

This is the average of the values for WSS in the VMPRF Resource Utilization by User report, (found in the Sum/Avg line).

WKSET (V) Server. Total working set of a related group of server virtual machine(s). This is calculated by:

$$\sum_{s \in \text{ server Logged Users}} Avg\_WSS_s$$

Avg_WSS is found in the Avg WSS column in the VMPRF Resource Utilization by User Class report for each class of server.

WRITES/SEC. The number of page writes per second done for system paging.

This is taken from the NSEC column for the PAWRIT field for the total RTM time interval on the RTM SYSTEM screen.

XSTOR IN/SEC. The number of pages per second read into main storage from expanded storage. This includes fastpath and non-fastpath pages. It is calculated by:

$$Fastpath\_In + NonFastpath\_In$$

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.

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.

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 the FS7F, FS8F, and VSECICS workloads:

$$\frac{\text{XSTOR IN/SEC + XSTOR OUT/SEC}}{\text{ETR (T)}}$$

For the PACE workload:

$$60 \times \frac{\text{XSTOR IN/SEC + XSTOR OUT/SEC}}{\text{ETR (H)}}$$

File Number: S390-34
Program Number: 5684-112

Printed in U.S.A.

GC24-5673-01